

Geoff Huston
June 2016

Open Season

In June 2016 the Organization for Economic Cooperation and Development (OECD) hosted a meeting of ministers to consider the state of the Digital Economy.

The central message from this meeting was the message that: "Governments must act faster to help people and firms to make greater use of the Internet and remove regulatory barriers to digital innovation or else risk missing out on the potentially huge economic and social benefits of the digital economy." All well and good, and as a piece of rhetoric it seems to strike an appropriately positive note without straying far from what appears to be bland truisms of our time.

The OECD also noted that: "the level of Internet openness will also affect the digital economy's potential. According to the Global Commission on Internet Governance (GCIG) "One Internet" report, presented at the Ministerial meeting, an open and accessible Internet should generate several trillions of dollars a year in economic benefits. A fragmented Internet on the other hand would weigh on investment, trade and GDP, as well as on the right to free expression and access to knowledge."

If the level of Internet openness is so crucial, then how well are we doing with the Open Internet? Two years ago the outlook was pretty bleak:

At best, I would label the "open" Internet an aspirational objective right now. We don't have one. It would be good if we had one, and perhaps, in time, we might get one. But the current prospects are not all that good, and talking about today's Internet as if it already has achieved all of these "open" aspirations is perhaps, of all of the choices before us, the worst thing we could do.

<http://www.potaroo.net/ispcol/2014-06/open.html>

Has the situation altered in any way since then?

Let's once more look at this topic again, and start with what we mean by the term "Open Internet." The expression builds upon a rich pedigree of the term "open" in various contexts. It gives the impression that "open" is some positive attribute, and when we use the expression of the "open Internet" it seems that we're lauding it in some way. But are we, and if so, in what way?

A useful characterisation of "open" is the Internet's use of free, publicly available standards that anyone can access and build on to. This represents a major shift from the world of closed vendor-specific technology standards of just a few decades ago. This was best exemplified in the time of mainframe computer technology. Back then if your enterprise was an "IBM shop" then exchanging data with a "DEC shop" was a major exercise in frustration. Everything that talked "IBM-ese" from the physical connectors through to the software protocols was different to what talked "DEC-ese". And woe betide the adventurous customer who even through of purchasing third party peripherals! Computer vendors were prone to making changes in the behaviour of their equipment to catch out these attempts to clone their equipment. That "closed shop" IT environment is well and truly over, and one to the

major levels that brought this to an end was the rise of computer networks. We now inhabit a network-centric world where it is the networks that define the common language of interoperability, and the attached computer systems are forced to adhere to these network standards. The result is that these days new providers can introduce goods and services that they can confidently expect to be fully compatible with their existing IT infrastructure and the larger Internet environment. Under this principle of open technology the Internet is not closed to new investments in the provision of good and services, and it minimises the barriers to entry. You don't need to reverse engineer the behaviours of an incumbent to access a market: you can simply adhere to a set of open standards, and from that derive confidence that your technology will interoperate with the installed base. In this respect “open,” as opposed to “closed, as in proprietary and private” is seen as an impressive step forward.

Another useful characterisation is that the Internet is "open" to all forms of traffic, and will treat all traffic that flows across the network in roughly the same way. This principle of the Open Internet is sometimes referred to as "Net Neutrality." Under this principle, consumers can make their own choices about what applications and services to use, and they are free to decide what lawful content they want to access, create, or share with others. This form of openness promotes competition and enables investment and innovation.

The Open Internet also makes it easy for anyone, anywhere to launch innovative applications and services, changing the way people communicate, participate, create, and do business. If you develop an innovative new tool that allows communication, you don't have to obtain permission to share it with the world.

We've learned some valuable lessons in the past few decades, and one is that technology is at its most effective when its accessible to open collaborative effort. When we are able to build upon the work of others, and contribute the outcomes back to the common pool of knowledge and capability we see the evolution of technology at its most capable and at its most astounding. At the heart of this is the concept of "openness."

But as usual there is a rather large gap between theory and practice, and while the theory of an open Internet appears to offer a bountiful set of opportunities, the practice is perhaps more prosaic. What we are learning is that this openness is extremely fragile, and what we have built with the Internet has exposed some very definite limitations that dampen this somewhat utopian ideal of openness in technology.

The underlying technology may be freely accessible, but that is not the same as freely usable. Free as in "freely available" is not necessarily the same as free as in "free beer" in this context. The Internet is built upon a towering structure of intellectual effort, but that does not mean that it free of various forms of Intellectual Property Rights. While a technology may be readily available, it may be that its use, particularly in commercial contexts, may entail license constraints, and yes, money changing hands. While it was fortuitous in the early days of the Internet that a royalty-free source code implementation of the internet protocol suite was released under the terms of a contract between DARPA and the University of California, Berkeley, this was perhaps one of a small number of exceptions rather than the general rule of the day. And these days the pervasive adoption of “open source” technology is not exactly the same as “free to borrow, adapt and re-package source”.

Nor is “open” the same as “unlimited”. The underlying network transmission and switching resources are not free to use and infinite in capacity, so when the aggregate sum of demand for access exceeds available capacity some form of arbitration occurs between these competing demands. "Free to use" a network to access content again does not imply that such use is free from cost, nor does it imply that such use with without constraint of any form. And “open to all” may be true to an extent, but when there is competition for access to a finite resource, accessibility and openness maybe arbitrated by a users' capacity to pay.

“Open” is not the same as “elastic” or “flexible”. Increasingly, we are shutting down the areas of potential innovation in communication models in order to concentrate our efforts to support a small number of service models. These days the means of communication on the Internet are limited to a conventional “client/server” transactional model, and it's limited to a behaviour that sits upon a web transaction, complemented by a name resolution service that uses the DNS protocol. Many other forms of interaction, such as peer-to-peer services, are often blocked by various forms of accreted network middleware. In this manner an openly available and accessible technology that uses an unconventional communications model may not necessarily prove to be a viable technology on today's environment.

And now that we have divided our environment into a distinguished set of servers and classified everything else as an undistinguished client, crossing the boundary and creating new servers and services is increasingly difficult. The barriers to entry are high and getting higher. With the hiatus in the supply of IPv4 addresses getting IP addresses for a service can be a challenge. With all the various forms of toxic attack, creating a service platform that is appropriately armored and defended can be a challenge. Increasingly services need to operate in a secure manner and eschew any form of open unencrypted operation. Again, this can be a challenge to set up and operate in a robust manner. The service and content delivery role is increasingly a specialized role performed by a small number of larger entities, and prospective service providers find themselves with few alternatives other than to use one of the service and content distribution experts.

Tensions are also apparent in the area of privacy and security. Should an open Internet support a user's choice to use tools, services and devices that preserve the user's personal privacy to the maximal capability of today's technology? Or should the considerations of security within a broader societal context place limits on the extent to which individual actions can be completely and totally concealed? Should we tolerate an increasingly exposed network that places all of us at increased risk of various forms of attack and exploitation?

At the same time an open network, open protocols, and open technology are all susceptible to various forms of malicious attack. How can we ensure that the millions of devices that people, businesses and public authorities use to connect to the Internet cannot be subverted and readily transformed into a catastrophic attack vector?

Perhaps the largest problem for open innovation in today's Internet is the Internet's own success and ubiquity. The incumbent providers can access economies of scale of operation that are inaccessible to all others, which allow them to gain positions of market dominance, and the stasis of the installed base means most forms of novel innovation fail to gain the threshold critical mass of acceptance needed to ensure a future. The larger the installed base the higher this threshold of critical mass of acceptance of innovation becomes.

None of these questions have clear answers. But they are pressing questions for public policy. Behind a dazzling veneer of high technology, the Internet is still just another form of the public communications space, and whether the Internet's various investment vehicles use private or public capital, the space in which we work and play on the Internet is always a public space. This means that while market forces strongly influence the day-to-day conversations about the Internet, the longer term debate needs the presence of a strong public voice to defend societal values. It is incumbent on us all to ensure that the open Internet continues to serve all of us, preserving essential qualities of ubiquity, accessibility, safety and utility that we should expect from every public common space.

The question is: Are we up to this challenge of preserving an open and vibrant Internet?

So far, the answers we've come up with are not looking all that promising.

Authors

Geoff Huston B.Sc., M.Sc., is the Chief Scientist at APNIC, the Regional Internet Registry serving the Asia Pacific region. He has been closely involved with the development of the Internet for many years, particularly within Australia, where he was responsible for building the Internet within the Australian academic and research sector in the early 1990's. He is author of a number of Internet-related books, and was a member of the Internet Architecture Board from 1999 until 2005, and served on the Board of Trustees of the Internet Society from 1992 until 2001 and chaired a number of IETF Working Groups. He has worked as an Internet researcher, as an ISP systems architect and a network operator at various times.

www.potaroo.net

Disclaimer

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