

June 2014
Geoff Huston

The Open Internet?

I'm sure we've all heard about "the open Internet." The expression builds upon a rich pedigree of term "open" in various contexts. For example, "open government" is the governing doctrine which holds that citizens have the right to access the documents and proceedings of the government to allow for effective public oversight, a concept that appears to be able to trace its antecedents back to the age of enlightenment in 17th century Europe. There is the concept of "open society," a theme that was developed in the mid 20th century by the Austrian philosopher Karl Popper. And of course in the area of technology there was the Open Systems Interconnection model of communications protocols that was prominent in the 1980's. And let's not forget "open source," which today is an extremely powerful force in technology innovation. So we seem to have this connotation that "open" is some positive attribute, and when we use the expression of the "open Internet" it seems that we are lauding it in some way. But in what way?

I hear the virtues of the "open Internet" being extolled so much these days that I can't help but wonder what exactly we are referring to. So let's ask the question.

What is an "open" Internet?

Is there a "closed" Internet out there somewhere, and we're trying to distinguish between the two? Or perhaps "open" means "good." In that case is there also a "bad" Internet, or even an "evil" Internet out there that we really should avoid. But how can we avoid these "closed," "bad" or "evil" Internets if we have no idea what an "open" Internet actually means in the first place. Is the term "open Internet" just one of those mantras that we constantly repeat to ourselves without really questioning what it actually means?

Maybe there is there some tangible attribute of "openness" that we are associating with the Internet when we use this term. In the same way as "open source software" implies some particular attributes about a software system when considering it's availability and terms of use of the software, does the term "open Internet" imply an Internet with quite particular attributes?

Let's ask Google.

My first effort to get Google to assist me in defining the open Internet led straight to a web site published by the Federal Communications Commission of the United States (<http://www.fcc.gov/openinternet>).

It's useful to quote from the opening paragraphs on this web page.

The "Open Internet" is the Internet as we know it. It's open because it uses free, publicly available standards that anyone can access and build to, and it treats all traffic that flows across the network in roughly the same way. The 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 are free to decide what lawful content they want to access,

create, or share with others. This openness promotes competition and enables investment and innovation.

The Open Internet also makes it possible for anyone, anywhere to easily launch innovative applications and services, revolutionizing the way people communicate, participate, create, and do business—think of email, blogs, voice and video conferencing, streaming video, and online shopping. Once you're online, you don't have to ask permission or pay tolls to broadband providers to reach others on the network. If you develop an innovative new website, you don't have to get permission to share it with the world.

<http://www.fcc.gov/openinternet>

But this does not seem to match the Internet that I am familiar with. Lets replay those two paragraphs and look into the detail, and see if they can withstand this level of scrutiny.

"The "Open Internet" is the Internet as we know it. It's open because it uses free, publicly available standards that anyone can access and build to"

Well not really. These days technical standards are subject to a considerable set of constraints relating to intellectual property rights. Many applications that operate over the Internet use proprietary protocols that are not described in free publicly available standards. Skype, for example. Or Apple's iMessage. And, of course there are many other examples where applications deliberately use closed proprietary protocols. Perhaps this sentence is referring more specifically to the technical standards published by the IETF, as the standards published by the ITU are not free, and many other forms of technical standards are distributed only to members or use other forms of qualifying access. What about IETF standards? Are they "open"? Can anyone use these standards? The answer is subtly nuanced, in that the standards are freely available, in that the technology specifications are published in such a way that they can be accessed without cost, but using them is another matter. The issue here is that IETF standards are not necessarily free of claims of Intellectual Property Rights (IPRs). And of course the rights holder may be making claims that effectively prohibit the use of this technology. The web site <https://datatracker.ietf.org/ipr/> contains the current list of disclosures of intellectual property rights relating to IETF standards. Its not an exclusive list, just those claims of intellectual property rights that have been disclosed in this public manner, and it is quite feasible that even when there is no published IPR claim on a technology, use of the technology could still generate IPR claims.

"and it treats all traffic that flows across the network in roughly the same way."

Well, only if your interpretation of "rough" is incredibly "rough!" The job of most firewalls is to treat traffic in widely differing ways. But perhaps that's a trite example. More seriously, we've seen internet service providers perform packet inspection and traffic profilers being deployed to identify traffic that was ostensibly associated with peer-to-peer networks. The rationale was that this particular technology was closely associated with movement of pirate content, but of course the outcomes was that traffic was not being treated uniformly. We've seen more recently claims that video streaming traffic is being deliberately damaged by access providers. The access providers have been complaining for many years that the content streamers were enjoying some form of "free ride" across their network. The content streamers took the view that the access network here in effect funded directly by the network's users, and the attempts to also extract money from the content providers was akin to "double dipping." The intensity of this debate has increased marked this year, and there are various allegations that access networks are indulging in selective treatment of traffic flows in order to place additional pressure on certain content providers.,

"The 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"

A solid case can be made that "net neutrality" is over. Certainly that appears to be the case in the United States at present, when in January 2014, the U.S. Court of Appeals sent the regulatory

framework of what is commonly referred to as “network neutrality” back to the US Federal Communications Commission (FCC), claiming that the Commission had overreached its authority in barring broadband network service providers from slowing or blocking selected content. In other words the court was saying that the framework that was intended to ensure that these carriers treated all content on an equal and non-discriminatory basis was beyond the authority of the FCC in this context.

What we are seeing is that the access providers are exercising their monopoly powers in the last mile access market to assert dominion over the users who reside in their catchment, and act as a broker of these users to content providers. Users have a decreasing ability to make their own choices about services.

"and are free to decide what lawful content they want to access, create, or share with others."

“Lawful content” or “censorship?” In some sense the distinction here is a matter of perspective, and one regime’s lawful content provisions are another’s effort’s to suppress criticism. But of course it extends well beyond governments and efforts to place boundaries on the activities of their citizens on the Internet. This also encompasses the efforts to preserve the monopoly interests of one sector over those of the users. If I purchase digital content, can I share it with others in the same way as I would lend you my book, or my CD? Or is this Internet a more restricted and paranoid environment, where sharing is frowned upon and labelled as theft? To what extent has the once-open Internet been turned into a shopping mall where the interests of the vendors prevail over any interests of consumers?

"The Open Internet also makes it possible for anyone, anywhere to easily launch innovative applications and services,"

Well, no.

The level of innovative freedom is incredibly limited. IP uses an 8 bit field to encode a transport protocol. So there are, in theory, 255 useable protocols. Only 1, TCP, is assumed to work most of the time, while 2 more, UDP and ICMP, could work some of the in some circumstances. The rest don’t work. In TCP not every port works. Sometimes some parts of the Internet are restricted to the web on port 80 and the secure socket service on port 443. A service on any other TCP port may, or may not, work. A similar picture exists with UDP, in so far as port 53 sometimes works, but its often intercepted and synthesised, as this is the DNS, and many of the efforts to enforce the provisions of lawful content involve filtering the DNS. And of course ICMP mostly does not work.

This means that you are free to innovate in the network to build innovative applications and services as long as you limit your innovation to TCP port 80 or port 443, and of course you must use the DNS as the rendezvous mechanism. It is surprising that we have the diversity in services and applications that we have given this incredibly set of underlying constraints. At the same time one can only conjecture as to what would have been possible if the underlying platform were as flexible as the Internet Protocol actually permits.

"revolutionizing the way people communicate, participate, create, and do business—think of email, blogs, voice and video conferencing, streaming video, and online shopping."

Think of a world where email is now an extensively analysed window into the user’s credit card. Think of a world where every device we use, and not just computers or this smart devices in our pockets, our cars, our credit cards, all report back to base about precisely where we are, when, and even what we are doing. Think of a world where the consumer is now the product. Where the considerable computation power of the world of information technology is used to assemble billions of individual profiles, taking each user’s purchases, monetary value, habits, desires, situation, health and any other exploitable aspect of their life and turning it into marketing collateral to be used to influence, cajole and lead you along

paths where your interests are subjugated to those of the retailer. Is this the world we wanted? Is this what we mean when we say “open”?

"Once you're online, you don't have to ask permission or pay tolls to broadband providers to reach others on the network."

I have yet to see a broadband provider who gives away their product for free. Yes, I pay my ISP to access the Internet in order to reach others. So do you. So does everyone else. Are they seriously suggesting that the Internet should be free? I haven't heard that particular catch cry since the very early 90's

"If you develop an innovative new website, you don't have to get permission to share it with the world."

Well, not any more. If you want to operate this website in a manner that is secure, then you will probably want a domain name certificate, and that certificate will need to refer to a unique IP address. But we've exhausted the supply of new IPv4 addresses, so the costs and effort required to secure your IP address has risen dramatically, and the longer this state persists, such costs will only rise further. So security of content is now a luxury good, rather than an open and readily accessible utility. So what if you resign yourself to the inevitable and use a virtual hosting provider? You then you need to worry about your virtual neighbours. Their indiscretions could land the common IP address on various national black lists, or on spam lists, and then your ability to share with the rest of the world is over.

So if that's an “open Internet” it has nothing to do with the Internet that you and I use today.

Can Wikipedia help here?

The idea of an open internet is the idea that the full resources of the internet and means to operate on it are easily accessible to all individuals and companies. This often includes ideas such as net neutrality, open standards, transparency, lack of internet censorship, and low barriers to entry. The concept of the open internet is sometimes expressed as an expectation of decentralized technological power, and is seen by some as closely related to open-source software.

http://en.wikipedia.org/wiki/Net_neutrality#Open_Internet

Let me hone in on one point here: “low barriers to entry.” Is this what we have in today's Internet?

Much has been undertaken in recent years in many countries the area of regulatory reform to increase the level of competition in the provision of goods and services in the communications sector. This effort entailed dismantling much of the regulatory paraphernalia associated with a single incumbent monopoly telco, and in its place encouraging the entrance of private sector investment and entrepreneurial innovation. There have been many successes here, and much of the the bloated inefficiencies of the incumbent telco have been removed through the onslaught of competition from new entrants, new ideas and new services.

But a number of problems still remain. The economics of last mile access has proved resilient to competitive pressure in many countries. We've seen the copper loop be the subject of competitive pressures, but twisted pairs of copper wire are not the best high frequency medium, so efforts to lift the speed of copper and at the same time allow for diverse access to each pair in a cable bundle is proving elusive. Speed and distance are also related, in an inverse manner, so in countries where the telephone operator saved costs by using extended copper loops are now paying a cost in very low achievable DLS speeds. Common cable infrastructure exists, but the nature of the cable, including the shared return path means that these systems support exclusive use by operators, and form natural monopolies. And fibre-based access systems have their own problems in cost, particularly where there is a requirement to retrofit new cable infrastructure in existing urban environments. Wireless was touted as the great hope

here, but again this is becoming a scarce resource, as the search for ever high data speeds results in ever increasing demand for spectrum, and each operator demands exclusive access to particular bands in the spectrum space, so there is a natural limit to the number of mobile service operators in terms of the underlying availability of spectrum. Are there low barriers to entry if you want to be a communications service provider in the Internet? No.

Although the topic of competition reform in the communications services sector absorbs a huge amount of attention by policy makers and public sector regulators. The studies on the evolution of the mobile sector, with the introduction of mobile virtual network operators, the concepts of spectrum sharing and the issues of legacy dedicated voice services and voice over data and similar topics dominant much of the discussion in this sector. Similarly there is much activity in the study of broadband access economics, and grappling with the issues of investment in fibre optic-based last mile infrastructure, including issues of public and private sector investment, access sharing, and retail overlays again absorbs much attention at present.

You'd think that with all this effort that we be able to continue the impetus of competitive pressure in this sector, and to continue to invite new entrants to invest both their capital and their new ideas. You would think that we would be able to generate yet further pressures on historically exorbitant margins in this sector and bring this business back to that of other public sector utility offerings. But you would be mistaken.

Open competitive markets depend on common platforms for these markets that support abundant access to many resources, and in terms of today's communications networks abundant access to protocol addresses are a fundamental requirement. But over the past decade, or longer, we have consistently ignored the warnings from the technology folk that the addresses were in short supply and exhaustion was a distinct possibility. We needed to change protocol platforms or we would encounter some serious distortions in the network.

We have ignored these warnings. The abundant feed of IP addresses across most of the Internet has already stopped, and we are running the network on empty. The efforts to transform the Internet to use a different protocol sputter. A small number of providers in a small number of countries continually demonstrate that the technical task is achievable, affordable, and effective. But overall the uptake of this new protocol continues to languish at levels that are less than 3% of the total user population.

The ramifications of this are, in general, completely invisible so far to the consumer. Web pages still appear to work, and we can all shop online from our mobile devices. But the entrance of new players, and the competitive pressure that place on the market is drying up. The lack of protocol addresses is an extremely fundamental barrier to entry. Only the incumbents remain.

Shutting down access to the Internet to all but existing incumbents should be sending chilling messages to regulators and public policy makers about the futility of generating competitive infrastructure in copper and in radio spectrum if the same cannot be maintained in the level of provision of Internet access and online services.

This is not a comfortable situation and continued inaction is its own decision. Sitting on our hands only exacerbates the issue and today's situation is assuming a momentum that seats incumbents firmly in control of the Internet's future. This is truly an horrendous outcome. Its not "open". Whatever "open" may mean, this is the polar opposite!

The Open Internet?

It should be pretty obvious by now that I'm sure that today we just don't have an "open" Internet.

The massive proliferation of network-based middleware has resulted in an internet that has few remaining open apertures. Most of the time the packet you send is not precisely the packet I receive, and all too often if you deviate from a very narrowly set of technical constraints within this packet, then the packet you send is the packet I will never receive. The shortage of addresses has meant that the rigors of scarcity has replaced the largesse of abundance and with this has come the elevation of what used to be thought of as basic utility, including privacy and security in online services into the category of luxury goods only accessible at a considerable price premium. Our technology base is being warped and distorted to cope with an inadequate supply of addresses and the ramifications extend out from the basic domain of the internet protocol upwards into the area of online services and their provisioning. From the crowding out of open technology by encroaching IPR claims, to the problems of the mass of our legacy base restricting where and how we can innovate and change, and the rigors of scarcity of addresses, the picture of the technology of the Internet is now far from “open.”

Maybe the “open” Internet is something entirely different. Maybe it’s about the policy environment, and the competitive landscape. Maybe it’s about the attributes of having no barriers to entry in the supply of goods and service using the Internet. This could be deregulation of the carriage and/or access regime, allowing competitive in packet transport. Or the ability to deliver content and services without requiring the incumbents’ permission and without extortionate price gouging on the part of providers of critical bottleneck resources. Maybe in the “open” Internet we are talking about the benefits from low barriers to entry, innovation, entrepreneurialism and competition in the provision of goods and services over the Internet platform.

But this is not "open" either. The fact that we’ve exhausted our stock of IP addresses impinges on this considerations of markets for the provision of goods and services on the Internet, and their open operation. Without your own pool of IPv4 addresses you cannot set up a packet pushing business, so that's no longer "open". And without your own pool of IPv4 addresses you cannot set up secure services as a content service provider. So as long as you are willing to offer goods and services over an open, insecure, untreatable channel, and as long as you are willing to put the fate of your enterprise in the hands of your virtual neighbours with whom you are sharing IP addresses and hosting platforms, and so long as the price of access to these shared address resources is not in itself a barrier to entry, then perhaps this niche is still accessible. But its not what we intended it to be. It’s not “open”.

Perhaps the “open” Internet, in the sense of being an "open" platform that can carry the hopes and aspirations of a socially transformative power of a post-industrial digital economy, is now fading into an ephemeral afterglow.

Maybe its not too late, and maybe we can salvage this. But there are many moving parts here, and they all deserve attention. We need to use an open common technology platform that offers abundant pools of protocol addresses. Yes, I’m referring to IPv6. But that’s by no means the end of this list. We need continuing access to software tools and techniques. We need open software. We need open technology. We need open devices and open access to security. We need to open competitive access to the access infrastructure of wires and access to the radio spectrum. We need open markets that do not place any private or public enterprise in overarching positions of market dominance. We need an open governance structure that does not place any single nation state in a uniquely privileged position. We need open dialogues that enfranchise stakeholders to directly participate in conversations that matter to them. Indirect representation is just not good enough. We need all of these inputs and more. And each of them are critical, in as much as we are aware from centuries of experience that failure in any of these individual aspects translates to catastrophic failure of the entire effort.

Yes, this is asking a lot from all of us. But, in particular, its asking a lot from our policy makers and regulators. The mantra that deregulated markets will naturally lead to these forms of beneficial outcomes that enrich the public good ignores a rich history of market distortions, manipulations and outright failures. An “open” Internet is not a policy free zone where market inputs are the sole constraint. Markets aggregate, monopolies form, and incumbents naturally want to set forth constraints and conditions that define the terms of any form of future competition. And in this space of market

behaviours our only residual control point lies in the judicious use of considered regulatory frameworks that encourages beneficial public good outcomes.

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.

Today’s Internet is many things, but it’s certainly not an “open” Internet. It could be, but to get there it's not just going to happen by itself. It’s going to need our help.

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. He has been closely involved with the development of the Internet for many years, particularly within Australia, where he was responsible for the initial build of the Internet within the Australian academic and research sector. 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.

www.potaroo.net