February 2018
Geoff Huston

Peak DNSSEC?

The story about securing the DNS has a rich and, in Internet terms, protracted history. The original problem statement was simple: how can you tell if the answer you get from your query to the DNS system is 'genuine' or not? The DNS alone can't help here. You ask a question and get an answer. You are trusting that the DNS has not lied to you, but that trust is not always justified.

Whether the DNS responses you may get are genuine or not is not just an esoteric question. For example, many regimes have implemented a mechanism for enforcing various national regulations relating to online access to content by using DNS interception to prevent the resolution of certain domain names. In some cases, the interception changes a normal response into silence, while in other cases a false DNS answer is loaded into the response. As well as such regulatory-inspired intervention in the DNS, there is also the ever-present risk of malicious attack. If an attacker can pervert the DNS in order to persuade a user that a named resource lies behind the attacker's IP address, then there is always the potential to use DNS interception in ways that are intended to mislead and potentially defraud the user.

DNSSEC is a partial response to this risk. It allows the end client's DNS resolver to check the validity and completeness of the responses that are provided by the DNS system. If the published DNS data was digitally signed in the first place, then the client DNS can detect this, and the user can be informed when DNS responses have been altered. It is also possible to validate assertions that the name does not exist in the zone, so that non-existence of a name can be validated through DNSSEC. If the response cannot be validated, then the user has good grounds to suspect that some third party is tampering inside the DNS.

From this perspective, DNSSEC has been regarded as a Good Thing. When we look at the rather depressing saga of misuse and abuse of the Internet and want to do something tangible to improve the overall picture, then 'improving' the DNS is one possible action. Of course, it's not a panacea, and DNSSEC will not stop DNS interception, nor will it stop various forms of intervention and manipulation of DNS responses. What it can do is allow the client who posed the query some ability to validate the received response. DNSSEC can inform a client whether the signed response that they get to a DNS query is an authentic response.

The Costs and Benefit Perceptions of DNSSEC Signing

DNSSEC is not without its own costs and the addition of more moving parts to any system invariably increases its fragility. It takes time and effort to manage cryptographic keys. It takes time and effort to sign DNS zones and ensure that the correct signed information is loaded into the DNS at all times. Adding digital signatures to DNS responses tends to bloat DNS messages, and these larger DNS messages add to the stress of using a lightweight datagram protocol to carry DNS queries and responses. The response validation function also takes time and effort. DNSSEC is not a 'free' addition to the DNS.

In addition, DNSSEC adds further vulnerabilities to the DNS. For example, if the date fields in the DNSKEY resource records expire, then the material that has been loaded into the zone that was signed with this key also expires, as seen by validating resolvers. More generally, if the overlay of keys and meshed digital signatures fails in any way then validating resolvers will be unable to validate DNS responses for this zone. DNSSEC is not implemented as a warning to the user. DNSSEC will cause information to be withheld if the validating DNS resolver fails to validate the response.

Attacks intended to pervert DNS responses fall into two major categories. The first is denial, where a DNS response is blocked and withheld from the query agent. DNSSEC can't solve that problem. It may well be that the DNS "name does not exist" (NXDOMAIN) response cannot be validated, but that still not help in revealing what resource record information is being occluded by this form of interception. The second form of attack is alteration, where parts of a DNS response are altered in an effort to mislead the client. DNSSEC can certainly help in this case, assuming that the zone being attacked is signed, and the client performs DNSSEC validation.

Is the risk of pain from this second class of attack an acceptable offset against the added effort and cost to both maintain signed zones and operate DNSSEC-validating resolvers? The answer has never been an overwhelming and enthusiastic "yes". The response to DNSSEC has been far more tempered. Domain name zone administrators appear to perceive DNSSEC-signing of their zone an representing a higher level of administrative overhead, higher delays in DNS resolution, and the admission of further points of vulnerability?

The overwhelming majority of domain name zone administrators appear to be just not aware of DNSSEC, or even if they want to sign their zone they cannot publish a signed zone because of limitations in the service provided by registrar, or if they are aware and could sign their zone, then they don't appear to judge that the perceived benefit of DNSSEC-signing their zone adequately offsets the cost of maintaining the signed zone.

There are a number of efforts to try and alter the combined issues of capability and perception. Some of these efforts attempt to offload the burden of zone singing and key management to a set of fully automated tools, while others use a more direct financial incentive, offering reduced name registration fees for DNSSEC-signed zones.

The metrics of signed DNSSEC zones are not easy to come by for the entire Internet, but subsections of the name space are more visible. In New Zealand, for example, just 0.17%of the names in the .nz domain are DNSSEC-signed (ldp.nz). It appears that this particular number is not anomalously high or low, but, as noted, solid whole-of-Internet data is not available for this particular metric.

It appears that on the publication side, the metrics of DNSSEC adoption still show a considerable level of caution bordering on scepticism.

DNSSEC Validation

What about resolution behaviour? Are there measurements to show the extent to which users pass their queries towards DNS resolvers that perform DNSSEC validation?

Happily, we are able to perform this measurement with some degree of confidence in the results. Using measurement scripts embedded in online ads and using an ad campaign that present the scripting ad across a significant set of endpoints that receive advertisements, we can trigger the DNS to resolve names that are exclusively served by this measurement system's servers. By a careful examination of queries that are seen by the servers, it is possible to determine if the end user system is passing their DNS queries in to DNSSEC-validating resolvers.

We've been doing this measurement continuously for more than four years now, and Figure 1 shows the proportion of users that pass their DNS queries though DNSSEC-validating resolvers.

Perhaps it's worth a brief digression at this point and look at exactly what "measuring DNSSEC validation" really entails.

Nothing about the DNS is as simple as it might look in the first instance, and this measurement is no exception. Many client-side stub resolvers are configured to use two or more recursive resolvers. The local DNS stub resolver will pass the query to one of these resolvers, and if there is no response within a defined timeout interval, or if the local stub resolver receives a SERVAIL or a REFUSED code, then the stub resolver may re-query using another configured resolver. If the definition of "passing a query through DNSSEC-validating resolvers" is that the DNS system as a whole both validates signed DNS information and withholds signed DNS information if the validation function fails, then we need to be a little more careful in performing the measurement.

The measurement test involves resolving 2 DNS names: one is validly signed and the other has an incorrect signature. Using this pair of tests, users can be grouped into three categories:

a) None of the resolvers used by the stub resolver performs DNSSEC validation, and this is evident when the client is able to demonstrate resolution of both DNS names, and did not query for any DNSSEC signature information.

b) Some of the resolvers perform DNSSEC validation, but not all, and this is evident when the client is seen to query for DNSSEC signature information yet demonstrates resolution of both DNS names.

c) All of the resolvers used by the client perform DNSSEC validation, and this is evident when the client is seen to query for DNSSEC signature information and demonstrates that only the validly-signed DNS name resolved.

The measurement we are using in Figure 1 is category 'c', where we are counting end systems that have resolved the validly-signed DNS name and have been unable to resolve the invalidly-signed DNS name.

Use of DNSSEC Validation for World (XA)

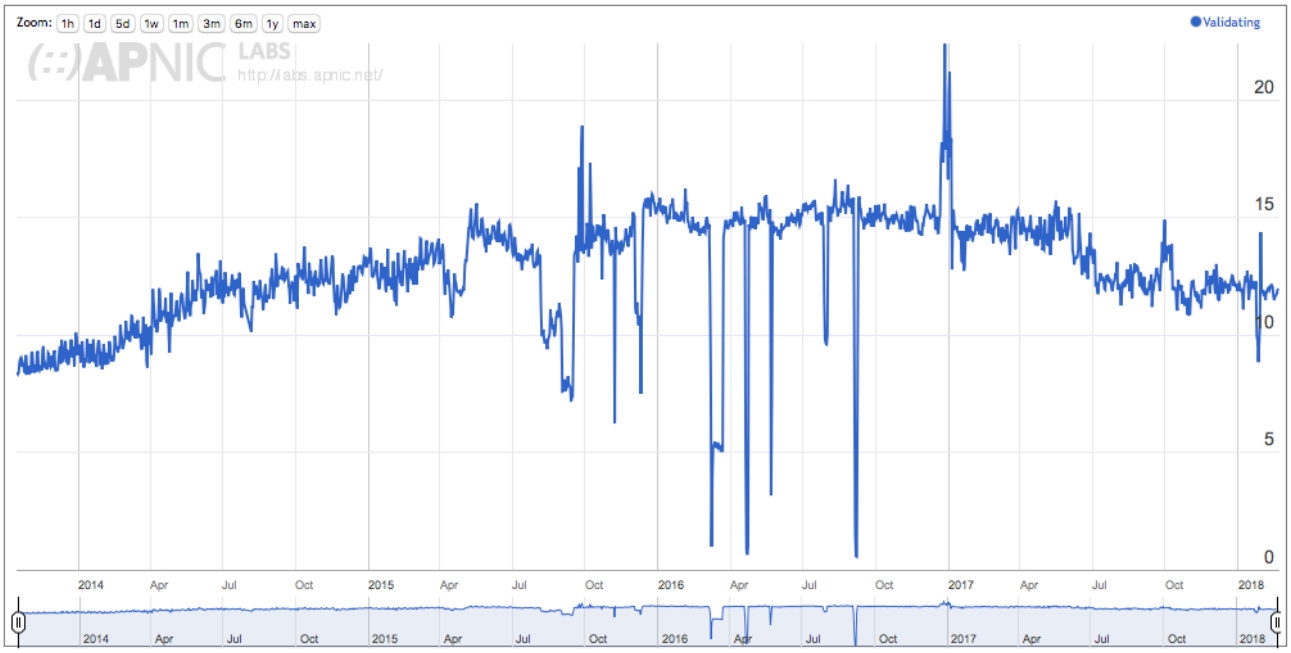

Figure 1 - Daily proposition of the Internet user population using DNSSEC-validating resolvers

Figure 1 shows a story that is consistent with an interpretation of "peak DNSSEC" from the perspective of DNSSEC validation. When we started this measurement in late 2013 we observed that around 9% of users passed their queries to DNSSEC validating resolvers. This number rose across 2014 and 2015, and by the end of 2015 some 16% of users were sitting behind DNSSEC-validating DNS resolvers. However, that's where it stuck. Across all of 2016 this number remained steady at around 16%, then in 2017 it dropped. The first half of the year saw the number at just below 15%, but a marked drop in July 2017 saw a further drop to 13%. At the time of the planned roll of the KSK the number dropped further to 12%, where it has remained until now.

If this number continues to drop then we stand the risk of losing impetus with DNSSEC deployment. If fewer users validate DNS responses, then the rationale for signing a zone weakens. And the fewer the number of signed zones, the motivation for resolvers to perform DNSSEC validation also weakens. Up until 2016 DNSSEC was in a virtuous circle: the more the number of validating resolvers the greater the motivation for signed zones, and the more the number of signed zones the greater the motivation for resolvers to perform validation. But this same feedback cycle also works in the opposite sense, and the numbers over the past 14 months bear this out, at least on the validation side.

From the validation perspective, the use of DNSSEC appeared to have peaked in early 2016 and has been declining since then.

Why?

Given that our current perceptions of the benefits of DNSSEC appear to be overshadowed by our perceptions of the risks in turning on DNSSEC, then the somewhat erratic measures of DNSSEC adoption are perhaps unsurprising.

I also suspect that the planned KSKroll and the last-minute suspension of this operation in October 2017 did the overall case for DNSSEC adoption no favour. The perception that DNSSEC was a thoroughly tested and well understood technical mechanism was given a heavy blow by this suspension of the planned key roll. It

exposed some uncertainties relating our understanding of the DNSSEC environment in particular and also of the DNS system as a whole, and while the measure was entirely reasonable as an operationally conservative measure, the implicit signal about our current lack of thorough understanding of the way the DNS, and DNSSEC, works sent negative signals towards both potential and current users of DNSSEC.

However, I can't help but think that this is an unfortunate development, as the benefits of DNSSEC are under-appreciated in my view. DNSSEC provides a mechanism to enables high trust in the integrity of DNS responses, and DANE (placing domain name keys in the DNS and signing these entries with DNSSEC) is a good example of how DNSSEC can be used to improve the integrity of a currently fractured structure of trust in the name space. The use of authenticated denial in DNSSEC provides a hook to improve the resilience of the DNS by pushing the resolution of non-existing names back to recursive resolvers through the use of NSEC caching. DNSSEC is not just being able to believe what the DNS tells you, but also making the name space more trustworthy and improving the behaviours of the DNS and its resilience to certain forms of hostile attack.

It would be a sad day if we were to give up on these objectives due to lack of momentum behind DNSSEC adoption. It would be unfortunate if we were to persist with the obviously corrupted version of name certification we have today because some browser software writers are obsessed by wanting to shave off the few milliseconds that need to be spent in validating the name against the name's public key when using DNS. It would be unfortunate if the DNS continues to be the weapon of choice in truly massive denial of service attacks because we are unable to deploy widespread NSEC caching to absorb these attacks close to the source. But a declining level of DNSSEC adoptions means that these objectives appear to fade away.

I'm hoping that we have not passed the point of peak use of DNSSEC, and the last 2 years has been a temporary aberration in a larger picture of progressive uptake. That would be the optimistic position.

Otherwise, we are being pushed into a somewhat more challenging environment that has strong parallels with the Tragedy of the Commons. If the natural incentives to individual actors do not gently nudge us to prefer outcomes that provider better overall security, more resilient infrastructure and a safer environment for all users then we are all in trouble.

If the network's own ecosystem does not naturally lead to commonly beneficial outcomes, then we leave the door open to various forms of regulatory imposition. And the centuries of experience we have with such regulatory structures should not inspire us with any degree of confidence. Such regulatory strictures often tend to be inefficiently applied, selectively favour some actors at the expense of others, generally impose additional costs on consumers and, as often as not, fail to achieve their intended outcomes. If we can't build a safe and more resilient Internet on our own, then I'm not sure that we are going to like what will happen instead!

Disclaimer

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

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