The ISP Column
A column on things Internet

May 2026
Geoff Huston

Encrypted DNS at OARC 46

OARC, the organisation for DNS operations and research met for two days in Edinburgh in May 2026. The topic of using encrypted DNS transports for communication between recursive resolvers and authoritative services in the DNS was discussed at this meeting, and I'd like to share my impressions of this work.

Much has been said about the use of the DNS as a means both of tracking the online behaviour of individual users and as a means of online censorship and control. Almost every online transaction starts with a DNS query, and if one were able to assemble the complete set of DNS queries generated by an individual user it would be possible to assemble a relatively complete profile of their online activity.

For many years this aspect of the DNS, as a means of observation into the activities of others, received little attention. The sensitivities over state and private digital mass surveillance were raised to a new level of public prominence by the material in the Snowden papers a little over a decade ago. This episode brought significant attention to the overall topic of DNS privacy, as the DNS name resolution protocol has been sadly lacking in some basic protections. By default, DNS queries and responses are in the clear, which makes the DNS prone to hostile man-in-the-middle manipulation as well as the noted issues related to eavesdropping providing copious information to onlookers.

Stub to Recursive DNS

There have been some responses to try and mitigate these DNS privacy issues. One is to make eavesdropping and manipulation harder by using encryption for DNS transactions. The issues around encryption and the efforts with DNS over encrypted transport were a topic of interest in the DNS world, and the DNS Privacy Working Group of the IETF (DRPIVE) generated a number of specifications that introduced channel encryption and remote server authentication through DNS over TLS (DoT), DNS over QUIC (DoQ) and DNS over HTTPS (DoH), collectively referred to as "DoX".

Compared to a simple UDP transaction the overheads of setting up a TCP session and then establishing an encrypted session context are considerable. Even if you combine these two tasks with QUIC, its still a considerable overhead. And this would be an unsustainable overhead were it not for one saving factor: the stub resolver is able to reuse this session for all of its queries to the recursive resolver. This means that the cost of the initial session establishment can be amortised over all subsequent queries. What remains is the trade-off between encryption and decryption on the one hand and the avoidance of UDP fragmentation, and DNS truncation and TCP followup on the other.

These DoX tools apply to the connection between the "stub" resolver operated in the user's environment and the recursive resolver. It can be applied when the recursive resolver is operated by the user's ISP (through the net change in the security stance is minimal, as the ISP is already in a privileged position with respect to the user's traffic), but it makes much more sense when the user chooses to use a remote open resolver, where the user's DNS traffic is passed across open Internet segments. It's perhaps unsurprising that when you look at the traffic profile of a large open DNS recursive resolver that the use

of DoH and DoT is significant, at 15% and 10% of the query load as seen by Cloudflare's 1.1.1.1 service (https://stats.labs.apnic.net/edns/XA). While I do not have access to data from large scale retail ISP's recursive resolvers, I suspect that the DoX query levels are much smaller in that localised context.

Recursive to Authoritative DNS

However, this is just one part of the DNS privacy story, and the second part is the DNS traffic between the recursive resolver and the DNS' authoritative servers.

One could argue that this second part is less of an individual privacy concern, as the end user's identity (in the form of the user's current IP address) in not carried in these queries. However, there is still some information exposure, both in the queries themselves and in the possibility of response interception or response substitution. It's still a privacy and resilience problem, but different to the stub-to-recursive step. If an attacker can successfully alter a response from an authoritative server, it will be loaded into the recursive resolver's cache and then will be used to answer subsequent stub queries for the ensuing cache lifetime (commonly termed "cache poisoning").

An approach to reducing the information exposure risk in the recursive-to-authoritative path is to perform DNS Query Name Minimisation (RFC 9516, "DNS Query Name Minimisation to Improve Privacy", November 2021). Here, during the process of nameserver discovery, each authoritative nameserver is presented with a query name that is stripped to just one label more than the longest matching domain name for which the name server is known to be authoritative, and uses a QTYPE selected by the resolver to possibly obscure the original query's QTYPE. This has proved to be quite popular in DNS recursive resolvers, and a measurement study performed in 2020 indicated that some 18% of the Internet user population passed their queries through recursive resolvers that performed Query Name Minimisation (https://www.potaroo.net/ispcol/2020-09/qmin.html).

Another approach is to use the same approach as DoT and DoQ (RFC 9539, "Unilateral Opportunistic Deployment of Encrypted Recursive-to-Authoritative DNS", February 2024) (DoH is not included here as the existing NS delegation information does not convey the path part of a DoH endpoint URL). This opportunistic form of ADoX (Authoritative DNS over encrypted transport) uses encryption without authentication (RFC 7435, "Opportunistic Security: Some Protection Most of the Time", December 2014). A recent study, presented at OARC 46, looked at the support of ADoX in authoritative nameservers. This study found that the Authoritative Servers for 3,074,281 (0.93%) registered domains support some form of ADoX, and if you look at the IP addresses of these nameservers, then 2,585 (0.32%) uniquely IP-addresses nameservers support ADoX. A comparable study of resolver behaviour using Atlas probes found no recursive resolvers used ADoQ and just 44 probes were behind recursive resolvers that used ADoT, mostly operated by the Quad9 open resolver service. To all practical purposes deployment of ADoX has stalled. It has become somewhat of an impasse. Authoritative services see little benefit in supporting ADoX is no recursive resolver will use it, and the incremental support costs on the server are far higher than stateless UDP transactions. Recursive resolvers see no point in supporting ADoX queries if no Authoritative servers support it as well.

This opportunistic approach to ADoX is not without its downside. It takes a number of round-trip intervals to attempt a TLS/TCP or QUIC connection with the authoritative nameserver, and if this is with the intention of performing a single query then the encrypted transport overhead is considerable, as compared to a simple UDP transaction. Unlike encryption on the stub-to-resolver path, there is no certainty of session reuse, so while subsequent queries can be almost as efficient as UDP, there is no assurance that there will be subsequent query to this authoritative nameserver from the recursive resolver. While the channel is encrypted, the server is not authenticated in this opportunistic encryption framework. It would therefore be prudent to also perform DNSSEC validation for signed domains to provide assurance of the authenticity of the response. The upside of this approach is that the use of a ADoX transport means that large responses can be handled, and experimental approaches to DNSSEC Validation such as CHAIN Query Requests (RFC 7901, "CHAIN Query Requests in DNS", June 2016).

Another approach is to signal the capabilities of the nameserver in the same DNS resource record as the record that denotes a label as a delegation point. This is essentially the DELEG record, which builds upon the SVCB construct that places a number of information fields in a single DNS response. DELEG is not a minor tweak to the NS record. It represents a major change to the DNS architecture, as it alters the DNS framework by loading parent zones with not only the information that the label is a delegation point with the delegated zone's nameservers, but it also may add the nameserver's cryptographic keys, IP addresses, and protocol capabilities into this single record. It is also part of the parent zone, which implies that for DNSSEC-signed zones these DELEG records are signed by the parent's signing key. In this context the role of DELEG is very straightforward: to communicate the ADoX capabilities of a delegated zone's servers to a recursive resolver, and do so in a "strict" manner that if resolver does not the matching ADoX channel then no resolution may be performed. In DELEG there is no need for opportunistic probing. The capabilities of the authoritative servers for the delegated domain are enumerated in the DELEG record.

The DNSSEC-signed DELEG record raises an interesting question about the role and purpose of DNSSEC validation in the DNS. The role of DNSSEC validation in the DNS is to provide assurance about the DNS response, and not to provide assurance about the resolver's traversal within the DNS name space from parent to child when resolving a name, if such a traversal has even occurred (the response may have been generated from the resolver's cache). Unlike the parent-size NS record, the DELEG record is DNSSEC-signed in the parent zone, but the DNSSEC model does not necessarily make resolution any more robust if each DELEG delegation is independently validated before use.

Is this ADoX work little more than an academic exercise? Shifting the entire DNS from a set of lightweight UDP transactions to a sequence of encrypted sessions is not a step to be taken lightly, and dire warnings about "catastrophic consequences" were parts of the discussion on this topic at the OARC workshop.

Right now, most of this effort, with the specifications of opportunistic ADoX and DELEG, is indeed a paper exercise. It's still the case that the Internet's framework to protect the authenticity of online transactions is based on X.509 domain name certificates, and this approach circumvents any reliance on assurance of the DNS name resolution outcomes. It is certainly feasible to inject false information into the DNS in many ways, and the consequence is that end users may be misled into contacting a false service endpoint. Channel encryption in ADoX does make some potential attack vectors far harder, but other attack vectors remain, and the overall problem of assurance in the DNS resolution process remains.

The constraint here is that if the application service is using TLS (such as HTTPS), then the server is forced to demonstrate its knowledge of the private key associated with the domain name, and if it cannot do so then the transaction should be abandoned. Does ADoX or DELEG alter this picture in any way? Absolutely not! In a weaker formulation of the same question, could ADoX be used as a substitute for DNSSEC-signing of resource records? Again, the answer is absolutely not! ADoX is about how a recursive resolver resolves a name, whereas DNSSEC is about how a user can be assured of the authenticity of the result, irrespective of how it was obtained! They are complementary approaches, and mutually substitutable.

It's possible to adopt a security purist stance, applying strict controls over every step in the process of resolving a DNS name, validating the authenticity of signed DELEG records, and assiduously opening up an encrypted and end-point validated transport session with each authoritative nameserver, and DNSSEC-validating the result of the resolution operation. But the pragmatic question is whether service providers, and ultimately users, are willing to pay the incremental time and transaction cost penalties associated with the application of such stringent controls. So far what we observe is that fast and low-cost options are heavily favoured by users and service providers alike. Rather than attempting to validate every step in the DNS name resolution process we favour a fast and cheap DNS resolution process in the clear (with its attendant risks) and push the authentication function back to the application layer in the form of the TLS handshake. The DNS provider and DNS resolution operators bear none of the incremental costs of the function of assurance of authenticity. Attempting to push the cost of assurance

back to the name resolution infrastructure seems to me to be a regressive step that will strike considerable resistance from the name infrastructure operators. To put it in different words, the DPRIVE agenda stopped at the hurdle of recursive-to-authoritative queries simply because there was insufficient interest in running with this infrastructure model in the public Internet.

Is this work going anywhere useful?

What's the future for encrypting the recursive-to-authoritative hop in the DNS? I'm not sure that I can point to any assured future here for this work. The underlying issue is that the economics of this space are working against DNS infrastructure operators. Having to absorb the additional costs of running encrypted sessions here in the absence of any prospect of incremental revenue is not a sensible position for any enterprise. Incurring cost without any matching revenue stream is never a sensible business decision!

For the user, the DNS remains largely free of incremental query cost when using the DNS. We've been able to preserve this property by ensuring that service authenticity assurance is largely provided by TLS as an application service, rather than as an attribute of DNS resolution. I'm not sure that this picture will change in the forseeable future.

The presentations at OARC 46 can be found at:

https://indico.dns-oarc.net/event/56/timetable/#20260516.detailed

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