Up and Down the Stack

Through a Nerd's Eyes: Making The Internet Better the Internet Way

February 24, 2019

Hi everyone. I'm super pleased to be invited here to speak to you today. I've been trying to make it to an APRICOT meeting for many years, and I finally did it! Thank you so much for welcoming me, and to Philip and the Program Committee for inviting me, and thanks to our hosts and sponsors for making this excellent meeting possible. I must say that my ability to get here is very much to the credit of the Internet Society team here in the region: Raj, Noelle, Subhashish, Naveed, Olivia, Aftab, and Adrian. Adrian has just joined us, but the rest of the team made sure this was on my calendar months ago, so thanks to them.

Who am I, you might ask, that anyone should invite

me here? I'm Andrew, and I'm a nerd. For those of you who don't know me, during most of my career I worked on technical stuff. I was a database guy who helped set up the .info domain name registry in 2001. I moved over to the DNS and starting to work on that. At the Internet Engineering Task Force I was co-chair of some working groups. I was one of the primary offenders behind RFC 6141 (DNS64), and I was an Internet Architecture Board member and chair. From 2012 until 2018 I worked for Dyn, who provide a significant amount of DNS infrastructure.

While I was IAB chair, I had to be an Internet politician. At that time, we were going through the IANA stewardship transition, and I had to talk to people inter-

ested in policy, and the general public, and even (gulp!) politicians. Today, I am the President and CEO of the Internet Society, which is an organization many of you know and work with. We've been around for more than 25 years, trying to make sure that the Internet is for everyone. We support and promote the development of the Internet as a global technical infrastructure, a resource to enrich people's lives, and a force for good in society.

Today I want to talk about what I have learned about networks over the years, by travelling up and down the stack – all the way to layer 9! There's a reason I want to talk about this. It is because I am alarmed about what could happen to the Internet. We are at a meeting of

Internet operators. Internet operators understand, because they have to, what the Internet model of network deployment is all about, and why it works the way it does. But now that the Internet has become so widespread and so important for others, those others want to manage it. And it turns out, many of the people who have ideas about managing the Internet misunderstand how the Internet works, and so they're proposing policies that are harmful. So, I'm not here to offer you advice, or to tell somewhat-amusing anecdotes about what it's like for a nerd to talk to a Prime Minister. I'm here, instead, to try to enlist your help. If we're going to keep the benefits of the Internet we have, then we're all going to need to pitch

in to make sure that the Internet is not undermined.

You may be asking why should you care about things like "infrastructure" or "tech policy"? Surely it's someone else's problem – maybe the Internet Society's. The policies, the lofty theories, and the lawyers come and go. Netops remains, right?

Well, I'll tell you why. Nobody Internets alone. Or, to state that in proper English, you can only build the Internet with all the other networks: together.

The magic of the Internet really comes from its architecture. It is, as you all know, a network of networks (of other networks, all interconnected). But it works because of what seems like a magic trick. Every network that

participates does its own thing: each one implements the things that it wants to implement to support local needs. Yet, to get the big advantages, participating networks use common protocols, which work from the edge of one network all the way through to the edge of another network, to permit communication. Even better, that communication works without everyone having to have contractual relationships with one another all the way along. So, you *can't* Internet alone. The way you get any kind of internet is people building networks using common building blocks that permit open communication, and the way we all get the global Internet is for us to use those common protocols everywhere. In this sense, things like "open protocols"

and "open networks" are not a *value*, as though there is some kind of moral or political claim. They're instead a necessary condition for having the Internet at all.

This way of building a large, global network has big advantages that are, I would argue part of the reason the Internet has ended up displacing most other networking technologies. For instance, the Internet does not require a lot of centralized or global co-ordination to make it work, because the only real prerequisites are that you have an end point that can speak the relevant protocol, and a way to send traffic. This makes the Internet cheaper to operate than other large networks.

The Internet style of engineering is also more respon-

sive to change than other styles, because the costs of changes are aligned with the local operator of the network who is making the changes (and who is presumably getting the benefits of change).

And the Internet is a wonderful promoter of opportunity and innovation, because this style of networking permits you to invent something and share it with others – maybe just your friends – without getting any permission to do so.

But this way of building a large, global network of networks also depends on people sharing the assumptions behind it and working within those assumptions. Those who participate need to believe, at bottom, that the co-

operative model will work because each of us has enough self-interest to keep it working. Your interest in our communication means you'll route my packets, and my interest means that I'll route yours.

Now we all know that network operators – yes, even some of us here! – deploy middleboxes that are designed precisely to foil the end-to-end design. That's ok, because the end-to-end network is mostly a spectrum of behaviour. It's perfectly normal that not everything in every network can be contacted from every other node on the Internet. All we need is for all the different networks to be interconnected in a more or less open way.

I don't need to tell you that the Internet is not the only

way to build networks. There is nothing inherently wrong with centralized networks. The telephone network was a beautiful piece of engineering, even though it was not an internet. It's not that other networking techniques are bad. They're just not as useful as the Internet's technique, because they are more expensive or less flexible or both. Still, the value to people in the Internet style of engineering is getting lost. Some people are trying to take the success of the Internet, and turn that to their own purposes in promoting other network services that are not like the Internet at all. Various countries, of course, are busily trying to turn themselves into modern versions of the old information services like America Online and Compuserve.

These systems provide a *gateway* to the Internet (or parts of it), but they're not actually alternatives to the Internet.

Depending on the deployment model, 5G may also be a technology that undermines the ability of the network edge to control its destiny. Slices might turn out to be a great way to use the available bandwidth more efficiently. But they also look like an excellent opportunity for carriers to impose greater controls on the end points.

And, of course, many web applications are elaborate walled gardens, rather than ways to reach the Internet. The business model is dependent on keeping you inside the garden.

Every single one of these approaches to global network-

ing depends, today, on the Internet. Country-firewalled networks depend on the Internet to provide what they need for their desired network and for their users. The entire promise of 5G is dependent on ubiquitous connectivity, demanded by businesses that have thrived on the Internet. It remains to be seen whether those business models will still work if, as some critics have warned, 5G is used to put all the control back in the hands of the carriers. And the large, enormously profitable, near-monopoly providers of "walled garden" applications often help to build Internet infrastructure. They certainly depend on the Internet to deliver their users to them.

The current environment comes with some challenges,

and we will need to face those challenges. But the right way to face the challenges is not to adopt rules and strategies that will replace the Internet with some other system. There is only one architecture recognizable as "the Internet", and it is an open network of networks. Everything else is just some other kind of network. We have to design responses to problems that take into account the real nature of the Internet. And to do that, we who understand the Internet need to ensure that those who make the policies understand what is and is not possible. We need to show them that the Internet way is the best way to achieve their goals.

Clouds, edges, and the way we network

Let's look at an example to see what I'm talking about. Many of you will know that my former employer, Dyn, had a Bad Day in 2016. That event was in some ways a result of the kind of open architecture I am arguing for, so there are a couple things I want to highlight to help us understand the way forward.

Who's in charge around here?

The first issue seems easy to understand. Because of the open architecture of the Internet, a number of badly-

designed devices were connected to it. The devices were security cameras – Dyn said there were many of them in the Mirai botnet that was the source of much of the problem. Security cameras are almost perfectly designed to become sources of attacks. They need good bandwidth, so you can watch what's going on. They need lots of processing power to compress video, so a few additional processes doing nasty things won't be noticeable. And they need to be easy to turn on and hook up, so the chances are excellent that their default security profiles are a total shambles. And so, the video cameras were made into a botnet, and the rest is history.

In the rest of the world, where people walk around and

interact with physical objects, we have *lots and lots* of regulated things that cannot be sold without the necessary regulatory stamp. Want a car? There is an enormous pile of regulations that were worked out over many years in order to make that car acceptable on today's roads. Want a wire to install in your house? There's a whole *other* pile of regulations for that wire. How about a light bulb? Well, yes, there are regulations for that too, and national standards bodies that set what qualifies, and so on. Food. Clothing. Furniture in some countries has an amazing label that tells you it is illegal to remove the label!

It's no surprise, then, to see people – even quite respected technical people – calling for government regu-

lations of devices to connect to the Internet. But how could such regulations actually work? Governments are, necessarily, geographically limited. The Internet connects networks to other networks, not countries to other countries. There is no reason a network needs to end at a country border. Indeed, connecting across different national frontiers is part of what makes the Internet stronger and more resilient in the face of trouble, so giving up that connectivity pattern is not a good idea.

Well, the idea is that when a jurisdiction that happens to be an important market enacts good regulations, it will affect everyone positively because everyone will get the benefit of the improvements. No sane device maker will

have a "US-only" or "Korea-only" device. Instead, they'll build one device that meets the most stringent rules, and sell that everywhere.

The problem with this idea ought to be obvious: it assumes that no jurisdiction will enact regulations that conflict with another jurisdiction's. But conflict among different jurisdictions' regulations is quite common. Political processes naturally take into account the interests of the politicians' constituents, and not the interests of people in another country. The way that governments co-ordinate that kind of international inconsistency is through treaties. I sure hope our solution to the problem, "Security on the Internet is not happening fast enough," is not, "I know!

Let's get an international treaty!" Treaties are magnificent things, but they are not usually adopted quickly by every country in the world, and they're not always adopted the same way.

But there *are* things we can do, and they start with returning to the basic design of the Internet. The Internet means that, on your network, you are in charge. And that means that you need to think about what kinds of things are connected to the network, and how. Remember that the Internet style of networking puts most of the "intelligence" about decisions at the edge of the network. We do that because the applications at the end are really the things that are in the best position to know what

they need. Now, one of the important problems with the Internet of Things devices is that they're often not very "smart". They have limited capabilities and usually primitive interfaces. This means that the *application*, which is where the "intelligence" is supposed to live, is not the same as the *device*. Consumers need to be able to select devices that are safe, yes; but they also need network capabilities that enable that safety, and applications that provide the necessary safety. Something like the Manufacturer Usage Description (or MUD) is needed, to ensure that network traffic that is intended to be *only* local stays local. That approach does not require international treaties or new regulatory authority, and it certainly

doesn't require remaking the engineering of the Internet. Instead, it relies on network operators doing things to make their own networks less vulnerable, which produces a virtuous circle that makes the Internet work well.

This is why the Internet Society keeps working on Internet of Things efforts with consumer groups and through multistakeholder approaches. But it's also why we need network operators to remain engaged in these topics. The voice of reality about how the networks actually behave is important to counter the views of those who would prefer a different architecture. The low layers of the Internet community need to send that message to the high (political and economic) layers.

In addition, there are things *you*, as a network operator or consumer, can do. If you deploy access networks, consider how consumer-oriented gateways could use MUD as part of a way to keep traffic that ought to remain local (or otherwise constrained) appropriately controlled. (As an aside, while it might seem like there is some irony in proposing middleboxes to protect the end to end network, we have to remember that the end point is really an *application*.) If you are purchasing IoT devices, either for yourself or as part of your duties, make sure they meet at least minimal standards such as being updatable and maintained, and configurable so as to avoid dangerous defaults. Private trustmarks for devices and systems are

starting to roll out, and it makes sense to track them and try to depend on them where they provide the information you need. And of course, if you are a device manufacturer, we urge you to adopt the principles expressed in the Online Trust Alliance IoT Trust Framework. Distributed action by many players will make us more secure than a single, regulator-imposed solution. Finally, of course, help us educate and inform those who might be making policy without the benefits of understanding the Internet. I will naturally suggest co-operation through your local Internet Society chapter, but there are lots of ways to help with this!

Does the cloud mean the old architecture is dead anyway?

The second issue in the Dyn attack is really a difference between the way we often talk about the Internet, and how it is actually deployed today. It is useful to reflect on this because it is likely to affect what options we have in the future.

When I spoke earlier about the Internet and all the independent networks that make it up, you no doubt thought about the differences among those networks. As we all know, not all Autonomous Systems are the same. Some are tiny. Some route a lot of traffic. Some receive practically no traffic and send lots. And some, of course,

provide services for others.

In the early days of the Internet, people ran things themselves. Even if I'd had a giant DDoS against my DNS server, it wouldn't have been news because I wasn't going to affect everyone else. The reason attacks make the news now is mostly not because one really important site goes down, but because so much of the Internet infrastructure is now provided by a small number of operators. We all use the cloud; so when a cloud operator has a bad day, lots of people are affected.

Like many trends, cloud services have gone through fashion cycles. Just a few years ago, you could take approximately anything, say "cloud" in front of it, and

call that a business plan. The cloud was subject to insane quantities of hype, and everyone needed to have a "cloud strategy".

It has long seemed to me that as soon as popular publications start talking about some trend in computing and putting the word "strategy" after it, that trend will soon be in trouble. Sure enough, people are now criticizing cloud services as being too centralized, too profitable, too big, or too powerful. The cloud is, in fact, part of a general trend of consolidation and concentration. Whether it be the consolidation of transit into fewer, larger players; or the increasing power of web applications; or even the smaller, more narrow collection of protocols that we rely

upon (overwhelmingly, https), there's a lot of evidence of consolidation and concentration on the Internet.

Still, it is important that we recognize that this is all not really as new as it seems. It is true, of course, that Amazon Web Services (or Alibaba or whoever you like) has altered the deployment of many services online. And indeed, the basic design of online services – now mostly web services and APIs instead of open protocols – does change what the application layer of the Internet looks like. There are good reasons to be worried about this.

But various kinds of concentrations have emerged before on the Internet, at multiple layers. The NSFNet represented a concentration of policy authority in the

earliest Internet days. Internet Exchange Points (IXPs) are, by their very nature, concentration points, yet they clearly provide a lot of value in exchange for the risk of concentration, because the concentration also promotes open peering and effective interconnection. There have often been worries about the diversity of software code bases and the danger of monoculture – a danger that continues to haunt us as we move from open standards to APIs. The thing is that, over time, particular choke points move around, risks change, and mitigations to old issues emerge. One big issue in the late 1990s, for instance, was the total dominance of Microsoft's web browser. That may tell us something about permanent favourites.

Concentration is not obviously always bad. Internet exchange points really work by concentrating traffic in some places, and this lowers latencies and costs and encourages the development of many interconnections more effectively than other models. This undercuts a different kind of concentration that would result from overwhelming dominance by one network. Similarly, it's true that AWS currently looms large in every market in the world. But it's also true that, 10 years ago, hardly anybody on the planet could afford some kinds of facilities that AWS will now rent to any one of us by the hour. That is no small thing.

The application layer also looms large now. Big web

services, living on top of the Internet, often function as a proxy for the Internet itself – both to the users and in the minds of regulators and policy-makers. Yet these applications are in fact themselves dependent on the underlying Internet infrastructure. That is why the application operators are often keen participants in communities like this. To a regulator, Facebook might look like it is running everything; but Facebook obviously knows they need the Internet infrastructure to remain healthy and strong. To do that, just like everybody else, they must collaborate and interconnect. The way that interconnection works from year to year might change – everyone here knows about Geoff Huston's observations on the death of tran-

sit. Yet the interconnectedness continues, even if in a changed form. And keeping that fact apparent to the policy and regulatory minds requires, again, engagement of the Internet community.

This is why the Internet Society's 2019 Global Internet Report is more an opening in a conversation than a definitive answer on the topic of concentration and consolidation. Our investigation in preparing this year's Report led us only to the conclusion that we did not know enough about how the Internet is changing. It is tempting to reach for simplistic answers about the Internet's future. But it has been through several changes in the past, and many of those changes surprised even those who were watching

closely. I think if you had asked most observers around the time when the Web first appeared whether http would become the default protocol for the Internet, you'd have received a baffled look. I remember demonstrations of hiding http messages in DNS messages, in an effort to circumvent filterns and firewalls. Nowadays, many people regard https as the potential saviour of the DNS, because you can almost always get https messages through.

It's easy to make a mistake in "picking winners". Instead of *picking* a winner, it's possible to alter the environment to *create* certain winners, and permanently cement them in place. It's pretty easy to imagine regulation that could create systematic advantages for incumbent services

on the Internet, just as many telephone regulations created built-in advantages for the existing ways of doing things in the past. Those regulations sometimes turned out to be significant barriers to innovation. So, I hope to entice you to read the next Global Internet Report, and to help us in the next year answer the questions it raises.

The fundamental value of interconnectedness is also why we continue at the Internet Society to believe in efforts like MANRS – Mutually Agreed Norms for Routing Security. In a network of networks there is no centre, so there is no centre of control, so reliable interoperation is everyone's responsibility. We need norms.

Interconnectedness is also why we continue to work,

in the communities that make them, on IXPs. They are an example of the network community, represented here at APRICOT, working out solutions that are good for each of the constituent networks, and also good for the Internet as a whole.

For the most effective way to ensure we don't break the Internet is to improve it in Internet-like ways. And the most effective way to ensure global connectivity for everyone is to use the Internet. But this means that network operators need to show that their self-interest provides the necessary protections for the Internet, lest we get more poorly-considered regulation that favours other network technologies. We need people to sign up for (and

show they are implementing!) efforts like MANRS, and to start validating routes. We need to make sure that IXPs remain neutral, effective and attractive alternatives to "building your own" and not exchanging traffic. These are real, effective actions that network operators can take to discourage regulatory overreach in the Internet.

Nobody Internets alone

The reason meetings like this are so important is because they carry forward the most important part of the Internet: the community of operators who make it happen. The Internet needs a strong, technically-informed community

to exist, because understanding how our mutual interests give us all the global Internet is a better way to create global connectivity than a centrally-planned system would be. But that means we need to make sure we keep alert to the drift we see in the world today, away from the Internet style of engineering toward systems that do not place both the gains and the responsibility for deployment in the same place. We know this system works, thanks to communities like this. We need to ensure that the whole world understands what they might give up if they give up this architecture.

I hope you will join us in ensuring that the Internet is for everyone. Thank you very much for having me here

today.