October 2016

Geoff Huston

DNS OARC 25

DNS OARC is the place to share research, experiences and data primarily concerned with the operation of the DNS in the Internet. Some highlights for me of the most recent meeting, held in October 2016 in Dallas, were:

DNS DDOS attacks: This presentation was about using an authoritative server exhaustion style attack. i.e. launch large volumes of *random*.target_name style name queries into the DNS to target the authoritative name servers for the target name. Queries came in large volumes, and shared a common signature pattern (in this case the random names did not contain the ascii characters 'z' and '9'!) Query volume was recorded at some 20M per sec initially and then tailed off. Volume of these queries by end point is variable. They started looking for the attack, and then found that the source code was posted to hack forums. The source code looks to use Google, Version, Level 3 and Hurricane Electric as the recursive resolver intermediary. The attacks are typically 10 minutes in duration. The bot end population was of the order of 380,000 endpoints, but this has been dropped to 300K and falling due to reactions to the Krebs DDOS attack.

Privacy and anonymity on the DNS. This presentation started with the observation that privacy is hard in the DNS and the recourse is via various forms of encryption on the query/response channel. However, side channels are prevalent in the protocol, and implicit signals (timing, query and response sizes and the outer headers) can be informative about the query even if the payload is opaque. This looked at padding, interleaving, resolver delays and chaffing as a means of further obfuscating an already encrypted channel to prevent sophisticated attempts to reveal the payload without breaking the encryption. To some extent I guess that this is a relatively academic exercise with rather limited areas of application. No doubt there may be some parts of the world that share such extreme sensitivities, but it's not clear to me that this is all that relevant to the average Internet user!

IDNs - the issues of allowing the universe of non-ascii characters into the DNS environment was given an eclectic twist when Shane Kerr decided to register a domain name consisting solely of characters drawn from the Egyptian hieroglyphic character set. He found a registrar willing to perform the registration via the equivalent punycode representation, and managed to delegate the name. However, getting a domain name certificate proved to be beyond the CAs that he tried! There is a deeper point about IDNs and the broader topic of so called Universal Acceptance - the idea that we can manage simultaneously hold two representations of the same name using completely different name forms and never get confused was always going to be extremely challenging topic and the solutions that we use today have a very ad hoc feel to them.

A little over a year ago the "Yeti" DNS environment was proposed as a means of operating an alternate DNS root structure in order to facilitate testing of aspects of DNS root operation that would not perturb the production DNS root. The project now encompasses some 26 distribution points (or "roots" in the Yeti context). They have experimented with multiple Zone Signing Keys, and found some issues with IXFR. They have also experimented with a large ZSK, and rolling their KSK. It's unclear if there has been any substantive discoveries so far, which is in some ways is a testament to the robustness of the mainstream DNS structure.

The presentation by Paul Hoffmann was of interest if only to highlight that the DNS appears to work in spite of a considerable level of neglect and cruft. He examined the state of name servers used in second level domains, looking at the issue of orphan glue records, IP address and name distribution and similar. For me the presentation gave the impression of "well that's cute … but what does it mean?" Any examination of a large set of DNS zone files leads one to conclude that DNS is largely broken, and there are a lot of folk that either get it wrong or allow bit rot to creep into their configurations. Equally there are a set of folk who appear to tinker with code that allows quite bizarre zone configurations to work almost in spite of themselves! But aside from concluding something that we are already know - namely that the DNS works by chance more than by design, I'm not sure what the message is. Perhaps that's all there is - folk are largely operating their little part of the DNS universe largely by using somewhat corrupted recipe books and getting parts of it wrong. But as long as it meets some locally defined objective then no-one really cares about the rest. When an external observer applies the blowtorch of precise correctness to the DNS standards bible then it is little surprise that one sees extensive brokenness! Yet it appears to work, in so far as either the name has fallen into disuse and the fact that is does not resolve is of no consequence, or despite all the configuration anomalies, the name still resolves!

In news from the DNSSEC area, CZ plans to shift their keys from RSA to ECDSA P-256. They have studied resolver capability using Atlas and APNIC Labs data, and have concluded that there is little impact in such a shift. Their approach will be to double sign with ECDSA and RSA to start, but the Unbound resolver was only patched to allow this in a code release as of October 2015, so they are in no hurry. They are aiming for June - July 2017 to start this process of change.

One presentation added to an interesting protocol debate. When DNSSEC was introduced the decision was made to introduce the validation tests in a manner that was backward compatible. When a DNSSEC-validating recursive resolver is unable to validate the response that it was about to pass back, it does not pass the response and instead indicates this with a SERVFAIL error. obviously it has not failed, and the true error is a validation failure, but unless DNSSEC used new error codes this was not an option. So what happens when the DNSSeC signatures fail is close to the question of what happens when the authoritative servers for a zone fail. Nominum reported a study of this failure case, and saw the query volume rise in response. It seems that when the DNS servers are reported as failing the resolvers appear to redouble their efforts and increase their query rate. Logically one would think that the opposite reaction, and a deliberate response of reducing the require rate would make more sense, but then this is the DNS!

The sport of measuring availability and adjacency of all of the elements of the anycast constellations of each of the root servers continues. My suspicion is that this starts out as an exercise in seeing if it were possible to mount such an observation, and now a number of folk do this. But to a very real extent the latency to any particular root server is irrelevant to end users! If the aim is to return good new fast (the IP address of a defined name that has an IP address) then the caching behaviours of recursive name servers largely eliminate the use of the root in name resolution in the first place! All the root can tell the recursive resolve is the name servers of the 1600 or so top level names, and no more. And for most popular names this information iOS held in a local cache simply because the names are popular. So while the data analytics are fun to perform, it seems to me to be an exercise in juggling numbers with no real practical application! The challenge, as far as I can see, is to demonstrate that there is some direct relationship between latency to the root servers and slow DNS resolution performance for the names that are resolved by applications.

Many authoritative DNS servers save their queries (and responses) for subsequent analysis. One approach has been to use the data for one ccTLD server to identify the relative use of various resolvers. The problem here has been the imposition of "noise" across the collected data. The noise includes many potential sources, including cache refresh, query log replays, automated systems and others. It would appear that the level of these "zombie" queries is now reaching the same volume as user-triggered queries, making the exercise of attempting to identify the resolvers that users actually use

an exercise in filter out out these unwanted noise sources. The work performed in NZRS used cluster analysis on the query patters to decide if a source IP address corresponds to a conventional; resolver or not. This work strikes me as work where the method is probably of more interest than the answer, because if the task is to identify the resolvers that users actually use, then there are more effective ways to answer this question. The most direct way to do so is to get as many users to ask a uniquely identifiable query against an instrumented authoritative name server, and then link the user against the resolvers they use.

No matter how well specified we think a protocol may be, and how well every corner case is identified, it always seems to be possible to device a new case that is not well covered. Afnic came up against the cast of so called "Empty Non-Terminal" names in the DNS. Somewhere between RFC5155 and RC7129 the case where a name itself is not defined as a terminal name, but has delegated names, and NSEC3 signing is used on the zone, then its possible that the result is not DNSSEC-validatable. The answer for Afnic? after seeing anomalous behaviour across a number of resolver code bases and a number of popular resolvers the simple response has been to change the empty status of the name to non-empty by adding a short TXT record! Problem fixed!

Many enterprises outsource running their DNS zone to a specialised DNS service provider, and there are number of such service providers that operate in this space. It's not surprising that from time to time these entities may change their DNS service provider. Prior to DNSSEC it was a case of coordinating a change of the name server NS records and all was done. With DNSSEC it's a little more involved, particularly if you are after a "zero fail" outcome where the zone is always valid for validating resolvers. RFC6781 proposed a standard way of doing this, but it presupposed that each operator could publish subtly different versions of the same zone file. If this is not possible then the steps of shifting the NS records, the DS records and the in-zone NS and DNSKEY records have to be handled in a particular order and in a particular way. The presentation from Rightside detailed their experiences in moving a couple of signed domains and the steps they followed to achieve this.

The increasing use of IPv6 in access networks has some interesting side effects. Sebastian Castro reported on a situation where two of the nameservers for .nz were hosted on an IPv4-only provider and therefore had no IPv6 address record in the glue records. It was noted that the query level for AAAA records jumped significantly at the start of 2015, and was only reduced once AAAA records were added for these two servers. This may well be a side effect of dual stack application behaviour where the application queries for both the AAAA and A records of a zone and the lack of AAAA in the glue records appears to trigger additional queries. The presentation noted that there are 532 authoritative nameservers for TLDs in the root zone that do not provide AAAA glue records, so this problem may be more widespread than just NZ!

OARC meetings are a solid dose of DNS! I found the two days of presentations stimulating, and at times thought provoking. Despite its overt simplicity the DNS is a very subtle piece of critical Internet machinery, and when one considers the diversity of the way in which folk manage their zones and the way in which resolvers and authoritative servers behave, the fact that the DNS works as well as it does, let alone works at all, is perhaps the true wonder of the Internet!

All these presentations (and more) can be found at:
https://indico.dns-oarc.net/event/25/timetable/ - 20161015.detailed and
https://indico.dns-oarc.net/event/25/timetable/#20161015.detailed

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 building the Internet within the Australian academic and research sector in the early 1990's. 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 and chaired a number of IETF Working Groups. He has worked as an Internet researcher, as an ISP systems architect and a network operator at various times.

www.potaroo.net

Disclaimer