

# Three kinds of concentration in open protocols

A position paper for “Design Expectations vs. Deployment Reality in Protocol Development Workshop 2019”

Andrew Sullivan, [ajs@anvilwalrusden.com](mailto:ajs@anvilwalrusden.com) or [sullivan@isoc.org](mailto:sullivan@isoc.org)<sup>1</sup>

There are at least two ways that one can observe concentration affecting deployment patterns on the Internet. One is the tendency to deliver functionality over the Internet as a web service or web application. This style of deployment means that the open protocol approach is countered by a bias toward APIs controlled by the web service operator. These APIs, which in another deployment model might be natural points for open protocols to develop, tend to favour large, closed systems deployed atop the theoretically open Internet. Importantly, however, these services depend (at least initially) on underlying open Internet infrastructure. There is an open question, in any case, about whether these kinds of services are natural monopolies<sup>2</sup>.

The second pattern is more curious, however, because it involves concentration of network services that offer the standard capabilities of the Internet. Prominent examples of this pattern include the provision of authoritative DNS, the provision of DNS resolution, and SMTP and associated mail services. Examination of these cases, however, suggests that there are three different stories to tell about them. One is that there is not a real problem, because the provision of services this way is a simple concentration of a commodity service. The demand in this case will find a natural equilibrium. The second is that the services provided amount to extensions of the functionality not available using the standard protocols on the Internet. This pattern may be troubling because of what it says about protocol standardization; but the issues could presumably be addressed with some care in future protocol designs. The third is that the services provided are concentrated because of business patterns related to the web-service operation mentioned above. This third case may be something of a threat to the Internet architecture.

## Commodity service, not a real problem

The least troubling of the cases involves the provision of standard services to customers on the Internet. It is true that this amounts to a kind of concentration, but it is not clear that this concentration is especially harmful. It represents the tendency of certain kinds of services to be provided more efficiently by a service provider than for everyone to do everything themselves. This is the same reason that, for instance, web hosting or even Internet service provision

---

<sup>1</sup> The author is an employee of the Internet Society. This paper does not represent an Internet Society position.

<sup>2</sup> See Andrew Sullivan (2017) Avoiding lamentation: to build a future Internet, *Journal of Cyber Policy*, 2:3, 323-337, DOI: [10.1080/23738871.2017.1400083](https://doi.org/10.1080/23738871.2017.1400083)

happens. It is of course *possible* for every person who wants a web page to operate their own web server (and indeed, that was a common model at the dawn of the web). It turns out often to be more convenient for people to hire another company to do most of the usual administration work of a set of web servers, and to use the benefit of open standards to enable a market in these standard services. One web server can be useful for many different web sites if they are all lightly loaded, and so it is better to amortize the cost of administration across several customers instead of every customer paying the cost of the administration themselves. Because the service is built up from standard components, concentration in this case is not likely to be a serious problem except insofar as it represents a concentration of intellectual capital (i.e. that nobody else ends up with the experience to operate the same protocols effectively). If the market gets too concentrated then new entrants are likely to show up: since this is a commodity service, switching costs are low, and there is likely always to be some demand for vendor diversity. So, the market will likely reach some equilibrium, but there is no obvious monopoly advantage in this market.

Understood this way, it is neither surprising nor particularly interesting or important than specialist companies offer mail or authoritative DNS hosting. These are Internet services that scale well, so it is natural there is a market that exploits efficiency gains. And the chances that a specialist organization can afford to hire specialist expertise seem higher than the idea that every or even most organizations can afford to do that independently.

#### [Special services fill standards gaps](#)

One difficulty that comes from the above market is that the services are all commodities. A commodity service can be trivially replaced. For someone consuming a network service, that is a benefit, but to the vendor of the service it is a problem. Naturally, then, specialist service vendors find ways of differentiating their services, diverging from standard features and thereby making them “stickier” for their customers.

Service vendors are sometimes aided in this by the gap between the facilities offered by standard protocols and the ways people wish to use those services. In the authoritative DNS market, for instance, the DNS protocol has long been maintained with an assumption that its answers come from a uniform global database that in theory will deliver the same answer to everyone. But operators found that it was useful to give different answers to the same query depending on where that query came from. The means to express “where it came from” took a very long time to standardize, and means to encode the alternative answers in a standard way have still not been settled.

A similar story can be told about “deliverability” measures and SMTP service. In theory, standard SMTP provides a mechanism whereby any mail server can send mail to any other (and indeed, this was historically a deployment pattern). As a practical matter, however, mail systems are aggressive about filtering and use the source of mail as one criterion for filtering. Arguably, the lack of strong antispam mechanisms within the mail standards produced the conditions under which some operators are more likely to be able to ensure their mail gets through than other operators do.

All of this suggests that a factor that needs to be considered is the gap between what standard protocols can or do provide (for whatever reason) and the kinds of functionality that people want. But it also suggests that some of these cases could be ameliorated with protocol development that is more flexible in operation and growth. (It may be significant that the examples I've used here are mostly old Internet protocols. But see below.)

### [Business patterns and linked protocols as an architectural threat?](#)

As noted above, a lot of deployment on the Internet recently tends to favour large, closed systems deployed atop the theoretically open Internet. These systems quite often rely on outside services, which ought to be a natural place to find common, open protocols. But often, instead, the services are offered over closed API that do not permit interoperation and that tend to make customer relationships “sticky” (see above for more on this). In addition, the widespread acceptance of services under high levels of central control has created a consumer market that may be less sensitive to concentration. At the same time, the service market (and the investments in it) appear to have a current bias in favour of mining user data. This encourages the development of services (and even protocols) where the service design tends toward concentrated deployments. This is particularly acute in services (and protocols) that depend on the Web for their functionality, because applications deployed as web services often have interlinkages that put control in central hands. That is, individual web services need not be centralized, but because of the way they can depend on and reinforce each other there is user data shared across the individual services which, when those services are controlled by the same back end provider, can reinforce advantages of the back end provider. The deployment model, then, reinforces a business model that profits from cross-linkages and consolidation.

This suggests that, if web applications do tend toward natural monopoly, it may well be that standard services that are to be delivered over HTTPS naturally cause a drift toward high degrees of concentration. That is, suppose a given web service wishes to depend on a standard protocol, and that protocol has two choices for deployment: it can use its own standard transport and so on, or it can be encapsulated inside HTTPS. The live question, then, is the extent to which the HTTPS encapsulation provides any of the concentration-promoting benefits of other web services, such as the wealth of user data available in HTTP headers, the potential for JavaScript use, openness to cross-origin use, and so on<sup>3</sup>.

If HTTPS encapsulation of standard protocols really is a factor in promoting concentration, a depressing reality might be that commercial incentives for such encapsulation will gradually increase. If true, there seems little that one can do about it at a technical level, because the issue actually has to do with economic consolidation. A standards community that refuses to accede to such encapsulation will find it has no influence to prevent the encapsulation, locking more Internet traffic inside web applications.

---

<sup>3</sup> Some concerns along these lines appear to underlie some current worries about DNS over HTTPS (DoH), though it is not clear to me how reasonable those concerns are given the actual specification of DoH. It may be that operational experience with DoH is too little so far to draw any conclusions about this, since the nature of web applications themselves did not seem obvious until high-concentration web applications were already a feature atop the Internet.