The ISP Column
A column on things Internet

Geoff Huston
July 2017

Notes from IETF 99 – The IEPG Meeting

It's hard to classify the IEPG meetings that occur at the start of the IETF week. Many years ago they had a role in allowing network operators to talk to other operators about what they were seeing and what they were thinking about. Those days are long since over, and today the IEPG meetings present an opportunity for an eclectic set of diehards to listen to an equally eclectic collection of presentations that wander over much of the topics of today's Internet, without any particular common theme or filter.

George Michaelson presented on a study he had done in trying to identify if there was a pool of users who were incapable of using transport security services (TLS - essentially an encrypted transport session service) to retrieve web objects. These days there is a strong push to place all content on secure servers, and rightly so, but the niggling doubt remains as to whether we are leaving some users and some equipment stranded if we were to say "this content is only accessible over a secure session". The exercise was one of attempting to measure the level to which users showed a behaviour that was consistent with being unable to load a web object over a secure transport session. The problem is that in this case you are not measuring a behaviour - you are measuring the absence of a behaviour. Measuring the absence of a signal is hard, and in this case there are many reasons why a user may not fetch a web object, particularly when executing a script embedded in an online ad. The result is that the signal that a web fetch has failed to occur is a very noisy signal. This presentation showed the extent to which advanced statistical tools can assist in trying to extract a signal that would otherwise be buried within the normal levels of noise.

Some weeks ago I tried to get my domain name registrar to add DS records to my name record so that I could enable DNSSEC on the name. The conversation unfortunately did not go far, as the registrar does not support this. But why should I have to make this a conversation with the domain registrar anyway? RFC 7344 and RFC 8078 describe a method to automate the entire process. The child publishes the intended DS and DNSKEY records associated with the new key in the signed child zone, and the parent may subsequently collect this data, either by periodic polling or in response to an explicit push notification. All well and good, but where are the tools to support this automation? The CZ.NIC folk have released a couple of tools that support this automation of DNSSEC key management. FRED, their open source registry tool, will periodically poll for the CDNSKEY record and if found commences the process of uploading the key and installing the new DS record. The KNOT resolver supports KSK rollover with automated CDNSKEY submission. There is still some concern that the population of DNSSEC-signed zones is relatively small, and tools such as this that try and make the entire process of key management simpler can only be applauded.

Configuring name servers is often a black art. Are more name servers better? Or is it a wasted effort? Should the name servers be widely distributed over the net? Is it batter to configure name servers behind the same IP addresses, using anycast, or use an explicit list of servers in a unicast manner? Giovane Moura reported on some research work in looking at the distribution of queries to authoritative name servers using RIPE's Atlas system to probe into an experimental name server configuration. The first result they found is that most recursive resolvers are seen to query all of the authoritative name servers over time. Resolvers do not simply latch onto one server and ignore the rest.

Some 60% - 70% of resolvers have what they call a “weak” preference for a single authoritative name server. A smaller cohort of recursive resolvers, some 10% - 30% or so, appear to have a “strong” preference for a single authoritative name server. The preference for the relatively faster authoritative name server is stronger when the authoritative name servers are closer to the recursive resolver, as measured by RTT. The conclusion is that while resolvers will query all authoritative name servers over time, but the queries will aggregate on the name server that is seen to be closest to the resolver. So how can one improve performance of authoritative name servers? This study suggests that rather than leaving the entirety of the selection process to the DNS, better results are achieved by letting the routing system pick the “closest” authoritative name server. The recommendation in this study is to use anycast on all of the authoritative name servers.

The Key Signing Key (KSK) of the root zone of the DNS will be rolled on the 11th October. There are a number of unknowns in this exercise, so we can’t tell exactly what problems may arise. There are potential problems with the larger DNS responses sizes associated with the various phases of the introduction of the new key, but previous measurements point to an expectation that this will have a marginal impact on the overall DNS system. The larger unknown is the issues with resolvers that have used manually managed keys and fail to apply the key update at the right time. The Root Canary project is intended to perform close observation on those resolvers that appear to perform DNSSEC-validation, looking for signals that might indicate some problem being experienced with the KSK roll. The preliminary results indicate that there is not a lot of support for new crypto algorithms these days and some errors about the way resolvers handle unknown algorithms (returning SERVFAIL rather than an insecure response. The web site for this work is https://rootcanary.org

I reported on the stats or more specifics in BGP. The presentation is based on the work described in http://www.potaroo.net/ispcol/2017-06/morespecs.html. The basic conclusion is that, as expected traffic engineering more specifics tend to show higher levels of routing instability, while overlay more specifics are the more prevalent. The surprising result is that IPv6 has some notable routing stability issues, and unstable prefixes in IPv6 show up to 100 times the update levels as compared to Ipv4. This obviously merits further investigation.

These presentations can be found on the IEPG’s web site: http://www.iepg.org/2017-07-16-ietf99/index.html

Author

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

www.potaroo.net

Disclaimer

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