ICANN and Technical Work: Really? Yes!

Steve Crocker

DNS Symposium, Madrid, 13 May 2017

Welcome, everyone. I appreciate the invitation to say a few words here. This is an important meeting and I think we will all learn quite a bit.

I'll talk mainly about technical aspects of DNS, but I chose the title to highlight an important and welcome development at ICANN. I've been watching where ICANN focuses its resources over the years. We've always had a strong technical foundation at ICANN, but, of necessity, political and contractual issues dominated the agenda. These issues haven't gone away, but I am now seeing much greater strength and much greater attention on the technical side. The chief technology officer is now part of the top-level management team, alongside the chief information officer, and the team under the CTO is noteworthy. Almost all have noteworthy accomplishments and reputations prior to joining ICANN. We have a team of stars, which is good for ICANN and good for the community. I hope this trend continues and that ICANN becomes known as a place for the best and the brightest in our field to consider joining.

The DNS layer is peculiar within the Internet ecosystem. There is a vibrant business in the selling of names, but almost no market for the operation of the lookup of queries. This means an awful lot of the development and research is done by the academic and non-profit community. From an economic perspective, this is unbalanced and I worry about the long-term health of this ecosystem. I'll return to this point at the end of my talk, and it's something for discussion over the next few years. Meanwhile, on the technical level, there is a LOT of activity, and this is what brings us here today.

By Internet standards, DNS is an old system, yet it is continuing to expand and evolve. The Internet of Things, the connection to DANE and other forces are likely to expand the number of DNS records in the entire system by an order of magnitude or more. The core design remains strong and can handle the load, but there will be pressure on various aspects.

Much of what you will hear today is about measurement. Measurement is essential and the increase in measurement activity and the development of better measurement and reporting tools is essential. Measurements tell us how well the system is working and often bring to light aspects we had not expected. Measurements are one part of a classic trio of modeling, metrics and measurements, and one of my messages this morning is that in addition to measurements, we need more extensive models and we need to work with models to find and predict how the system will behave under various forms of stress, with attention to three topics in particular: DNSSEC, performance, and software reliability.

DNSSEC

DNSSEC has progressed steadily. A significant event, the rollover of the key-signing key, or KSK, of the root is taking place in stages, with the big event scheduled for 11 October this year. Your tee-shirts have the public part of the new key, though I hope all of your systems are able to receive and install the new key without you having to copy it from your tee-shirt.

With respect to deployment, the progress is mixed, though compared to the deployment of IPv6 the progress is lightning quick. Here’s the status of DNSSEC deployment in the ccTLDs.

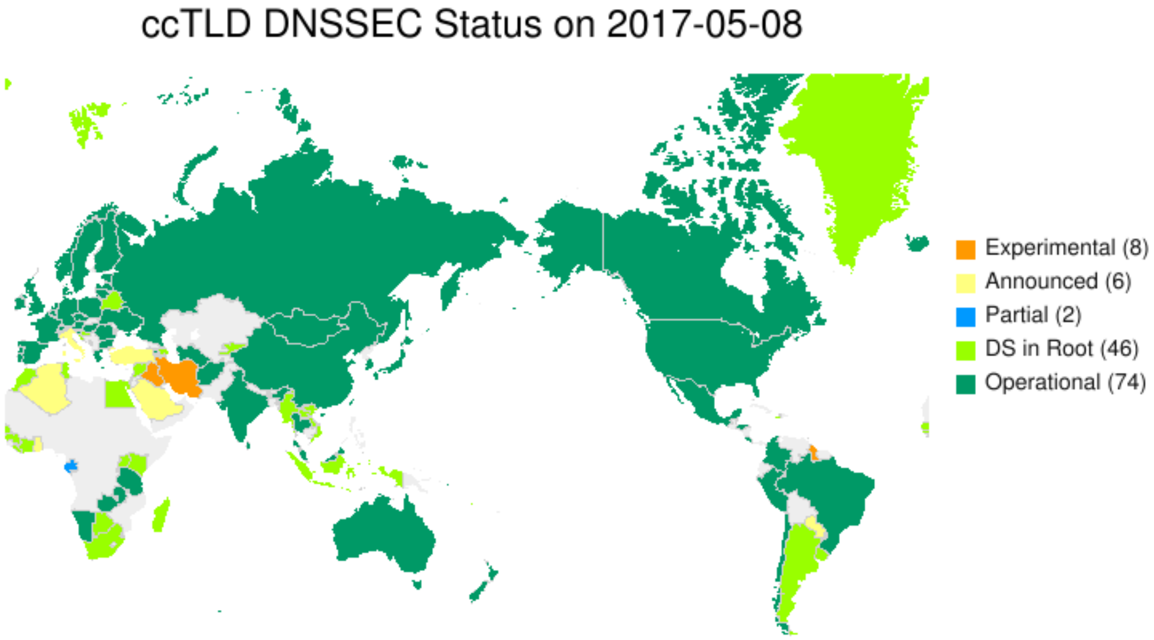

Dark green means the TLD is signed and it accepts DS records from its registrants. Light green means it's signed but is not yet accepting DS records from its registrants. This is usually a transitional stage. As you can see, with the exception of Africa, we're doing quite well, and even in Africa there is noticeable progress. The gTLDs are mostly signed as well. This was expected since we made it a requirement for all of the new gTLDs.

Signing of enterprises and actual checking of signatures are not as far along and need some help. We are also seeing some issues that have to be addressed:

- DANE integration
- Size of responses due to IPv6, key rollover over and long RSA keys
- Speed of validation – browser vendors care about this
- Local trust anchors, e.g. homenet, which is just one of many local trust anchors

PERFORMANCE

DNS, like all of the layers of the Internet, is full of complex interactions. One of the key elements in the DNS architecture is the caching resolver. Without these, the DNS system would not scale properly and the authoritative servers would break under the ever-increasing load.

The rule for caching resolvers is simple: the TTL says how long to keep prior responses, and it's ok to discard sooner if there's no space.

What actually happens is more varied. Some have minimum times, so they effectively raise the TTL. Some extend the TTL as long as lookups for that name continue to come in. Some generate "ghost" queries to the authoritative server for the negative queries they've seen recently. (I have to credit Geoff Huston for this tidbit, and I hope I've conveyed it accurately.)

These and other variations are motivated by the individual developer’s or local operators ideas of the best policy. What’s hard to tell, however, is the overall impact of these separate decisions.

The kind of measurement activity being reported today and being carried out across the world will shed some light. Let me suggest that in addition to the measurements, it would also be helpful to look at the various policies in the abstract to see what behavior is predicatable and where the stress points are likely to be.

I’ll give one small anecdote from the earliest days of the Arpanet. The Arpanet IMPs – that’s what we called the routers then – accepted a “message” of up to 8,000 bits from a host, broke it up into 1,000 bit packets, sent the packets to the other end, where they were reassembled into a full message for the receiving host.

In the receiving IMP, there was small buffer allocated to keep track of the several packets in a message – a "handle," and there was a small pool of these handles. This wasn't necessary for a single packet message, i.e. a message that was smaller than 1,000 bits, so this storage was only allocated for messages over 1,000 bits. In those early days, the actual traffic on the Arpanet tended to be either very short messages for interactive traffic, or very long messages for file transfers.

One other form of traffic was generated internally by the IMPs, and that was snapshots of the statistics of each IMP's interactions with each other IMP. These were created as internal messages and sent to the network monitoring center, and they looked the same as all other traffic.

At first these measurement messages fit into a single packet, but after the network grew past a certain point, the messages consumed two packets.

Do you see where this is going? Short messages didn't consume handles. Long messages consumed one handle for every eight packets. But the measurement messages consumed one handle for every two packets. And it turned out the pool of handles wasn't big enough.   And the entire network crashed!

A little analysis ahead of time might have helped.

Well, the interactions across the network are far more complex these days, and the analysis challenge is much harder. There are lots of parameters and not enough science about how to set them and how they interaction with each other. Bring the policies into the light and study their interactions from an analytical perspective, not just by measurement.

SOFTWARE RELIABILITY

We all know how vital the DNS layer is. If it breaks, the results are bad. One of my strongest fears is a latent bug in widely deployed DNS software and the threat that it might be exploited in a way that causes widespread outage or some other form of disruption.

Each of the developers does its level best to write quality code and to test it before deployment. But the main pressures on developers – remember I mentioned funding? -- are on adding new features or improving performance.

I'd like to see a strategic effort to transform the development process to provide much higher assurance, including better tools for analyzing the attack surface of this software. This is not an easy subject, but I think it's important, and, better yet, I think it's possible to make significant progress. There have been big advances in tools to analyze software and to provide assurance that software does what it's supposed to do and doesn't do what it's not supposed to do.

Summing up, my messages this morning are this:

1. ICANN is growing stronger technically and will make increasing contributions to the community.

2. Security is vital and will ever more so. DNSSEC is one of the main forms of protection. It's partially deployed and used; more is needed and there are some emerging challenges.

3. Performance issues are important and require a lot of attention. Use modeling as well as measurement to find the bottlenecks and other issues.

4. The reliability of the DNS software is vital. We cannot afford a large-scale disruption. We need to improve our analysis techniques and we need to insist on the highest quality standards.

Thank you for your listening. You have a fully packed program for a very substantive day.