

November 2021
Geoff Huston

NANOG 83

The network operations community is cautiously heading back into a mode of in-person meetings and the NANOG meeting at the start of November was a hybrid affair with a mix of in-person and virtual participation, both by the presenters and the attendees. I was one of the virtual mob, and these are my notes from the presentations I found to be of personal interest. I hope you might also find them to be of interest as well.

Famous Internet Outages

The year 2021 has not been a good year for Internet outages. The widely reported Facebook outage in October 2021 came after outages from Akamai in July 2021 and Fastly in June. All these events received widespread coverage in the media at the time. The thing is that 2021 does not stand out all that much from Internet outages in previous years, and we can confidently expect to see more outages in 2022 and beyond. A panel session at the start of NANOG 83 discussed some of the more memorable outages and their causes.

One of the early outages occurred on the 2nd of November 1988 when a virus code devised by Robert Morris ran rampant across the Internet. It exploited a buffer overflow hole in a resident daemon process to inject and then execute code. The code was tailored to DEC VAX and Sun-3 systems, the predominant host systems on the Internet at that time. The issue was the very high replication rate of the virus, injecting multiple copies on systems which then brought the host to its knees.

Hostile malware taking out large parts of the Internet have continued over the ensuing years. There was the 2003 SQL Slammer worm that exploited a buffer overflow bug in Microsoft's SQL server implementation. The worm code was tiny (just 376 bytes), the transport was via a single UDP packet and it targeted servers, so it infected the high-capacity core of the Internet extremely rapidly. In 2008 there was a similar replicating malware system in the form of Conficker, again targeted at Microsoft Windows platforms, with an aggressive spreading algorithm.

Such malware incidents that have a significant replication factor and easily discernible network signature raises the question of the role of ISPs in such attacks. To what extent should ISPs take on the role of the Internet's universal firewall? Or should they position themselves as totally neutral carriers and trade all traffic in precisely the same manner. Most ISPs would be keen to assist and protect customers, so they often undertake moves to block malware, but in many instances it's not quite so black and white. What is their position when the attacker and intended victim are customers of other ISPs and they are acting as a transit? What is the position where there is a request from a law enforcement authority from a foreign country?

It's not just hostile attacks. Sometimes it's a case of accidental damage cause by operational slip-ups. One of the well-remembered incidents was the AS7007 event from 1997. What happened here was that AS7007 leaked a large volume of more specific routes (all /24s), all with a new origin AS, namely AS7007. AS it emerged from one of the war stories at the NANOG panel was that AS7007 used RIPv1 to pass a collection of routes learned from the eBGP instances to the internal routers in the network. Folk with a long memory might remember that in 1997 Classless routing was still a bit of a novelty, and RIPv1 only used Classful prefixes. When it learned a set of routes from a Classless eBGP session it dropped the prefix length information and added the Class A, B and C prefix lengths derived from the value of the first octet of the address prefix. At this point the eBGP speaker learned these more specific routes and promptly propagated them out into the

inter-domain routing system. By today's standards it was not a big leak (some 6.000 routes) and didn't last very long (2 hours), but the more specifics cause a large-scale traffic redirection, which was noticed at the time!

Outages have been caused by submarine landslides, such as the multi-cable outage in the Luzon Strait close to Taiwan in December 2006 which disrupted six of the seven east-west cable systems linking east Asia with North America. In this case the repair of 18 individual cable faults took some time due to the complexities in locating the correct cables due to the extent of the cable movement on the seafloor following the event.

There have been train derailments that have taken out fibre runs, as the railway right-of-ways have always been attractive as potential fibre paths for long distance runs.

As well as such natural events, collateral damage events and similar there are also a set of deliberate outages where the local regime acts to shut down all Internet access in a country. Sudan is in a period of Internet shutdown since late October 2021. Similar shutdowns have occurred in Egypt, Syria, and Bahrain. Over the course of 2020, 29 countries intentionally shut down or slowed their internet communications at least 155 times, according to a new report published by Access Now, a digital rights group.

If practice makes perfect, we are probably only getting better at generating Internet outages!

Who Controls the Internet?

Bert Hubert provided a keynote presentation on the topic of who controls the Internet.

There is no shortage of contenders to claim such control. In the light of the previous comment about the increasingly common use of Internet shutdowns or slowdowns by national governments it certainly looks like a number of governments operate under the perception that their national corner of the Internet is under their control. This is a far cry from John Perry Barlow's 1996 Declaration of the Independence of Cyberspace (<https://www.eff.org/cyberspace-independence>), although this message is still being carried by activists of various shapes and forms. The internet has been built on a foundation of private sector investment, and there is a strong case to be made that the Internet is controlled by these corporate interests from the private sector, but this is likely to be an over-simplification. The massive shift of advertising expenditure from print media and television into the Internet could sustain the case that the Internet is controlled by advertisers. The perversion of many of the technologies of the internet into tools to enhance digital surveillance, all in an effort to improve the effectiveness of advertising and thereby increase its value of the advertiser, tends to sustain the case that all of this environment controlled by advertisers, and without advertiser-funded services the Internet would probably be a rather hollow shell, devoid of any forms of compelling content. One could also see the Internet as the expression of a deregulated market, where ultimately, it's the collective summation of user preferences that drives the entire show. From this perspective the case can be made that it's you and I, as consumers of these digital services, are in control of the Internet. Without you and I as subjects of such intense scrutiny and as targets of advertising there would be no advertisers, and with advertisers there would be no funded services, and without this universe of compelling digital content provided through these services there would be no Internet, or at least no recognisable Internet.

Let's go through these contenders one by one and let's first look at governments. Within some regimes the role of government in controlling the Internet quite low key, while in other countries it's far more evident. The North Korean situation is a relatively extreme case where the digital environment is highly curated and the tools and services, including email, social networking, search and video conferencing have been heavily adapted to meet local requirements. The surveillance capabilities used by the private sector elsewhere are still evident, but in the case the client is the government rather than the advertising broker. The situation in China has some similarities to North Korea, although on a far larger scale and at a level of technical sophistication that is one of the more advanced. While there appears to be the capability for China to essentially disconnect from the rest of the Internet, there are certain windows of visibility that permit Chinese entities to participate in the broader digital environment. The control measures used in China are shrouded within a cone of silence. There have been periods where Chinese enterprises have attempted to join in with the larger global market with their own approaches and technologies, but such efforts have been recently rebuffed by Beijing, or

rebuffed by the broader Internet market. There has been a long-held belief from other quarters that much of the Chinese technologies are either strongly derivative or unashamed clones of foreign technology.

But these two cases appear to be exceptions to the norm. The efforts in other countries have been less successful in imposing a government agenda on the national digital environment. In Turkey, Indonesia, and Iran the experience has been that active censorship and exerting control is hard work and there is a continuing escalation between block and counter-block. Much of the issue lies in the efforts by global technology vendors to increase the level of privacy and obscuration of user transactions and traffic, and national governments that are attempting to stay one step ahead of these moves find themselves in increasingly challenging positions calling for additional technical resources that simply may not be available.

Most governments do not have sufficient local capabilities to effectively block access to proscribed digital services. This is the case in Russia, where efforts to block some services provided via AWS and Digital Ocean were so broad and sweeping that it had negative impacts on the national economy, and the government found itself backing off.

Western Europe has its own issues, with protection of intellectual property rights and preventing digital privacy in the United Kingdom having a mixed track record of effectiveness. The European Union with GDPR and more recently NIS2 have had some impact on the Internet, but in many ways the measures are more cosmetic than substantial. Blocks on domain names and blocks on file sharing are sporadic with variable results in terms of their effectiveness.

The United States has always had a vexed relationship with the internet. For many years the pre-eminent position of the US government with respect to the carriage infrastructure elements of DNS names and IP addresses was part of a story that went, informally, along the lines of: "We're here to protect you from the vagaries of other governments. Some of those other governments are seriously scary!" At the same time the US has been captured by the intellectual property rights lobby, and the draconian provisions of the DMCA, the unilateral ability to intervene in international payments and the broad reaching powers that are encompassed in national security measures all point to a US set of interests that clearly extend far beyond the lines of the map that bound the geography of that country.

Now let's look at the case for corporate control of the Internet. Facebook's Content Oversight Board was not the first, nor will it be the last corporate tribunal to determine who can use a service, how they can use it and whom they interact with. Google, YouTube, Microsoft, Apple, Twitter, Flickr and many others use similar control mechanisms. These oversight bodies are often arbitrary, their decisions are final and binding, do not necessarily adhere to a consistently applied set of procedures, and are sometimes enforced without warning or recourse to appeal. For example, these days losing access to Facebook or WhatsApp can be a dire situation for a company that relies upon such platforms to sustain their relationships with their customers.

The corporate interest is always present in such actions, and respect for societal and human values, such as privacy, respect, remediation of harms and the forms of recourse that are found in many legal systems are either missing or provided only at the whim of the corporate entity. In this respect their control structures and behaviour is intentionally both absolute and unaccountable.

What about the case for technology control? The evolution of our technology has been a fascinating path. Early computers were "open" in the true sense of the word. Anyone could write code and feed that code to the computer. When we look at portable devices, or browsers, or many applications the best description of the openness of these platforms is "hermetically sealed". The behaviours of the platform, be it a device or an app, are set by the manufacturer and hidden behind defensive barriers including layers of encryption and obfuscation. A good case in point was the Skype app, where the essential asset of the application was the code itself. It was shipped in an encrypted binary, and almost everything in the binary image was obfuscated (<https://www.blackhat.com/presentations/bh-europe-06/bh-eu-06-biondi/bh-eu-06-biondi-up.pdf>).

The transport applications that interact with the network are collapsing into a single protocol, QUIC, which exposes just a UDP stream where both the payload and the session control channels are completely obscured

from third parties outside a very limited circle of visibility. This includes governments, network operators, the device platform, and other applications. Incidentally it also includes you, the end user. From this perspective, all forms of control are passing to the application.

This question of control can be phrased as a tussle between contenders, including governments, corporates, and the technology itself.

What powers will governments exercise to compromise the device and open it up for inspection, and what powers will the application use to hide its activity from both the device and the network?

While application behaviours have not normally been the subject of regulation will we see the scope of regulations broaden to attempt to curb how applications may behave, particularly with encryption and obscured data.

When a device vendor exercises complete control of the device's app developer ecosystem, preventing the device from loading unauthorised applications, and the control has an in-built not just a tithe on sales revenues, but a whopping 30% extortion racket (yes, I'm referring to Apple if you hadn't guessed) then who is really in control here?

It's unclear that governments have any effective control these days. The pace and scope of change has left many national regimes following in the wake of change of control trying to regain ground through regulation and taxation rather than setting out a rule-based framework in advance. The big technology companies have gained much in the way of control, including centralising the points of control into internal functions that have little in the way of broader public accountability. Privacy is good, particularly when the lack of privacy has been weaponised in this surveillance economy. Encryption is good as a means of strengthening this privacy and freely available information is good. But neither governments nor technology companies are effective in protecting the public interest through this process of change, and the benighted user is very much a victim in all this. We really need clearer role accountability in this sector.

Internet Routing Registries

ARIN's Brad Gorman presented on the topic of IRR Spring Cleaning. Internet Route Registries (IRRs) are a means of listing the combination of address prefixes and routing intention, intended to be used to validate the information being passed through the routing system. IRRs contain route objects, linking an address prefix to an originating Autonomous System Number (AS) and aut-num objects, describing an AS, its adjacent AS's and the routing policies it applies to these AS's. There are set objects that allow groups of route objects or aut-num objects to be grouped together and treated as a single entity.

So, what's the spring cleaning at ARIN? Historically the IRR operated by ARIN operated as a common resource for network operators, and there was little in the way of validation of entries that were placed in the IRR. Over time the IRR contained both accurate and up-to-date information and earlier information that had lost relevance over time, and clients of the IRR have issues in discriminating between the two. The response from the Regional Internet Registries has been two-fold. The first is the RPKI where address and AS holders can associate their public key with the resource in an RIR-issued digital certificate. However, the RPKI framework is missing explicit attribution and routing policies, both of which have their uses in the network operations community. In response to this ongoing use of IRR and a desire to explicitly label IRR entries that are current and authentic ARIN has introduced the concept of "authenticated resources" where the association of the address prefix or AS number with a resource holder is based on current information in ARIN's resource records. For some time now ARIN has been operating essentially two databases: one describing Authenticated ARIN resources and a second legacy database of resources that are not authenticated by ARIN. ARIN has indicated that it will be withdrawing the IRR-NONAUTH database on the 31 March 2022. You've been warned!

Deeper Peering

Lumen's Guy Tai presented on Lumen's proposed changes to its peering policies. Peering in the Internet ISP sector has been a thorny problem ever since we started using a multi-provider network. When a packet transits multiple networks on its journey from source to destination then who pays whom? In multi-provider activities, including the telephone industry and the postal system, each transaction was individually valued and funded by the sender and the revenue was apportioned to each contributing provider according to either a mutually agreed tariff or a universal settlement tariff. But the Internet was never able to establish a common definition of a transaction or an associated tariff model. Packets are just too small and too random. Instead, the retail model quickly adopted a simple flat fee model. The implication of this move was that inter-provider settlements needed to reflect the retail model. The ISP industry adopted an equally simple inter-provider model: Either you're the customer and you pay me, or I'm the customer and I pay you, or we just don't pay each other anything at all. This latter option, peering, is used when the two parties are of equal size.

It sounds fine in most respects, except that the attribute of "equality" is not really an objective judgement. Instead, it's a negotiation, and here there are obvious rewards for a smaller provider to peer with a larger provider. It's in the interests of the larger provider to make their "equality" threshold as high as it can. It's in the interests of the competitive smaller providers, and the market regulators who have an interest in continued competitive presence in the market, to make these "equality" thresholds far lower.

Lumen's latest effort is a play right out of the old telephone book. Instead of setting up a small number of peering points and exposing the entire network to the other network, the intent is to divide the network into a large set of, for want of a better term we will call "local call zones", and if you peer in a local call zone you only get to see the routes for the customers of the peer who are also in the same call zone. Now the provider is not exposing its transit network to leverage through peering, as smaller entities derive no competitive advantage as they need to have the same breadth of coverage in order to avoid being treated as a customer. As Guy concluded in his presentation: "Lumen has updated peering requirements to mandate peering in all Tier 1 interconnectivity markets and 2/3 of Tier 2 interconnectivity markets in the US." If Lumen's idea is to place pressure on regional or smaller ISPs and wean them off peering with Lumen by exposing them to the additional costs of transit then yes, this might just do that.

Scanners

Scanning the IPv4 address space has been a popular sport for years. It became popular when the Zmap scanner code was released by researchers at the University of Michigan. There are many such tools about these days, but the ZMAP tools are still popular, and as they say: "we've built nearly a dozen open-source tools and libraries for performing large-scale empirical analysis of Internet devices." (<https://zmap.io/history>).

Many of these scanners operate blindly, walking through the entire IPv4 space, but others appear to be more selective in their scanning, targeting networks within a particular geography, or targeting networks associated with a category, such as residential, enterprise or educational campus, according to a presentation from Max Resing of the University of Twente.

I'm a little sceptical about these conclusions, as the way the experiment is conducted can have a strong impact on the resultant data. Earlier work in setting up entire /8 prefixes as honeypots for scanners showed a strong pattern for blind scanning across the entire address range, which is inconsistent with targeted scanning.

Anycast

Is anycast an art or a science? Anycast is used extensively in today's Internet. It is used extensively in the DNS for both open recursive resolvers and authoritative name servers. It's used by a number of content data networks. It's used in mitigation of DDOS attacks. But how good is it?

The ideal is that anycast distribution should relate to geography, and each anycast distribution point should attract nearby clients. The reality is a lot messier and the anycast boundaries are blurred and do not correlate well to either distance or latency. Previous studies have indicated that around one third of end points are routed to a sub-optimal service point. It's also been observed that simply increasing the population of the

anycast constellation does not necessarily improve the outcome is the metric is average latency of an anycast service point. One factor behind this is that BGP minimises AS hop counts in its route selection process, not latency. A global transit network still counts as one AS. The second is that on the Internet AS paths are, on average, very short. What this implies is that the routing system is not capable of fine-grained discrimination between anycast service points.

This implies that the construction of a large scale anycast service constellation poses some challenges. If one wanted to predict the impact on an anycast service by adding or remove an anycast service point then the most effective way to do this is through a brute force excise of setting up a testable anycast service constellation on the Internet, setting up a large set of test points distributed across access networks and perform the *mxn* set of reachability tests on each configuration. The research work being undertaken by a team at Duke University is attempting to provide similar outcomes using a far more constrained testing setup. In an experiment using 15 sites, each peering with one of six transit providers, their prototype, AnyOpt predicted site catchments of 15,300 clients with 94.7% accuracy and client RTT's with a mean error of 4.6%. AnyOpt identified a subset of 12 sites, announcing to which lowers the mean RTT to clients by 33ms compared to a greedy approach that enables the same number of sites with the lowest average unicast latency.

Password Recovery Weaknesses

If you have ever forgotten your password, and that's most of us by now, then the common approach is to click on the “forgotten password” link and the service sends a password recovery link to your nominated mail address. What if the provider's DNS resolver does not perform DNSSEC validation, or the domain name of the client's email address is not DNSSEC-signed? Then the attacker can perform a DNS attack and cause the recovery message to be sent to the attacker's mailbox.

The underlying messages are simple. To service providers with user accounts: Passwords alone are a lousy security mechanism. Use Two Factor Authentication at a minimum! Oh, and use a DNSSEC-validating recursive resolver for all your domain name queries. No exceptions. The clients and client service providers: use DNSSEC to sign your domain name.

We've known these messages for years, but for many service providers they prefer to make it “simple” for their customers and persist in password-only account protection mechanisms. These days I see this as no different to treating these same customers with careless contempt!

Predictive Routing

A lot of effort has gone into the repair mechanisms used in BGP. The nature of a distance vector algorithm is that a BGP speaker may not have all the information at hand to immediately repair the routing state when it receives a withdrawal message. Instead, the BGP speaker withdraws the route internally, announces this withdrawal to its BGP neighbours, and then patiently waits for a neighbour to announce a replacement route.

There have been a number of approaches to mitigate this BGP repair delay, BGP Add Path being one of the more recent approached. In this approach BGP note only propagates its “best” path to its neighbours, but also any viable alternate paths that were viable. In this way a BGP speaker can process alternate paths when processing a withdraw and recover quickly.

In predictive routing this is taken a further step. Why not just allow the BGP speaker to remember previously learned viable next hop forwarding decisions and when the primary next hop decision is withdrawn just assume that a previously learned route still exists and just use it while awaiting an explicit BGP based repair? If this sounds a little dubious it should be remembered that the use of a default route has a similar outcome, which is telling the BGP speaker of a universal forwarding action to take when there is no more explicit routing information available.

IPv6 – The Next 10 Years

If you had asked anyone back in the year 2011 or so whether we would still be running IPv4 to support the Internet in 2021 I'm pretty sure that answer would be a very definite “No!” We were chewing through some

400 million IPv4 addresses every year at that point, and we were very dubious of the ability of NATs in IPv4 to scale up year on year. Ten years later when the same question is asked “Will we still be using IPv4 on the Internet in 10 years from now?”, then the answer from Comcast’s John Brzozowski is a surprising “Yes!”

The efforts in transitioning to a full dual stack environment are proceeding, albeit at a somewhat ponderous pace. Services are migrating over to dual stack platforms, and John’s analysis of the popular web sites points to a continuing uptake on IPv6 capability. At the same time the deployment of dual stack in access networks is following a similar steady pattern. One the hole user deployment of IPv6 follows the pattern of economies with a higher BGP per capita level, although Sweden, Norway and Spain are exceptions at one end of the GDP per capita spectrum and India’s ranking as the highest level of IPv6 deployment is a marked exception at the other end.

The end point of all of this should be that once the level of dual stack deployment reaches some yet-to-be determined critical mass it will then be viable to sustain IPv6 clients and IPv6-only services. At that point the pressure will be on for remaining IPv4-only services and access networks to install IPv6 services, and this would see the rapid decline in the use and value of IPv4. If we can’t determine in advance when we reached that point, then what indicator would signal to us that such a point is nigh. Counting services or user populations is fine, but it does not really assist in identifying the threshold point of change.

My preferred metric is the price of IPv4 addresses on the transfer market. When this price starts to decline then that’s a clear signal that the market as a whole has made the transition and now does not see IPv4 as a critical-to-have resource. Are we there yet?




Figure 1 – IPv4 Market price 2020-2021 - \$USD per address – Source: ipv4.global

No, we’re nowhere near that point using that line of reasoning. Over the past 12 months the market price of IPv4 addresses has doubled to around \$50 per address, according to the reports from IPv4.Global (Figure 1).

How long will this protracted transition continue? Nobody knows.

Quantum Key Distribution

In December 1926 Albert Einstein wrote, in response to a letter from Max Born describing the random and uncertain heartbeat of quantum mechanics, that “The theory produces a good deal but hardly brings us closer to the secret of the Old One. I am at all events convinced that He does not play dice.”

I will not pretend that I have any useful clue in the area of quantum mechanics, quantum computing and quantum cryptography, and while I truly appreciate the effort that Melchior Aelmans put into his presentation on Demystifying Quantum Key Distribution, I think the best I can do here is to note that in the area of cryptography at the very least this is an important topic and secondly you really need to review Melchoir’s material yourself and not any second hand interpretation that I could offer here. It can be found at https://storage.googleapis.com/site-media-prod/meetings/NANOG83/2395/20211102_Aelmans_Demystify_Quantum_Key_v1.pdf. His presentation pack has a useful reading list as well for those who are motivated to dig deeper into this fascinating topic.

NANOG 84

And that was NANOG 83 for me! Next time its NANOG 84 at Austin, Texas, February 14-16, 2022. And I guess we are all saying hopefully to each other “see you there!”

Disclaimer

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

Author

Geoff Huston AM, B.Sc., M.Sc., is the Chief Scientist at APNIC, the Regional Internet Registry serving the Asia Pacific region.

www.potaroo.net