October 2014

Geoff Huston

NANOG 62

NANOG 62 was held at Baltimore from the 6th to the 9th October. These are my observations on some of the presentations that occurred at this meeting. The presentations are available at https://www.nanog.org/meetings/nanog62/home

Measuring Broadband America

In 2010 the FCC took up a program that attempted to objectively measure the performance provided to customers from the various retail access providers. With the support of the industry actors, the FCC contracted with Sam Knows, and a program of sampled performance was commenced. Some 7,000 volunteers host measurement units that performs regular measurement burst tests to measure up/down speeds, latency and loss. The program uses Level3 servers, and MLab servers as measurement targets. The FCC program is now looking at augmenting the measurements to include aspects of quality and consistency of measurement, which presumably relates to video streaming. The measurement program currently is used on the DSL and cable-based access infrastructure. They are expanding this to the high speed satellite KA band services and the fibre access services.

The FCC claims that this program coincides with more honest reports of service quality, and companies are increasingly delivering no less, and sometimes more than what they purport to claim to deliver. Publishing this information benefits the consumer in terms of better consumer offerings and increased competitive pressure on the supply side of the industry. The FCC has observed that on cable-based infrastructure access speeds have been increasing by, surprisingly, some 30% per year, and the major holdback when this has not been the case is the use of legacy CPE. Upload speeds in particular have been improving. DSL access speeds have been largely static.

They are interested in research programs that can leverage off this deployed set of measurement units. There is a program with Georgia Tech to measure "in-home" congestion. They are also interested to see how this measurement exercise can be extended to measure access on mobile networks, particularly involving the 4 major MNOs (AT&T, Sprint, T-Mobile and Verizon). They are thinking about using a crowd sourcing-style of measurement, in a manner similar to speedtest, for example. The acknowledged issues with crowdsourcing is repeatability and consistency of measurement, but the FCC has released iPhone and Android versions of the SamKnows measurement code.

For 2015 they are looking at extending this program to measure DNSSEC and IPv6. (http://www.fcc.gov/measuring-broadband-america)

Understanding User Experience on Mobile Devices with the ICSI Netalyzr

One of the more memorable sides in this presentation was a reference to "map" drawn by Charles Minard in 1869 describing the statistics relating to the Napoleonic military campaign in Russia, and the subsequent retreat (http://en.wikipedia.org/wiki/File:Minard.png)

The presentation looked at the extent to which their measurement technique was able to infer the presence of middleware that impaired the performance of data services over 3G mobile services.

Burning Man - Network Automation

This was a presentation that described some of the aspects of setting up a relatively large scale temporary network in a desert location. Aside from the logistics of deployment of equipment there was an interesting story about their efforts to automate the network configuration and operational processes. They use NCG to take a description of the network in YAML and generate specific device config files for uploading.

SPAM in IPv6

The starting premise is that while there are many DNSBL lists for IPv4, the equivalent tables in IPv6 are far more sparsely populated. It is infeasible to do this at a level of /128 addresses, but /64 entries are very coarse. Previous studies (2009, 2010) found little SPAM in IPv6. One approach is to use connection attempts in IPv4 and IPv6 and see if reputation information in IPv4 can relate to an inference of reputation in IPv6. The data sets are still small, and the inference paths are uncertain. Of course this tie of V4 to V6 lasts as long as dual stack transition - the longer term approach of address-based reputation has some interesting translation issues when looking at privacy addresses and other forms of address agility. Even in IPv4 the issues of reputation in a world of increasing sharing of addresses has its own problems.

Project Turris

I like the CZ NIC folk! They see a problem, they look to see if they can help to solve a problem. Their work in the Knot DNS resolver and the Bird BGP implementation are both great examples of this simple approach of direct action. This presentation is another example of this process. The main objectives of this work is security research, end user security, and a proof by demonstration of the possibility of improving the state of the art of SOHO routers. The project is one of a SOHO router that acts as a distributed anomaly detector/honeynet using data capture points located in many end user networks, experimentation with a model of coordination of firewall rules based on collected data, and the use of a SOHO router that has integrated high quality IPv6 support, DNS and DNSSEC validation, update versioning control. CZNIC plan to distribute 1,000 such probes for a 3 year term. They opted to develop the units in the Czech Republic, using WiFi, LAN ports, USB, all on a 9 - 14W platform and all built in an open source context. The software platform would be based on OpenWRT, with NETCONF configuration. They plan for continual back to base firmware refresh. There is a lot of good thought in this approach, in that its an open source construction project that attempts to lift the state of the art CPE away from the variable to the consistently useful. As well as the integrated unit, there is Turris Lite (http://lite.turris.cz)

Creating carrier-grade Wifi Experience

Many network operators are coopting a rich deployment of access infrastructure and launching an overlay Wifi access network. HKT, Orange, Telstra are just a few. Comcast is not a MNO int eh US, but is looking to leverage its large deployed base of DOCSIS 3.0 modems and complement home hotspots with a small complementary outdoor deployment to provide a WiFi service for Comcast customers. The slide pack provided a high level of visibility into some of the engineering aspects of this deployment. One interesting observation was that 2.4G is now so heavily overcommitted that its now dead. One wonders if we are heading to a similar fate for the 5G band in some of the deniers area of population and technology concentrations.

DNS Privacy

There are many forms of enhancing the level of privacy in today's Internet. While we normally think of privacy in terms of the use of HTTPS and such, there are also concerns related to the unencrypted nature of DNS query and response. The response has been to look towards HTTPS as well, on the basis that port 443 is about the only end-to-end protocol that is not irretrievably mangled by middleware. The downside of the overhead of a TCP handshake, a crypto key exchange, added delays to DNS resolution and maintenance of potentially large scale TCP session state on the server is

countered, to a limited extent, by the potential of channel persistency to distribute the session establishment overhead across multiple queries. Above HTTPS the typical approach is DNS over JSON. This is not DNS 2.0, and not a rewrite of the DNS, nor any change in the query response model, but simply a means to focus on a way of protecting queries and responses in the DNS today. Much of the preliminary work has been on the DNS Privacy list, and DPRIVE has a BOF scheduled for IETF91.

TCP and Middleware

One quote I liked in this presentation was "the number of middle boxes is on par with the number of routers in a network" (Sherry, SIGCOMM'12, https://www.eecs.berkeley.edu/~sylvia/papers/fp150-sherry.pdf) In this case the experiment was to embed a hash of a number of middleware-mutable TCP fields into the IPID and received window size fields and then used Planet Labs and CAIDA ARK to generate a bunch of TCP sessions across the net to see to what extent TCP fields are being manipulated by middleware. They observed widespread manipulation of MSS values, timestamp options, SACK values, window scaling and window size, stripping of capability (such as MP capable), and ECN / TOS confusion. Their experiment found an incidence of TCP header modification at a level of some 50% of sessions. (http://tcphiccups.org)

100G and beyond

This was a dense and highly specialised presentation about achieving 100G and 400G on optical systems using coherent receivers. It's a dense presentation that is not readily summarised here, and the best I can do here is refer you to the presentation.

UTRS

This is a distributed automated filtering system that is intended to block hole attack traffic across a distributed set of networks, hopefully closer to the various attack sources. The idea is simple. The victim sends a notification of an address under attack to a distribution point, who, in turn, sends out an incremental filter update to the participants to add this network to their local discard filter list. UTRS uses an eBGP distribution system, so that the local router adds what is in effect a null route for the announced prefix set. In some sense the technology part is the trivial element - the trust and integrity component are the major parts here, and like all forms of remote trig erred black hole systems in the past with spam, the problem of fragmentation and partial approaches stymies the potential effectiveness of the approach.

Of course the other aspect of this response is that it actually fulfils the objective of the attacker! If the objective of the attack is to take the target victim offline, then this response takes the victim offline.

So from that respect it really only has effectiveness when the sources are listed in the BH feed. In a truly massive DDOS attack each source contributes a negligible component of the overall attack, and its not easy to reconcile the overheads of operating such a service with the scale of benefit it provides.

Cellular Out Of Band Management

A presentation that looked at using cellular data services as either a backup, or a full statute for wired services, focussing (not surprisingly) on the US market. In many ways its not anything special, as the cellular modem is just another interface on a router, and the uses of hot swap, primary and backup, VPNs and similar are equally possible on cellular as they are on any other IP transmission medium. The market offerings split between a conventional data offering, and a OOB management access system. The presentation covered a number of deployment scenarios, equipment and service offering and pricing.

SDN and Geni

Yes, the project is underway, yes, it relies on SDN to offer the researcher "slices" from a common elastic platform where the researcher can effectively define a custom view of forwarding paths and link together computation and transmission resources in ways that match the experiment's requirements. I'm still undecided of the utility of SDN outside of this particular context, but at the same time this particular requirement is one where SDN is remarkably apposite. Well done NSF to fund this!

Submarine Cable Infrastructure Trends

It may seem surprising, but yes, there are actors out there who undertake billion dollar investments into submarine cable infrastructure for what Tim Stronge of Telegeography describes as "stupid" reasons. The underlying message however is that the new cables, and the regrooming of older cables, are both causing continued competitive price pressure on cable infrastructure and most routes are seeing these conditions result in lowering unit prices. The speculation in the presentation included an examination of the price differential of transit prices for an example 10G transit service in New York and Tokyo, and concluded that there were signals that over time the transit prices will converge to a global constant (which is a somewhat surprising theoretical outcome in that scenario, in so far as it effectively devalues long haul transit infrastructure).

RPKI Deployment Considerations

This presentation was trying to expose some of Wes Georges' issues when considering the adoption of Route Origination Validation of BGP advertisements in a relatively large ISP, and highlights some of the challenges he has encountered. There are certainly some real issues here with deciding how to manage the generation and publication of Route Origination Attestations, concerning both the management of crypto keys and the management of an ISP's routing information, particularly as it relates to PI and PA, customer and internal routing origination, and the interaction between aggregates and more specifics from the perspective of ensuring that more specifics are adequately handled when a covering aggregate has an associated ROA. The current defined mechanism of integrating ROA validation outcomes as a local pref input, and placing this local pref into the context of an ISP's existing local pref settings can prove challenging. And of course if this fails is the failure mode "hard" or fail "open". Early adopters of a technology that has many moving parts, where the technology is not fully understood in terms of ubiquitous deployment, and is certainly subject to further changes is always going to be a "courageous". However there is only so much that can be done in the absence of practical deployment experience.

Disclaimer

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