Introduction Round

SEMI Workshop

26-27 January, 2015, Zurich

Mary Barnes

MLB@Realtime Communications, LLC

„I think this is an extremely important topic (obviously).
We've being wrestling with this for way too long
(my initial work in IETF was around middle boxes in the long since retired MIDCOM WG).
Many existing protocols continue to be defined to run over yet another transport
to get around some of the problems.
That's a good thing in that it allows some applications to actually work in certain environments. I t's a bad thing because it's yet one more option to be negotiated/agreed to ensure interoperability.“

Current work area(s): IETF DISPATCH WG (chair), WebRTC (standards and product development), SIP Forum (board member and NNI, WebRTC Task group member), IAB programs: Stackevo, Privacy&Security, Emergency Services

Desired (key) outcome:
My primary desire is that we come out with some action items for work items.

David Black

EMC Corporation, tsvwg co-chair

„UDP encapsulation is crucial to the future of network transports,
so it needs to be straightforward to use in IETF protocol standards
(congestion and UDP zero checksum usage with IPv6 are known problem areas)“

Current work area(s): Storage, Network Virtualization, QoS

Desired (key) outcome:

- "UDP Encapsulation for Dummies" checklist creation (might belong in the 5405bis draft):
  This might include a list work that needs to be done, e.g., specify a reusable ECN-based "circuit breaker" (might be a rate-limiter in practice - analogy would be a "current limiter").
- Framework draft on network applicability:
  if one is creating a protocol that is only intended for the "well-behaved" subsets of the Internet, what are candidates for "well-behaved"?

Marc Blanchet

Viagenie

„Policies in access networks, specifically public access networks, heavily restrict the use of outgoing ports, where often the only real one without any processing is 443. This situation creates many issues and has a big impact on the engineering of the Internet. For example, application developers have to either create complex connections establishment algorithms to test ports, and always have to have the last resort being 443, or to just only use 443 to make connections faster. Therefore, while this trend is growing, this makes the Internet running on a single port, under a single transport protocol. This has major impacts on application design, transport protocols use, network policies, IETF protocol development, etc… In the IETF context, many application protocols, such as webrtc recently, are designed with defined ports and then later are redesigned with a complex connection establishment algorithm and negotiation. The paper proposes some possible actions regarding this issue.“

Current work area(s): Network Engineering, Applications

Desired (key) outcome:
IETF/IAB document explaining the issue and possible actions.

Bob Briscoe

BT

„L4 extensibility by tunnelling a control channel within the payload of vanilla TCP (the one wire protocol that does traverse all middleboxes) to switch in different transport behaviours with zero added latency (SCTP, transport parts of http2, tcpinc, etc); at the same time freeing up regular TCP option space for cooperative end-to-middle signalling.“

Current work area(s): Transport extensibility, Latency, tcpinc, Virtualisation of network functions (NFV, ie. virtualised middleboxes), Signalling end-2-middle

Desired (key) outcome:

A new IETF WG to work on how hosts talk with middleboxes, alongside a way to prevent all middleboxes (legacy and future) interfering when hosts talk with each other - in the timescale needed for tcpinc (ie. yesterday).

Ken Calvert

University of Kentucky, Lab for Advanced Networking

„Let us not waste this opportunity to introduce missing mechanisms into the architecture, so that future tussles play out in policy space instead of protocol space.“

Current work area(s): routing and forwarding, future Internet architecture, network security

Desired (key) outcome:
Efforts to define:

1. a framework + wire format for composing end-to-end features/functions, including security, in a way that admits future evolution and addresses current problems [I hope this will happen in taps]; and
2. mechanism(s) for determining path characteristics, supporting explicit end-middle-end information flow (and possibly/eventually negotiation).

Spencer Dawkins

Huawei, TSV AD

Desired (key) outcome:

Improve our understanding of ways that transport protocols can/must evolve.

Aaron (Yi) Ding

University of Helsinki / Columbia University

„Given the benefits of middleboxes,
the middlebox detection mechanisms are yet to improve and this requires
the efforts from protocol designers, end users, vendors, operators, and researchers,
in terms of measurement, tools, experience collection, deployable architecture design.“

Current work area(s): Traffic management, SDN for network management, Energy awareness on mobile systems

Desired (key) outcome:

A set of guidelines on how to detect and debug hidden issues of middleboxes and solve them in a cost-effective manner.

Benoit Donnet

Université de Liège (Belgium)

„We need a « bestiary » (or taxonomy) of middleboxes, based on their policies and unexpected complications, for allowing newly introduced protocols (or TCP extensions) to confront themselves to the actual network and anticipate connection issues due to middleboxes"

Current work area(s): network measurements, network modeling, Internet architecture

Desired (key) outcome:

1. developing/deploying middleboxes is nice but we have to keep in mind that we have to troubleshoot/debug network connections (in a large sens) due to complications issued by middleboxes; therefore, network measurements that is oriented to middleboxes is of the highest importance
2. quantifications/measurements of middleboxes complications in network.

Lars Eggert

NetApp & IRTF

„ Am afraid we're trying to re-boil an ocean we tried to boil before.“

Current work area(s): Datacenter networking, cloud networking

Desired (key) outcome:

A single, concrete, work item of small scope that has the correct (partial) deployability incentives. Because if we can't find one, we're done.

Gorry Fairhurst

University of Aberdeen

„Working over an increasingly heterogenous Internet may need a transport layer that can figure out path characteristics and is more than just a bunch of transport protocols.“

Current work: Mostly things transport, but including satellite transmission and a variety of Internet-related topics

Aaron Falk

Akamai Technologies

„My primary position is making sure the TAPS working group takes advantage of relevant insights from this workshop, however I think it is critical that we consider the full 'dynamic range' of the Internet and make sure the Internet doesn't break on paths with long-latency, highly multiplexed links with less than 1 pkt/RTT/flow, or between two resource limited hosts.“

Current work area(s): transport protocols, distributed systems design, Internet architecture

Desired (key) outcome:

1. Catalog what has changed that will allow us to make progress on these problems.
2. Get running code: Encourage experiments to demonstrate proposals --before-- chartering standards activity.

Mat Ford

Internet Society

„It would be cool if we could ‘solve the stack evolution problem once’ by promoting technological approaches and specific technologies that let innovation win the tussle between innovating end-to-end and adding value inside the network. Aligning incentives between operators and users will be key to the deployment potential of any middlebox signalling protocol.“

Current work area(s): Internet measurement, community building

Desired (key) outcome:

A well-structured report that helps to contextualise the problem, encapsulate the challenges and potential ways forward, to help with socialising to a broader audience. And next steps, with owners, to identify follow-up actions, e.g. measurement activities, standards requirements, guidance for application developers.

Ted Hardie

Google/IAB

„*Using NFV service chains as a replacement for middleboxes creates a challenge for path characterization and active queue management, as the scaling methods can change the number of queues or alter the network path.“*

Current work area(s): WebRTC, SDN, Privacy and Security

Desired (key) outcome:

Obviously there is more than one project possible that would impact network stack evolution--I'd like to identify the simplest thing we can do that would have a positive impact.

Joe Hildebrand

Cisco Systems

„Session protocol in UDP, with many new transports inside“

Current work area(s): Internet-scale architecture

Disired (key) outcome:
One or more WGs formed to work on concrete protocols.

Russ Housley

Felipe Huici

NEC Europe Ltd.

Current work area(s): Virtualization, high performance networking, user-level stacks

Desired (key) outcome:

Mostly going as first step/sanity check that we're not the only ones thinking along these lines.

Christian Huitema

Microsoft

„To restore the balance between end-to-end and middle-boxes, new transport protocols based on UDP shall encrypt transport control and transport payload.“

Current work area: Networking privacy

Jana Iyengar

Google

„"End-to-end"-ness cannot be best effort, it has to be enforced -- new and deployable transports are likely to be vertically integrated stacks that include authentication and encryption.“

Current work area(s): Transport, QUIC, TCP

Desired (key) outcome:

An articulated understanding of transport functions and how they are likely to evolve, so as to better allocate work/areas in the IESG.

Mirja Kühlewind

ETH Zurich, IETF rmcat co-chair

„To enhance transport extensibility we need
(i) a more flexible transport layer and potentially more functionality in UDP and
(ii) build-in middlebox detection and potentially a middlebox cooperation protocol.“

Current work area(s): transport protocols, congestion control, and ECN

Desired (key) outcome:
Charter for a middlebox cooperation working group

Eliot Lear

Cisco Systems

„It's important for transport development to occur,
and not stall due to inability to deploy new transports.“

Current work area(s): (mostly) applications & security

Desired (key) outcome:
Methods or general directions to break the log jam. Ideas that allow the compressing of layers, or redefining of layer boundaries in more flexible ways.

Barry Leiba

Huawei Technologies / Applications AD

„Application-layer protocols generally have to make do with what's available from TCP, or use UDP and roll their own transport, duplicating a lot of work; that needs to change.“

Current work area(s): Everything applications-layer

Desired (key) outcome:

A plan for making advanced transport-layer features reasonably available to application-layer protocols.

Xing Li

CERNET / Tsinghua University

„The Internet architectural should be improved in the transport layer to satisfy the high-performance, end-to-end and trust-to-trust requirements in a distributed fashion.“

Current work area(s): Internet/softwires/v6ops

Desired (key) outcome:

I would like to see the workshop produces a document with a better understanding of the R&D directions in this topic.

Szilveszter Nadas

Ericsson Research

„Wireless optimization benefits from communication between the network and endpoints, incentives for all actors are needed for this communication to happen.“

Current work area(s): Flow Metadata, Resource Sharing, Traffic Management

Desired (key) outcome:

Definition of light-weight functionality which allows communication with middleboxes.

Erik Nordmark

Arista Networks

„Evolving the protocol stack used on the Internet is difficult. Sometimes we end up with less than ideal solutions. But we should try to avoid re-inventing hexagonal wheels.“

Current work area(s): In my day job I work on Vxlan and similar encapsulations which use UDP for the purposes of entropy for ECMP. But my interest is much broader in terms of the evolution of the protocol stack that take into account evolution of the endpoints as well as middleboxes; the latter seems to be inevitable.

Desired (key) outcome:

A broad understanding of the issues and concerns that affect the evolution of the protocol stack. With that in hand we can explore what is common and different. Realizing that the solutions that get implemented and deployed might not be architecturally pure due to deployment considerations etc, we don't need to re-invent different such less-than-ideal wheels.

Colin Perkins

University of Glasgow

*„Ossification is an expected state in the development of the network;
future development must be mindful that some aspects of the network are frozen,
and may need to reinterpret the existing protocol stack in creative ways.“*

Current work area(s): networked multimedia transport

Desired (key) outcome:

- Work to extend the Internet architecture to make protocols running over UDP into first class citizens, to provide a new extension point for transport evolution.
- Clearer understanding of the scope of TCP semantics, and how that protocol can evolve.

Bernhard Plattner

Miroslav Ponec

Akamai Technologies

„It is entirely appropriate to design and deploy new protocols on top of UDP when TCP is inadequate to meet the needs of an application today.“

Current work area(s): Content Delivery Optimization (protocols - unicast, multicast, and peer-to-peer)

Desired (key) outcome:

Agreement on a plan how to help architects and engineers with approaching readily deployable protocol designs efficiently at scale.

Costin Raiciu

University Politehnica of Bucharest

„Endpoints need a predictable network,
and the current Internet is anything but predictable.
A predictable service can be achieved constructively (via new API)
but it may not get deployed despite its merits.
An alternative is to force a predictable service via aggressive and diverse tunnelling
- an approach we term ninja tunneling.“

Current work area(s): multipath transport, static analysis of middleboxes, protocol design

Desired (key) outcome:

A starting point for a more productive cooperation between middlebox deployers (ISPs), vendors and the rest of the suffering world.

Richard Barnes

Philipp Schmidt

TU-Berlin

„We need to extend the interfaces between the layers to make application intents and network capabilities visible to allow all layers to adapt where advisable.“

Current work area(s): Multi-Access Connectivity and Network Architecture

Desired (key) outcome:

A plan to incorporate a best effort based meta-information sharing like socket intents into the Socket API.

Martin Stiermerling

University of Applied Sciences Darmstadt / Transport AD

*„Skeptical about the success of any stack evolution
but curious to see what is going to happen.“*

Current work area(s): IETF Transport Area Director & teaching students

Desired (key) outcome:
A first idea about where we are in stack evolution.

Dave Thaler

Microsoft and IAB

„Protocols & apps should be not only IP-version agnostic
and crypto-algorithm agile,
but also lower-layer-transport agile,
capable of running over UDP and TCP and HTTP.“

Current work area(s): TCP/IP, APIs, IoT

Desired (key) outcome:

Agree on next steps given the divergent incentives of middlebox developers vs end system (host/app) developers.

Brian Trammell

ETH Zurich

„Engineering solutions to extending and evolving the transport layer need to be based on empirical observations of middleboxes deployed in the Internet; this measurement should happen out of band (as part of a path transparency observatory) as well as in-band (endpoints should sense the transparency of the paths they're using and adjust the protocol used accordingly). It would be nice if we could do this while cooperating with middleboxes that actually provide useful functionality.“

Desired (key) outcome:

A work plan for addressing this problem within the IETF/IRTF; I believe there's probably an IETF RG on middlebox cooperation hiding in here somewhere.

Michael Welzl

University of Oslo

„TAPS“

Current work area(s): Transport protocols

Desired (key) outcome:
I did think about it but I honestly don’t know.

Brandon Williams

Akamai Technologies

„Middle-boxes do provide value, and so direct collaboration between middle-box implementors and protocol designers is required in order to better support transport protocol evolution.“

Current work area(s): Cloud-based network, transport, and application optimization

Desired (key) outcome:

I hope that we will be able to identify some key existing or proposed drafts that could help to support in-band detection of path characteristics,
signing up a combination of middle-box implementors and protocol designers for collaborative work to move those drafts forward.

Dan Wing

Objectives

By Area

Architecture (3)

- Ideas that allow the compressing of layers, or redefining of layer boundaries in more flexible ways.
- Work to extend the Internet architecture to make protocols running over UDP into first class citizens, to provide a new extension point for transport evolution.
- A plan to help architects and engineers with approaching readily deployable protocol designs efficiently at scale.

Interface Broadening (4)

- A plan for making advanced transport-layer features reasonably available to application-layer protocols.
- A plan to incorporate a best effort based meta-information sharing like socket intents into the Socket API.
- "UDP Encapsulation for Dummies" checklist creation (might belong in the 5405bis draft): This might include a list work that needs to be done, e.g., specify a reusable ECN-based "circuit breaker" (might be a rate-limiter in practice - analogy would be a "current limiter").
- a framework + wire format for composing end-to-end features/functions, including security, in a way that admits future evolution and addresses current problems

Middlebox Cooperation (10)

- A starting point for a more productive cooperation between middlebox deployers (ISPs), vendors and the rest of the suffering world.
- Definition of light-weight functionality which allows communication with middleboxes.
- Key existing or proposed drafts that could help to support in-band detection of path characteristics, signing up a combination of middle-box implementors and protocol designers for collaborative work to move those drafts forward.
- I believe there's probably an IETF RG on middlebox cooperation hiding in here somewhere.
- A new IETF WG to work on how hosts talk with middleboxes, alongside a way to prevent all middleboxes (legacy and future) interfering when hosts talk with each other - in the timescale needed for tcpinc (ie. yesterday).
- Framework draft on network applicability:
  if one is creating a protocol that is only intended for the "well-behaved" subsets of the Internet, what are candidates for "well-behaved"?
- Mechanism(s) for determining path characteristics, supporting explicit end-middle-end information flow (and possibly/eventually negotiation).
- A set of guidelines on how to detect and debug hidden issues of middleboxes and solve them in a cost-effective manner.
- developing/deploying middleboxes is nice but we have to keep in mind that we have to troubleshoot/debug network connections (in a large sense) due to complications issued by middleboxes; therefore, network measurements that is oriented to middleboxes is of the highest importance; quantification/measurement of middleboxe complications in the Internet.
- Charter for a middlebox cooperation working group

Education (5)

- Catalog what has changed that will allow us to make progress on these problems.
- A document with a better understanding of the R&D directions in this topic.
- Clearer understanding of the scope of TCP semantics, and how that protocol can evolve.
- Improve our understanding of ways that transport protocols can/must evolve.
- An articulated understanding of transport functions and how they are likely to evolve, so as to better allocate work/areas in the IESG.

Outcome

1. WG/RG on **middlebox cooperation**/communication protocol/functionality
   - Mechanism(s)/guidelines for (in-band) detection of **path characteristics**
   - Measurements for middlebox **troubleshooting**/debug
2. Framework on **network applicability**
   - Better **understanding** on
     - how transport should/must evolve
     - functions, semantics, and applicability of TCP
   - **Interface** to make transport features available for apps
   - Framework for **composing** end-to-end services

Documents

- "UDP Encapsulation for Dummies"
- Definition of a **"well-behaved"** network from a transport standpoint.
- Middlebox detection and middlebox issue troubleshooting **guidelines**
- In-band path characteristic **detection mechanism(s)**.
- Definition of a lightweight endpoint-to-middlebox **protocol**.

Evangelism and Words of Warning

- "Get running code: Encourage experiments to demonstrate proposals before chartering standards activity." — Aaron Falk
- "Dispelling some of the myths that TCP is not suitable for application X because ..." — Stuart Cheshire
- "A single, concrete, work item of small scope that has the correct (partial) deployability incentives. Because if we can't find one, we're done." — Lars Eggert

Let’s start!

IAB Workshop on Stack Evolution in a Middlebox Internet (SEMI)

Stuart Cheshire

Apple

„The currently fashionable enthusiasm for new transport protocols layered on top of UDP is misguided; it is due, in large part, to widely held but incorrect assumptions both about (imagined) deficiencies in TCP and about the true requirements of application software.“

Current work area(s): International text, Service Discovery, Transport, Internet, Bufferbloat

Desired (key) outcome:

Dispelling some of the myths that, “TCP is not suitable for application X because ...”