Oh End-to-End, Where Art Thou?

Janardhan Iyengar, Google
jri@google.com

As the Internet has increased in scale and importance, middleboxes have proliferated in response to end-user and operator needs. Perhaps inadvertently, middleboxes have led to an ossification of traditional transport protocols, and the Internet's transport has now been ossified for well over a decade. Yet, new transport technologies are continually deployed and used[1]. This short position paper outlines this new world and the challenges it brings to IETF work.

A significant consequence of the Internet's ossification is that traditional transport protocols are no longer end-to-end. Middleboxes routinely inspect and modify TCP and UDP headers, and blackholing or modification of headers that do not conform to a middlebox's local policy or to expectations of protocol behavior is not uncommon[2]. As a result, traditional transport is restricted to TCP and UDP, and innovation inside these protocols must be conditioned on deployability -- any changes to TCP or UDP generally seek to not challenge middlebox expectations. MPTCP and Minion are examples of TCP modifications that seek to innovate while retaining backward compatibility with legacy middleboxes.

Contemporaneously, networked applications have deployed solutions to end-to-end problems, often as vertically integrated technologies that include transport functions, application functions, network traversal functions, and often end-to-end authentication and encryption as well. For instance, Skype handles reachability through various L4 blockages, Google Hangouts implements seamless migration across IPs and network interfaces within a session, and Adobe's RTMFP provides multistreaming within a single connection context (ala SCTP) among several other facilities.

Further, while latency has always been important[3], as user-facing applications increasingly move to an always-connected model, end-to-end latency has become a critical cost. Transport latencies impact application responsiveness, and new transport technologies also seek to reduce latency through the stack both in the protocols and in implementations. Such vertically-integrated stacks are currently successfully deployed on the Internet, and represent the market's response to the transport logjam.

This author's position then simply follows from the points above:

---

[1] Facetime, Skype, Quake, Adobe Flash, all implement their own solutions to transport problems inside the application.

[2] Honda et al. published some excellent initial work on measuring evolvabilty of TCP in IMC 2011, but much work remains to be done on detecting and measuring prevalence of various types of middleboxes.

[3] http://rescomp.stanford.edu/~cheshire/rants/Latency.html.

1. Middleboxes are part of the Internet architecture, and are here to stay, and are unlikely to change on a controllable deployment schedule. Consequently, traditional transports -- TCP and UDP -- have become *substrates* which are presumed to *not* be end-to-end, since middleboxes are, either passively or actively but universally, a part of these conversations.

2. As a corollary, middleboxes ought to be treated as first-class citizens in the network. New transport technologies must exist functionally above TCP and/or UDP, so as to be compatible with the extant Internet, and design of new transport technologies must consider delegation of trust and various end-to-end functions to middleboxes. Insofar as middleboxes do not fit the TCP/IP model of the Internet, a new model is needed to articulate a functional decomposition of the Internet's architecture. Our model of the Internet should be fitted to reality, not the other way around[4].

3. Any new transport technology that purports to be end-to-end must include strong end-to-end authentication and encryption to guarantee that the conversation is in fact end-to-end. As a result, new transport technologies are likely to be developed above TLS/DTLS/some equivalent isolation mechanism, and are likely to be vertically integrated along with application functions.

This position presents two broad and very different kinds of challenges to the IETF, which hopefully will be among the plethora of issues discussed at the workshop. First, working on TCP and UDP as substrates will require increased focus and energy on understanding the limits of TCP and UDP evolvability in the extant network. Middlebox behavior and TCP and UDP evolvability on the Internet are poorly understood problems, and to produce good designs that work on the Internet, relevant middlebox behaviors need to be better understood and articulated. Second, vertically integrated, or end-to-end, stacks pose challenges to the IETF's structuring of work into areas -- these solutions will likely span multiple areas. This first challenge is a structural challenge for the organization.

[4] See "Breaking Up The Transport Logjam" in HotNets 2008 for one such decomposition, which is still layered. Layering is not necessary to articulate placement of these functions however.