

Encouraging better deployments

Stéphane Bortzmeyer bortzmeyer@nic.fr

April 2019

The call for papers of the DEDR workshop rightly points three cases where the initial assumptions of some actors have been challenged in actual deployments. It then calls for studies to be done, to see why it happened, and what could be done when the deployment has serious flaws. I want to emphasize here that these deployment realities were not a simple consequence of the initial choices, and that most actions to do in order to change things are not in IETF core mission (which does not mean that IETF should do nothing).

Origins

It cannot be claimed that some deployment models of today (such as the three cases described in the call for papers) are direct consequences of choices made by the initial protocol designers. Although design decisions, as written in RFCs, do have consequences (for instance for the exercise of human rights, see RFC 8280, specially the case studies in section 5.2.3), there are many other forces that interact to push towards a specific deployment model, and the intrinsic characteristics of the protocol are only a part of it.

For instance, in recent discussions about DoH (DNS-over-HTTPS, RFC 8484), some people claimed things like “DoH leads to centralization” while nothing in the protocol could be read as promoting centralization. When Gmail and a few others are the dominant players for email, no one suggests that SMTP (RFC 5321) and IMF (RFC 5322) are at fault. We frame the problem of Gmail dominance as an email ecosystem problem, not a SMTP / IMF problem.

The protocols do influence the later evolutions (that’s why continuing human rights reviews of new protocols are very useful), they do not decide it completely. Other factors contribute, such as the strategies of the different players in the field.

The call for papers mention “presumed specific deployment models” as if all protocol designers had clear and conscious ideas about the deployment model of their brainchild. But I doubt it was often the case. So when the call for papers

say “initial system design assumptions turned out to be wrong”, I wonder if they were actually wrong or simply underspecified.

On a side note, the call for papers seems to imply that realities of the protocol use are always bad, when compared to original expectations. But it is not always the case. There are sometimes good unexpected results. I don’t think that the people who designed IP and TCP thought about the Web, and how it drove Internet deployments to very unexpected heights, or that people who designed the Web envisioned Wikipedia, and its incredible amount of freely available knowledge. As a general rule, users often “subvert” or “hack” the new services, and this is sometimes a good thing.

Causes

Then, what are the factors contributing to the actual deployments? In this paper, I won’t explore it further, the call for papers already mentioned several possibilities, I prefer to focus on IETF’s tasks. It is important to note that most of these factors are rooted in human decisions, typically inspired by business or operations, and are not purely technical. Whatever the IETF does, humans will adapt and will try to bend the Internet protocols to their own goals.

Also, as noted in the call for papers, the deployments changed sometimes a lot during the course of time. Email did not become “centralized” immediately, but after many years. The landscape of deployment is always changing.

Work

Then, what the IETF can do? As our main mission is to design and standardize protocols, we should obviously keep an eye on their consequences. Although they are certainly hard to predict, we should and must think about the consequences of what we design. For instance, RFC 6973 is a great document for checking privacy consequences of Internet protocols. This should continue.

Besides the existing (non-mandatory) privacy and human rights considerations sections, protocol designers could be encouraged to think about deployment, and to try to foresee factors encouraging or discouraging centralization. Documenting these factors is certainly useful, at least for future researchers who will work on the actual results. This effort needs to stay modest: prediction is hard and, in the end, humans (users, governments, corporations) will decide.

But, in my opinion, most of the work to avoid centralization, and to promote “good” deployment models, is outside IETF’s realm. It relies on the shoulders of software authors and people doing deployments. These two categories overlap (and also have a non-empty intersection with IETF) but they are not the IETF.

Decentralized deployments need the availability of software, since we cannot expect most people to write their own. It is therefore essential to encourage software developers to write high-quality, production-grade code, and to release it under a free software licence.

For instance, today there are very few implementations of DoH and most of them are experimental. If we don't want the DoH world to be dominated by a few big actors managing public resolvers, a pre-requisite is to have several good software implementations, ready to be deployed. Encouraging such developments is not core IETF mission, but it is a mandatory step to avoid centralization.

One point where the IETF may help them is to ensure that protocols are well-written and correct (our core competency) but also that they remain simple enough to help a vibrant and diverse software ecosystem to flourish. Outside of IETF (and sometimes inside), some protocols are so complicated that one may wonder if they were not written with the purpose of avoiding alternative implementations. . .

The second thing to do is to encourage people to deploy the software. Experience on the Internet show that a lot of people are ready to deploy Internet services, with various motivations, but there is still a need to support them.

Let's take a non-IETF example, ActivityPub. It is a federated social network, relying on "instances", local servers managed by individuals, non-profit organisations, corporations, state agencies, etc. Like any federated and decentralized service, it could become an oligarchy of a few big commercial actors. One of the reasons that prevents it, up to now, is that there are several good free software implementations, so a lot of people can deploy it, and do it in practice (there are thousands of active instances at this time).

So, to summarize:

- To decrease the risk of centralization, the IETF has some work related to its core mission but the important work is to be done by software authors and people doing deployments.
- we should do our best to encourage them and help them, and be among them.