

June 2019
Geoff Huston

Network Protocols and Their Use

4 – Where Now?

In June I participated in a workshop, organized by the Internet Architecture Board, on the topic of protocol design and effect, looking at the differences between initial design expectations and deployment realities. These are my impressions of the discussions that took place at this workshop.

In this final part of my report I'll report on some expectations for the IETF's protocol standardization activities.

The internet faces many challenges these days, and while many of these challenges are the consequence of the Internet's initial wild and rapid success, few of these challenges have the same intrinsically optimistic tenor as compared to the challenges of the earlier Internet. We see an increasing capable and sophisticated set of threats coming from well-resourced adversaries. The increasing adopting of Internet-based services in all parts of our world increases the severity of these threats. We also see increasing consolidation by a shrinking set of very large global enterprises. Social media, search, cloud services and content are all offered by a handful of service operators and effective competition in this space is not merely an illusory veneer but has disappeared completely. The increasing dominance of many parts of the Internet by a small set of entrenched incumbents raises the obvious questions about centrality of control and influence, as well as the very real questions about the true nature of competitive pressure in markets that are already badly distorted.

For the IETF this poses some tough questions. Is the IETF there only to standardize those technology elements that these entrenched incumbents choose to pass over to an open standardization process to simply improve the economies and efficiency of their lines of supply while excluding some of their more important technology assets? If the IETF feels that this situation of increasing concentration and the formation of effective monopolies in many of these activity areas calls for some remedial action, then is it within the IETF's areas of capability or even within its chosen role to do anything here?

Some ten years ago the IAB published RFC 5218, on "What Makes a Successful Protocol". Much, if not all, of that document still holds today. The basic success factor for a protocol is for it to meet a real need. Other success factors include incremental deployment capability, open code, open specification and unrestricted access. Successful protocols have few impediments to adoption and address some previously unmet need. RFC 5218 also used a category called *wild success*:

"... a "successful" protocol is one that is used for its original purpose and at the originally intended scale. A "wildly successful" protocol far exceeds its original goals, in terms of purpose (being used in scenarios far beyond the initial design), in terms of scale (being deployed on a scale much greater than originally envisaged), or both. That is, it has overgrown its bounds and has ventured out "into the wild"." [RFC 5218]

One view is that for the IETF, success and wild success are both eminently desirable. The environment of technology standardization has elements of competitive pressure, and standards bodies want to

provide an effective platform for protocol standardization that encourages both submissions of work to be considered by the standardization process and through its standards imprimatur is able to label a technology a useful and useable. For the IETF to be useful at all it needs to be able to engender further wild success in the protocols it standardizes. So there is a certain tension in the propositions that the IETF should pursue a path that attempts to facilitate open and robust competition and eschew standardizing protocols that lead to further concentration in the market and the position that in order to maintain its value and relevance the IETF should seek to associate itself with successful protocols, irrespective of the market outcomes that may result.

Some of the tentative outcomes of this workshop for me have been:

- Technologies get deployed in surprising ways, which can have unintended consequences in threat models, surveillance capability and user privacy
- The focal point of technology and service evolution is moving up the stack, and applications are now taking responsibility for their own services, transport, security, naming context and similar.
- Perceived needs drive deployment, not virtue!
- Interoperability continues to be important but what are the interfaces that require standardization?
- With the Internet now the mainstream of communications, the support ecosystem is populated with more diverse actors and interests. IETF commentary could be helpful at this point, but by whom and to whom?
- Specific subject issues, such as DDOS, IOT, Spam, DNS, regulation, and centralization, are the topic of many challenging conversations, but none of these issues have easy resolution, and none are resolvable solely within the purview of the IETF.

What should the IETF do?

It is highly likely that the IETF will adopt a highly conservative position to such challenging questions and simply stick to doing what the IETF does best, namely, to standardize technologies within its areas of competence, and let others act as they see fit. The IETF does not define the Internet, nor is it responsible for either the current set of issues or the means of their solution, assuming that solutions might exist. The IETF is in no position to orchestrate any particular action across such a diversity and multiplicity of other actors here, and it would probably be folly for the IETF to dream otherwise.

No doubt the IETF will continue to act in a way that it sees as consistent with the interests of the user community of the Internet. No doubt it will continue to work on standardizing protocols and tools that proponents in the IETF believe will improve the user experience and at the same time attempt to safeguard personal privacy. It is difficult to see circumstances where the IETF would act in ways that are not consistent with such broad principles.

Disclaimer

The above views do not necessarily represent the views or positions of the Asia Pacific Network Information Centre.

Author

Geoff Huston B.Sc., M.Sc., is the Chief Scientist at APNIC, the Regional Internet Registry serving the Asia Pacific region.

www.potaroo.net