September 2021
Geoff Huston

# Another DNS OARC meeting

These are some notes I took from the DNS OARC meeting held in September 2021. This was a short virtual meeting with six presentations, but for those of us missing a fix of heavy-duty DNS, it was very welcome in any case!

## DNS Security Mechanisms

There isn't a single approach to DNS Security. Perhaps it's because there is no single threat vector. The DNS has channel vulnerabilities where DNS responses can be substituted or altered in various ways. The DNS infrastructure can be turned into a DOS attack weapon through various forms of amplification attacks. Knowledge of the user's DNS queries can compromise the user's reasonable expectations of privacy. What we've done in response to this spectrum of threat is to devise a collection of security mechanisms that are each intended to provide as defence against a subset of the larger collection of threat vectors. Masanori Yajima of Waseda University reported on an effort to survey the level of use of these various mechanisms in the DNS.

In the name resolution protocol, there are DNS Cookies and DNS over various permutations of TCP are intended to mitigate the risks of DNS reflection DOS attacks that are based on UDP source address spoofing.

CAA records in the DNS are intended to limit the actions of Certificate Authorities in the issuance of domain name certificates. If no CAA record is present in a domain name's zone, any CA is allowed to issue a certificate for the domain name. If a CAA record is present, only the CAs listed in the record(s) are allowed to issue certificates for that domain name. Obviously, CAA records are best used in DNSEC-signed zones. In the same space there is the use of DANE, domain keys in the DNS, where a hash of the issued certificate, or a hash of the public key that is certified in the certificate can be placed into the DNS. Again, this is best use in conjunction with DNSSEC. As browsers do not support DANE there is not much prospect of widespread adoption at present for DANE.

The efforts to combat various form of mail abuse have included several DNS mechanisms. There is SPF, the Sender Policy Framework, that lists in a DNS TXT record all the hostname or IP addresses that are authorised to represent themselves as being authorised to act on behalf of this domain name when sending mail. Again, DNSSEC would help here, but it's not a strict prerequisite. There is also DKIM, which, like DANE, places a public key in the DNS. An outgoing message that its purportedly originated by a sender in this domain includes a digital signature generated using the private key, and the receiver can look up the DNS and confirm that this DKIM public key can decrypt the digital signature. It should be no surprise to learn that this also works best with DNSSEC. DMARC records unify the SPF and DKIM authentication mechanisms into a common framework and allows domain owners to declare how they would like email purportedly from that domain to be handled if it fails an SPF or DKIM authorisation test. Then there is the MTA-STS TXT record in the DNS designed to prevent tampering with the STARTTLS part of a SMTP session between SMTP servers, and the TLSRPT to provide diagnostic reporting of the use of TLS in SMTP contexts.

And there is DNSSEC itself. This uses three DNS Resource Records, one for the digital signature (RRSIG), one for the signing keys (DNSKEY) and one for the interlocking of parent to child keys (DS).

There was a useful summary of these mechanisms in the presentation.

| Mechanism | Configure | Target | RR | TXT Format |
|---|---|---|---|---|
| **DNSSEC** | Server | <domain_name> | RRSIG, DS, DNSKEY | n/as |
| **DNS Cookies** | Server | n/a | n/a | n/a |
| **CAA** | Server | <domain_name> | CAA | n/a |
| **SPF** | Server | <domain_name> | TXT | v=spf1… |
| **DMARC** | Receiver | _dmarc.<domain_name> | TXT | v=DMACS1… |
| **MTA-STS** | Receiver | _mta-sts.<domain_name> | TXT | v=STSv1… |
| **DANE** | Receiver | _25._tcp.<domain_name> | TLSA | n/a |
| **TLSRPT** | Receiver | _smtp._tls.<domain_name> | TXT | v=TLSRPTv1… |

*Table 1 - DNS Security Mechanisms*
*from* https://indico.dns-oarc.net/event/39/contributions/867/attachments/822/1481/orac35a-yajima.pdf

How widely used are these security mechanisms?

When looking at adoption of technology in the DNS there are two parts to such a question: Firstly, how prevalent is the technology in terms of provisioning? In other worlds how many domains include these records in their zones (and keep them up to date, of course). Secondly, how prevalent is the technology in terms of use? Here we need to look at DNS queries or look at application behaviour across a broad sample of applications or users. This presentation concentrated on the first question, that of provisioning. The approach used was to take the usual subjects, namely the root zone, the top level domains and the top 10,.000 domains from the Tranco list (https://tranco-list.eu/). The results are shown in Table 2.

| Server Set | DNSSEC | DNS Cookie | CAA | MX | SPF | DMARC | MTA-STS | DANE | TLSRPT |
|---|---|---|---|---|---|---|---|---|---|
| **Root** | 100.0 | 100.0 | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 |
| **ccTLDs** | 56.7 | 81.1 | 0.0 | 6.3 | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 |
| **gTLDs** | 100.0 | 45.4 | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 |
| **Tranco** | | | | | | | | | |
| **Top 10** | 0.0 | 20.0 | 30.0 | 90.0 | 100.0 | 88.9 | 33.3 | 0.0 | 33.3 |
| **Top 100** | 4.0 | 21.0 | 48.0 | 86.0 | 96.5 | 84.9 | 5.9 | 0.0 | 5.9 |
| **Top 1K** | 9.2 | 13.8 | 22.7 | 88.1 | 92.9 | 74.0 | 1.5 | 0.6 | 1.8 |
| **Top 5K** | 8.6 | 18.6 | 14.9 | 87.8 | 89.9 | 58.5 | 0.7 | 0.8 | 1.0 |
| **Top 10K** | 7.7 | 17.4 | 13.0 | 86.8 | 89.7 | 54.0 | 0.5 | 0.8 | 0.7 |

*Table 2 – Adoption of DNS Security Mechanisms*
*from* https://indico.dns-oarc.net/event/39/contributions/867/attachments/822/1481/orac35a-yajima.pdf

For me some results in this table are surprising. I am not sure why, but I had thought that at present fewer than 1% of domain names were DNSSEC-signed, and the value of 9% of the top 1,000 names is completely surprising to me. I have often heard the response that "I'll sign my name when Google signs theirs!" and it's true that none of the top 10 names in the Tranco lists are DNSSEC signed! Spam is the scourge of the Internet, and mail administrators have been far more effective in getting the message out about SPF and DMARC. DANE and TLSRPT are basically unused. As to why this is the case, they have devised an interesting explanation based on degrees of difficulty, shown in Table 3.

| No. | Description | Points |
|---|---|---|
| 1 | DNS Resource Records need to be configured | 1 |
| 2 | DNS server configuration needs to be changed | 2 |
| 3 | A Mail server configuration needs to be changed | 2 |
| 4 | A Web server configuration needs to be changed | 2 |
| 5 | A third-party action is also required | 3 |

| Mechanism | Indicators | | | | | Difficulty |
|---|---|---|---|---|---|---|
| | 1 | 2 | 3 | 4 | 5 | |
| **SPF** | 1 | | | | | 1 |
| **DNS Cookies** | | 2 | | | | 2 |
| **DMARC** | 1 | 2 | | | | 3 |
| **CAA** | 1 | | | 2 | | 3 |
| **MTA-STS** | 1 | | 2 | 2 | | 5 |
| **TLSRPT** | 1 | | 2 | 2 | | 5 |
| **DNSSEC** | 1 | 2 | | | 3 | 6 |
| **DANE** | 1 | 2 | | | 3 | 6 |

*Table 2 – Difficulty of adoption of DNS Security Mechanisms*
*from* https://indico.dns-oarc.net/event/39/contributions/867/attachments/822/1481/orac35a-yajima.pdf

This intuition about adoption being hindered by more "difficult" technologies appears to be correlated with the adoption data,

If have reported in the past about the longstanding conversations in the DNS about the slow rate of uptake of DNSSEC and DANE. Some of these conversations head to a depressing conclusion that these technologies are simply never going to attain *mainstream* status and the industry appears to have reconciled itself to live with the consequences of an Internet infrastructure that has significant vulnerabilities that are going to be exploited from time to time. Other conversations point out that with the increasing use of outsourcing in DNS infrastructure as part of the larger topic of centralisation, these DNS specialists are actually far better equipped to deploy these security mechanisms that have a high degree of deployment difficulty. For example, DNSSEC validation was given a huge impetus when Google's public DNS resolver, 8.8.8.8, decided to support it. This second conversation points to a more optimistic conclusion regarding the uptake of these technologies.

### Mysterious Root Query Traffic

It always seems to be the case that whenever you look at the DNS you'll encounter something completely unexpected! In this case Christian Huitema and Duane Wessels reported on a high query volume seen at a root server. The queries all followed the same pattern of what appears to be a random string of 12- or 13-characters upper case alphanumeric characters in length, followed by a valid top level domain name The queries are all NS queries. They are seen at the root servers, but not at the top level servers. These queries started in December 2020.

Subsequent conversation in the OARC forum pointed to a likely explanation in the use of "nonce prefixes". (https://developers.google.com/speed/public-dns/docs/security?hl=en#nonce_prefixes). The idea is that off-path potential attackers would find if harder to inject a false DNS response to such a query. As the Google note suggests: "It should therefore be safe to attach a random label to a query name to increase the entropy of the request, while not risking a failure to resolve a non-existent name". As long as a server is not authoritative for both a TLS and the second level domains (such as would be the case in a wildcard delegation) then this technique is effective in adding entropy to queries without altering the responses from the authoritative server.

### Edwards Curve Cryptography and DNSSEC Validation

I presented on the question: Is the DNS ready for Edwards Curve cryptography in DNSSEC? And the answer is currently a clear "No!" Too many validating resolvers don't support this algorithm. See this article (https://www.potaroo.net/ispcol/2021-06/eddi.html) for the measurement details.

### Anycast

The use of anycast in the Internet, while very much commonplace these days, is far from uncontroversial.

The initial debates on the use of anycast centred around both the purported instability of TCP sessions to any anycast server in the face of routing fluctuations and the potential for misdirected ICMP messages, but these concerns, while technically relevant, have largely been discounted in the light of operational experience. The larger concern is that by pushing the optimal path selection role to the routing system, and to the AS path length in particular, the resultant system can be highly inefficient. Short AS paths in BGP do not necessarily mean a short network path in terms of distance, delay, nor do they imply a higher capacity path. This mixed attitude to any anycast has been reflected in the literature in the topic, including a paper in SIGCOMM '18 that noted that "'While it is not surprising that IP anycast is suboptimal … we find [anycast's] inefficiencies to be surprisingly excessive." while just three years previously at SIGCOMM '15 a paper noted that "For most clients, anycast performs well despite the lack of centralized control." The OARC presentation by Colombia's Tom Koch is a continuation of extensive measurement work by a group at Colombia and USC/ISI in measuring the effectiveness of anycast in the DNS.

These days there are many large-scale users of anycast including CDNs from Microsoft, Cloudflare the Verizon. In the DNS anycast is common at the root and at lower levels of the name hierarchy. The 13 root servers all use anycast to straddle over 1,400 unique locations. The question posed in this work is to understand the efficiency of anycast. In this case they are trying to understand to what extent the selected path in BGP

correlates with the shortest possible path from a client to any of the constellation of root servers in terms of "route kilometers" from client to server, and at the same time to understand the extent to which the selected path represents the lowest possible latency. Now, as distance adds latency one would reasonably expect these two metrics to correlate to a high degree, but there are situations where other factors in the network add delay but not additional distance. Neither of these metrics are explicit metrics used in BGP path selection, and while there is some intuitive feel that longer AS paths imply both a longer path in route miles and a slower path in terms of delay, it is not necessarily the case in every situation.

Figure 1 shows the measured outcomes of the root server system in 2018 measuring each server instance against latency to reach the selected instance and efficiency in terms of using the closest instance.

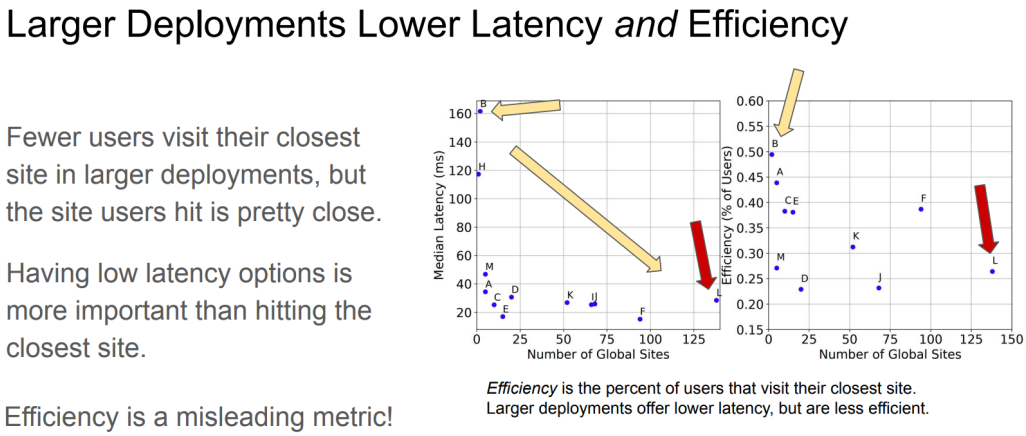


*Figure 1 – DNS Root System Anycast Deployment*
*from* https://indico.dns-oarc.net/event/39/contributions/866/attachments/824/1487/v2%20OARConline%20Workshop%2035a%20Submission.pdf

Why do we deploy anycast? Well lower latency, or a faster experience is a driving factor for many. But at the same time anycast can localise localised attacks and allow more servers to absorb an attack from a widely distributed set of sources, so resilience is another motivating factor. Anycast can also offer service resilience, where additional service elements can be added or revamped from the anycast constellation without the need for coordination - the routing system performs the joins and removals.

I found a useful observation from this presentation, namely that component performance is not representative of overall system performance. In such cases a system may perform as the sum of its components, but equally there are situations where the system may perform according to the best component performance, or in other cases at the level of the worst component performance. Anycast tends to perform well in most cases, as the best path in a routing sense has a reasonable level of correlation with minimising distance and delay, but it is at best an approximate rough level of correlation. The true advantage of anycast is that it takes an otherwise onerous and high overhead task of making regular periodic measurements and then reassessing the local selection and replacing it with a simple outsourcing of the entire process to the routing system!

## Automated Bootstrapping of DNSSEC

DNSSEC is not exactly a runaway success. The level of signing of domain names appears to be somewhere in the order of 5% - 8%, and the number of end users who sit behind DNSSEC validating recursive resolvers is around 30%. Now signing a zone can be challenging, particularly in the case of large dynamically generated zones, but at the same time coordinating the passing of the DS record to the parent zone can present challenges.

In our complicated world of the end client, their DNS provider, the contracted registrant, its chosen registrar and the zone operated by the registry, there is a lot of handling and forwarding requests and too many opportunities for the framework to be deliberately abused. The work behind the CDS/CDNSKEY records can help a lot here, but these mechanisms rely on having a trusted delegation in the first place, and the use of "old-signs-new" means that any key compromise can end badly.

Are there other ways of doing this that avoids the bootstrap vulnerabilities. deSEC's Peter Thomassen explored other approaches that leverage the trusted relationship between the registrar and registry, and instead of placing the CDS record in the delegated domain, it places the record in the registrar's zone, which is already a secured delegation.

It's a neat idea, but somehow I don't see that adding further levels of indirection is going to help here! The issue lies in the fact that the provisioning model in this devolved DNS landscape was devised in a world without DNSSEC and retrofitting the DS/DNSKEY linkage into this model was always going to be a problem! Some registrars have adjusted their tools and interfaces to deal with DS records. Some are gearing up to automate the process with CDS scanning.

One of the weaknesses here is that the NS record is unsigned, as DNSSEC protects only the integrity of the response. It does not protect the delegation chain from root to target domain. There is a conversation in the IETF's DNSOP Working Group to combine some combination of the DS and NS records into a signed record in the parent, protecting both the delegation and the interlocking of DNSSEC keys.

## Was that a DDOS or was it just the DNS again?

As an administrator of a DNS resolver or server, you may be confronted by a situation of a rapid escalation in the query volume (Figure 2)


*Figure 2 – Abrupt changes in DNS query Volumes*
*from* https://indico.dns-oarc.net/event/39/contributions/869/attachments/826/1486/DNS%20DDoS__Challenges%20and%20Mitigations.pdf

These days we are all very keen to blame such anomalies on a DDOS attack, as there are a lot of such attack scripts out there and they all appear to be used indiscriminately and frequently. But sometimes it's not a DDOS attack. Sometimes it's the DNS itself that is the problem. (For example, in our experimental measurement setup at APNIC we've observed at one point a single DNS "question" from an end-client stub resolver managed to generate 292,000 identical queries to the corresponding authoritative server in the ensuing 30 seconds!)

The issue is that the elements of DNS infrastructure can be both persistent and enthusiastic in their approach to resolving a DNS name, and the use of large scale server "farms" to create high capacity DNS service platforms can sometimes act as a massive amplifier and the resultant query explosion can bear all the hallmarks of a DDOS attack, yet it's not!

As Damian Menscher pointed out in this presentation on this topic: "Don't assume DNS outages are caused by DDoS even if you see an increase in traffic Blocking the wrong queries can make the situation a lot worse [...] Consider the system as a whole, including recursives: TTL and ECS have a huge impact on the effective capacity of an authoritative server."

## Disclaimer

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

## Author

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

*www.potaroo.net*