

June 2019
Geoff Huston

Network Protocols and Their Use

2 – Deployment Considerations

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 second part of my report I'll report on tensions between initial expectations of protocol design and standardization and subsequent deployment experience

Do we really understand the expectations of our Internet protocols?

What do we expect? Are these expectations part of a shared understanding, or are there a variety of unique and potentially clashing expectations? Do we ever look back and ask whether we built what we had thought we were going to build? Did anyone talk to the deployers and operators and competitors to understand their expectations, requirements and needs? In many working groups the loudest voices and the strongest held opinions might dominate a group's conversation, but this is not necessarily reflective of the broader position of interested parties, and not necessarily reflective of the path that represents the greatest common benefit. The strongest supporters of a single domain of interoperable connectivity are often new entrants and incumbents may have an entirely different perspective of the scope and expectations of a standardization effort.

Not only is this a consideration when embarking on standardization of a new protocol or a new tool element, but similar considerations apply to efforts to augment or change a standard protocol. Existing users may oppose the imposition of additional costs to their use of a protocol that appear to unfairly benefit new entrants. Change by its very nature will always find some level of opposition in such forums.

Perhaps one possible IETF action could be to avoid working on refinements and additions to deployed protocols, as this works against the interests of the deployed base and also sends a negative signal about the risks of early adoption of an IETF protocol. On the other hand, the IETF is not working in isolation, and the market itself would resist the adoption of protocol changes if those changes had no substantive bearing on the functionality, integrity or cost of the deployed service. In other words, if the augmentations offer no benefits to the installed base other than opening up the service realm to more competitors it is entirely reasonable to anticipate resistance towards such changes. A direction to the IETF to stop work on protocol refinements may well be a direction to stop working on ultimately futile efforts, and instead spend its available resources working in potentially more productive spaces, as the market will perform such choices between sticking to an existing protocol or adopting change in any case. But many items or work are started in the IETF with confident expectations of success, and "no" is a very difficult concept in an open collaborative environment. It does not need complete agreement, or even a rough consensus of the entire community to embark on activity. The more typical threshold is a cadre of enthusiasm. Whether its individuals or some corporate actors make no substantive difference in such circumstances.

This lack of critical ability to select a particular path of action and make choices between efforts has proved to be a liability at times. The standardization of numerous IPv6 transition mechanisms appeared to make a complex task far harder for many operators. The continuing efforts to tweak the IPv6 protocol appears to act against the interests of early adopters and a sense of delay and caution has become a widespread sentiment among network and service operators.

Scale has been a constant factor in deployment considerations. Protocols that can encompass increasing scope of deployment without imposing onerous costs of early adopters who are forced to keep up with the growth pressures being imposed by later entrants tend to fare better than those that impose growth costs on all. The explosive growth of Usenet news imposed escalating loads on all providers, and ultimately many dropped out. The broader issue of the scalability and cost of information flooding architectures cannot be ignored as an important lesson from this particular example.

Many protocols require adjustment to cope with growth. A good example here is the size of the Autonomous Number field in BGP. The original 16-bit field was running out and it was necessary to alter BGP to increase the size of this field. One option is a “flag day” where all BGP speakers shift to use a new version of the protocol. Given the scope of the Internet this has not been a realistic proposition for many years and probably many decades. The alternative is piecemeal adoption, where individual BGP speakers can choose to deploy a 32-bit ASN capable version and interoperate seamlessly with older BGP speakers. In general, where change is necessary for a deployed protocol, piecemeal deployments that are backward compatible with the existing user base will have far better prospects than those which are less flexible. In the early days of designing what was to become the IPv6 protocol there were various wish lists drawn up. “Backward compatibility” was certainly desired in this case, but no robust way of achieving this was found, and the protracted transition we are experiencing uses a somewhat different approach of coexistence, in the form of the dual stack Internet. Coexistence implies that any network cannot rid itself of a residual need for IPv4 services while any other network is still only operating an IPv4-only network. The entire transition process stalls on universal adoption, where the late adopters appear to claim some perverse form of advantage in the market through deferred cost of transition.

Is the IETF's conception of "need" and "requirement" distanced from the perspectives of operators and users? Should the IETF care when operators or users don't? Transport Layer Security (TLS) is a good illustration here. While the network was largely a wired network it was evident that users trusted network operators with their traffic, and efforts to encrypt traffic did not gain mainstream appeal. TLS only gained traction with the general adoption of WiFi, as the idea of eavesdropping on radio was easy to understand. And at this point the message of the need for end-to-end encryption had a more receptive audience. Should the IETF have waited until the need was obvious, or were its early actions useful in having a standard technology specification already available when user demand was exposed? It is hard to believe that the IETF has superior knowledge of the requirements of a market than those actors who either service that market or intend to invest in servicing that market. Having the IETF wait until it makes a clear judgement as to need runs the risk of only working on already deployed technology. At this point the value proposition of an open and interoperable standard is one that exists for all but the original developers and early adopters.

How do standards affect deployment? HTTPS is an end to end protocol that can be used to drive through various forms of firewalls and proxies. Packages that embed various services into of HTTPS sessions, including IP itself have existed for years, although the lack of applicable standards have meant that their use was limited to those who were willing and able to install custom applications on their platforms. The recent publication of RFC 8484 that described the technique of DNS over HTTPS (DOH) was more a case of formalizing an already well understood tunneling concept than representing some new invention. The existence of an IETF standard document effectively propelled this technology into a form of legitimacy, transforming it from just another tool in the hacker's toolbox into something that some mainstream browser vendors are intending to fold into their product. The standard in this case is seen as a precursor to widespread adoption. That should not imply that there is broad agreement about the appropriateness of the standardization or broad agreement with the prospect of broad deployment. DOH has been a story of emerging difference of expectations. Some browser vendors appear to be

enthusiastic about DOH as an enabler of faster service with greater control placed into the browser itself, lifting the name resolution function out from the platform and placing it into the application. However, the DNS community is not so clearly on board with this, seeing DOH as a potential threat to the independent integrity of the DNS as a distinct element of a common and consistent Internet infrastructure. Once the name resolution function is pushed deeply into the application what's to prevent applications from customizing their view of the name space? An important value of a single communications network resides within the concept of a single referential framework, where my reference to some network resource can be passed to you and still refer to the same resource. Should the IETF not work on technology standards that head down paths that could potentially lead to undermining the cohesive integrity of the common Internet namespace? Or are such deployment consequences well outside the responsibility of the IETF?

Deployment of technologies has exposed many tussles in the Internet. One of the major issues today is the tussle between applications and platforms. Today's browsers are now a significant locus of control, exercising independent decisions over transport, security, latency, and the name space, which collectively represent independent control over the entire user experience. Why should the IETF have an opinion one way or the other on such matters? If you take the view that a role of standards is to facilitate open competition between providers, then the issue in this space lies in the inexorable diminution of competition in the Internet. It appears that if one can realize unique economies of scale, and greater scale generates greater economies then the inevitable outcome is concentration in these markets. One of the essential roles of the IETF is diminished through this concentration within the deployment space, and the IETF runs the risk of being relegated to rubber-stamping technologies that have been developed by incumbents.

How can the IETF we measure the level of concentration in a market? If the IETF were to claim that they had an important role in supporting competition in decentralized markets, then how exactly would the IETF execute on this objective? What would it need to do? Is protocol design and standardization relevant or irrelevant to the industry composition of deployment that breeds centralization? Can the IETF ever design a protocol that would be impossible to leverage in a centralized manner? This resistance to concentration within the Internet appears to be an unlikely mission for the IETF. The Internet's business models leverage inputs and environments to create advantage to incumbent at-scale operators. It would be comforting to think that the protocols used, and their properties are largely orthogonal to this issue.

However, there is somewhat more at play. Standardization occurs during the formative stages of a technology, and this may be associated with deployment conditions that include early adopter advantages. If such advantages exist, then the rewards to such early adopters may be disproportionately large. This engenders positive market sentiment which motivates the early adopter to defend its unique position and discourage competition. Early adopters head to the IETF to shape emerging protocols and influence their intended entrance into the market. Their interests in the standardization process is not necessarily to generate a technology specification that facilitates opening up the technology to all forms of competitive use. Often their interests lie in the production of complex monolithic specifications replete with subtle interdependences and detail. Trying to position the IETF work to encourage competition by producing simple specifications of component elements that are readily accessible runs counter to the interests of early adopters and subsequent incumbents.

There is an entire world of economic thought on market dominance and competition, and it becomes relevant to this consideration about protocols and centrality in the Internet.

Is *big* necessarily *bad*? Is *centralization* necessarily *bad*? Or is the current environment missing some key components that would've controlled and regulated the dominant incumbents?

In many ways it seems that we are re-living the Gilded Age of more than a century past. There is a feeling of common unease that the Internet, once seen as a force for good in our society, has been entirely captured by a small clique who are behaving in manner consistent with a global cabal. The response to

such feelings of unease over the ruthless exploitation of personal profiles in the deployment space is to seek tools or levers that might reverse this situation. The tools may include law and regulation, the passage of time, new protocols, educating users, or new vectors of competition. In many ways this common search for a regulatory lever is largely ineffectual, as the most effective response to market dominance often is sourced from the dominant incumbent itself.

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