The ISP Column
A monthly column on things Internet

April 2022
Geoff Huston

AusNOG ’21

AUSNOG 2021 was held in ~~September 2021~~ ~~December 2021~~ April 2022 in Sydney over two days. Here are a few notes from presentations at the meeting that I found interesting.

Automating network management

Network management was never the ’sexy‘ part of the Internet. For decades, network operators worked with ASCII command line interfaces, using a language that looked like it was borrowed from Digital's RSX command language of the 1970s. A whole new level of sophistication was meant to be introduced with Simple Network Management Protocol (SNMP) operated in write mode, but this was quickly discounted as a security nightmare. So, we concentrated on various forms of crude *Expect* scripts to automate the tedious parts of typing in commands, but it was not much of a change.

Then along came Computer Science and increased the level of sophistication of the tools. Ultimately, it’s still stuffing a sequence of commands into an interpreter on the device, but the automation tool is attempting to translate a higher-level objective into the specific device configuration sequence. At this point, the flood gates opened and today we have Napalm, Salt, Ansible, Pepper, Chef, Terraform, Cloudify, and so on. They all seem to behave like Bablefish of command-line interfaces (CLIs).

One interesting observation here is that while the sysops community has enthusiastically embraced the DevOps culture and taken up the use of such tools, the NetOps community has evidently been more reticent in this respect. Jacob Taylor's presentation at AusNOG explored some of the reasons for this. There is a visible disconnect between vendors and consumers of these tools. The vendors have tended to augment tools with an entire ecology and style of management that extends well beyond the operation of a configuration management tool and heads in the direction of a vendor-specific environment with little consideration of interoperation across tools. It was reported that open standards are largely ignored in this space.

There is also a human problem, in that while we talk about DevOps there are very few folk who are thoroughly conversant in both worlds. Individuals who know how to develop such systems are not often skilled in devising robust and effective operational procedures, and folk with operational expertise are often not skilled software developers, so the coupling of expertise in both subject areas is not all that common.

The network operational environment is also not exactly a tolerant and forgiving one. A configuration error can not only impact the local network, but also impact the network’s neighbours, and further. While automation tools can alleviate aspects of the operational workload, it is one more potential thing to go wrong.

There is also the issue of incrementalism vs reload. Over in the configuration world, we’ve gone overboard in our use of filter lists, and they are often built incrementally over many years, and when that occurs, the result is not pretty. How to manage this is a challenge. One approach is to allow filtering rules to be described in a consistent way and have the tool translate this filter snippet into a set of filter entries in the syntax of the target device. This automates incrementalism, and often replaces one source of complexity for another. The other approach, reload, places the set of filter objectives into the management tool using a high-level language, and have the tool generate a complete filter list which is used to replace the in-device set when the configuration is loaded.

It's also relevant to ask about the overall approach to network management. The original CLI approach uses a device management style where individual devices are configured, and the orchestration of the network response is intended to be an outcome of these individual device settings. The other approach, which has proved to be more challenging to implement in simple ways, is to describe the network's intended responses, and allow the tool itself to determine what specific configuration needs to be loaded on each device in order to achieve the intended service outcomes.

As was pointed out in a related forum a couple of years ago; how is it that we are working on automation systems that are capable enough can support driverless cars, yet the seemingly simple objective of network automation is still elusive?

DDOS detection and mitigation

DDoS attacks continue to plague the Internet Service Provider (ISP) space. In this space, the old and simple attacks are still the best attacks, so it should be no surprise that when an attacker is attempting to direct a significant amount of traffic to an intended victim, UDP amplification attacks are very common. The technique is simple: send a UDP query to a server using the DNS, NTP, SSDP, Memcached, or CHARGEN service. The only prerequisite is that the server does not make any attempt to authenticate the source address of the query, and the response is far larger than the triggering query. Most of these attacks are short and intense. The intensity is such that they can not only impact the intended victim, but also saturate the network infrastructure close to the victim.

If a network operator wants to develop a network response to such attacks, then they typically need to configure their network to pass flow profiles to an analyser and use a flow analyser to detect anomalous traffic signatures and then pass this signature back to edge routers in the form of a drop filter or a redirection rule. There are off-the-shelf solutions for this task, such as *fastnetmon*, or you can write your own. An Australian retail ISP, Swoop, has gone down the write-your-own path, and presented their approach that they say allows for rapid detection and filtering of these types of DDoS attack.

Whenever I hear about mitigation methods for source address spoofed UDP-based attacks, I can't help but wonder if widespread implementation of source address validation and filtering would be easier for everyone in the long run.

The bigger picture sees malicious attacks and malware as a long-term growth industry. If the intent of the attack is nothing more sophisticated than simple service disruption, then the prognosis is not looking good. The inexorable rise of high-volume low-quality devices with Internet of things (IoT) devices simply places more stress on the task of attack detection and mitigation. When devices are both unmanaged and unable to defend themselves, and in some cases are so poorly engineered that they are wide open to hostile exploitation, then the outlook is pretty grim. If we are relying on networks to detect and prevent such attacks, then it looks to me like a sad case of misplaced hope.

The Australian broadband environment

It looks as if every national regime has struggled with upgrading the common access infrastructure from legacy twisted pair copper tails used by telephony to some form of digital infrastructure. It can be a capital-intensive undertaking, particularly in low density suburban contexts that are a feature of Australian cites, and the mismatch of the demands of conventional capital markets and the low returns that are achieved in cable infrastructure have created various customised solutions in each national environment.

In Australia, this upgrade of the national infrastructure was undertaken as a public program funded in the first instance as an item of capital works. After some intense lobbying between cable and wireless proponents, the government of the day took the decision to fund a program of installing a fiber optic connection to almost every residence. The government of the next day revised that decision to use a cheaper hybrid approach, installing a fibre tail to a kerbside node, and reusing the copper loop for the connection into the residence, using VDSL2 technology. This access infrastructure was meant to be a provider-neutral last-mile system, and the larger role of national and international transit was the role of service providers.

This means, that if you are an ISP in Australia and you are servicing *National Broadband Network* (NBN) customers then you need to connect to some, or all, of the Points of Interconnection (POI) with the NBN. There are 121 of them.

How does an ISP get to these POIs? You can buy a POI access service from someone else who had already connected to these POIs, or you need to build your own. These POIs are generally old telephone exchange buildings, and there was generally a single ducting setup into and out of the building. So, building your own presents a whole new set of issues in ducting, trenching, urban landscapes and such. It's not for the feint-hearted, and its little wonder that the ISP market in Australia has once more been the subject of intense aggregation in recent years.

What was the point of this large-scale taxpayer-funded exercise? If it was simply to bypass Telstra, the incumbent telco who was showing a very definite aversion to investing in new urban infrastructure, then to some extent it worked, as public funds were used to build this access infrastructure and Telstra was cut out of the provision of basic access business for residential customers as a consequence. But if the aim was to dilute Telstra's encompassing dominant position in this space, then it failed comprehensively. Telstra is still the dominant provider in this space, and it now has access to this taxpayer-funded fibre infrastructure on highly advantaged terms, as the POI access infrastructure for Telstra was already in place. One AusNOG presenter noted that Telstra already had 370,000 km of ducts and tunnels and any competitor either had to rent access from Telstra or struggle with building out its own POI access. So, has the NBN bypassed Telstra? Well, yes, but in so doing it relieved Telstra of some of the more onerous costs and burdens of running a last mile access network and allowed the company to further entrench its dominant position in the Australian retail market. It is extremely hard not to be very cynical about this entire exercise!

Speed and Quality

The retail ISP landscape has always fixated on speed. Consumers expect that faster services will cost more, and this greater speed will directly relate to a superior service experience. ISPs quote bandwidth, or access speed to consumers ass the metric that denotes the quality of the delivered service. Faster is better, right?

Latency, or end-to-end delay, also has a lot to do with the quality of the delivered service. Transactions across a local network seems to be more responsive to users simply because of the lower latency, and not necessarily because of a difference in capacity. In addition, there is the stability of latency, or jitter. Highly unstable systems with a high jitter rate tend to confuse feedback-controlled protocols such as TCP and drive the protocol into slower data rates because of the difficulty in sustaining a stable round trip time estimate in the face of high jitter. Loss and bit error rates are also critical elements of the delivered service.

What are the factors that have a bearing on delivered service quality, and how do they relate to speed?

In addition to the factors of latency, jitter, and loss, location of the measurement endpoints for a speed test has an impact on the result. On the client side, it's the termination of the access service or a host connected inside the client network which may possibly have a Wi-Fi access link. On the server side it's a case of how deep inside the ISP network and how much contention there is between the service elements at the time of the test. Is this a test on an otherwise idle access service? What's the variation between this idle measurement and one where the service is loaded with other traffic? There is the consideration of retransmits, and the measurement of delivered capacity compared to delivered throughput.

The presentation exposed for me just how much we don't know about these kinds of systems we use trillions of times each and every day!

Why should ISPs care about DNS Privacy?

I must admit that this is a fine question, but I'm really not sure how to answer it. Andrew Campling's presentation did a good job of describing the various permutations of DNS over encrypted sessions, but I'm not sure that it managed to provide a good answer to this basic question. The issue for many ISPs is that the DNS is a business overhead, not a source of revenue.

The introduction of DNS over encrypted sessions has attracted much attention in DNS circles, but there's a distinct possibility that it's more noise than substance. Yes, we've proved that its possible to operate DNS queries over an encrypted session, but in some ways that not a big achievement, as its essentially a minor variant of DNS over TCP. So, if the question is why we all aren't using DNS over TLS, QUIC, or HTTPS, then it's a similar question as to why we all aren't using DNS over TCP all of the time. TCP is the plan B of DNS queries, used when UDP can't cope via setting the Truncation bit in the UDP response. We've built a massive amount of DNS infrastructure based on the efficiencies of unencrypted DNS over UDP. A number of studies has shown that a server's capacity is reduced by two thirds if the queries are placed over TCP rather than UDP. Or, to put it another way, the cost of DNS server infrastructure would need to triple to replicate the capacity of the existing UDP-based infrastructure. Who would be willing to pay?

So, at the moment, DNS privacy measures are not in the mainstream of the DNS, and while there are efforts underway in the browser world to detect encryption capability in the existing recursive resolver infrastructure and switch to an encrypted channel if it is detected, there is an obvious level of reluctance from the ISP-operated recursive resolver infrastructure to adopt this technology. It's a case, to put it crudely, of the ISP incurring additional costs without either additional revenue or competitive market advantage.

The current position for ISPs is quite pragmatic. If a customer wants to redirect their DNS queries to an external open resolver, then in most regulatory regimes the ISP is off the hook. It's not just off the hook in terms of passing the costs of the DNS service to a third party, but it's effectively off the hook in terms of regulatory obligations for content blocking. If the customer elects to tunnel through the ISP-provided infrastructure using a secured channel and bypass the ISP's implementations of its regulatory obligations, then the ISP simply cannot look into this encrypted session. Whatever national content regulation may apply to the ISP's DNS server, if the customer or the application chooses to bypass that ISP-provided DNS service, then that's not up to the ISP.

But this may change. There are moves in a number of national regimes to rephrase the content blocking obligations in a way that does not identify the DNS as the means of implementation, inferring that the ISP still has an obligation to block access to such content, even when the user is using encryption for their DNS queries and encryption for their content access.

How the industry is going to respond is unclear.

Encryption is used in all kinds of situations, and these days it's an integral part of the Internet ecosystem. It's a way for innovative applications to tunnel through legacy infrastructure. It's a way of creating a secured group environment for enterprise use. It's a way of protecting our environment against a diverse range of hostile attacks. For consumers, encryption has been represented as a major step in maintaining some vestige of personal privacy in the face of widespread unauthorised surveillance.

However, it seems that encryption in the DNS is being viewed as a step too far by some regulators. This leaves ISPs in a conundrum. Insisting, as a regulatory provision, that ISPs use some mythical power to break encryption in the DNS and in content delivery is a triumph of wishful thinking.

We've not seen the last of this.

Regulation, smegrelation!

For much of the 1990s the Internet community was struggling to get out of the overarching shadow of the incumbent telcos, and for this activity be recognized as a communications activity in its own right. It seemed to some folk at the time that the telcos entered the Internet market as an ISP to allow it to continue with a dominant position of representing the entirety of the telecommunications realm, and the ISPs found that they were being treated in an arbitrary and dismissive manner by the incumbent telcos and they were unable to find any regulatory relief!

How much things have changed! Australia is not alone in having its Internet sector now the subject of intense regulatory scrutiny. There are the various critical infrastructure provisions, data capture and retention, data protection, personal privacy measures, and various consumer protection regulations. There are cyber security measures, content filtering, and similar mandatory-to-adopt measures imposed on Internet Service Providers. Doubtless we will see various security measures, including DNSSEC, BCP 38, TLS-secured content, and malware blocking also being swept up into this mandatory-to-implement space. There are also some emerging geopolitical themes that already encompass bans on some equipment vendors on providing equipment into the mobile space. It seems likely that apps and services will be swept up into this wave with the borderless Internet featuring the same geographical boundaries that have already been erected in the services and content world.

If the Internet was largely ignored in its early years by regulators and politicians, its surely making up for it now! However, I cannot see this ending up in a good place. Increasing the regulatory overhead adds a burden to the industry that is prone to stifling various forms of innovation and investment. However, it is a large leap of faith to believe that a completely deregulated industry will not only protect the interests of the provider while working in the interests of consumers, and also work in terms of national interests. As an American regulator once said: 'It takes a lot of courage to trust in deregulation.' If that's the case, then it appears that the level of courage required to take the hands off the regulatory wheel is one that is not shared by Australian politicians and regulators.

So yes, the regulatory environment for this industry is only increasing in terms of number and scope of regulations. The Internet Association of Australia and fellow travellers in this space certainly are not short of work to do here!

AusNOG '21

These are my impressions of just a few of the presentations at AusNOG '21. The conference program can be found online. The organisers assure me that they will have the presentations uploaded to this site very soon!

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