Council of European National Top-Level Domain Registries

Report on IETF109

Virtual Meeting
16-20 November 2020

CENTR vzw/asbl · Belliardstraat 20 (6th floor) · 1040 Brussels, Belgium
Phone: +32 2 627 5550 · Fax: +32 2 627 5559 · secretariat@centr.org · www.centr.org

Contents

I. Introduction

The final meeting of the Internet Engineering Task Force (IETF) this year was yet again a virtual one, making 2020 the first year the standards body has had only online meetings. With the IETF Network Operations Center still finding its feet with meetings at this scale, technical snags were to be expected, but hardly took away from the full agenda of working and research groups that saw a total of more than a thousand participants.

II. Adapting to the reality of encrypted DNS deployment

What do computer scientists, behaviour economists and cognitive psychologists have in common? They all appreciate the power of the default effect, i.e. whatever people get without making an active choice is what is likely to be the most popular. In the world of network protocol development, the story of deployment of encrypted DNS protocols is arguably centered around what will become the default.

With traditional clear-text DNS still being the most common, the future of the default choice of encrypted DNS is still up for grabs. The Adaptive DNS Discovery (ADD) working group now has a variety of proposals from internet service providers, cloud service providers and web browsers.

Background

The Domain Name System (DNS) is the way in which human-readable names (like centr.org) are converted to their network address (e.g. 178.208.52.35 or 2a00:1 c98:10:60:ffff:ffff:ffff:10) so you can connect to them. Notably, such queries have traditionally happened over plain text and therefore lacked security and privacy guarantees. Internet service providers, which have traditionally provided these services to users, can see what websites one is visiting. On-path attackers could also easily see this information, and even block certain websites based on it.

The possibility of more privacy in these queries finally opened up with the standardisation of protocols like DNS over TLS (DoT) and DNS over HTTPS (DoT) in 2016 and 2018 respectively. While there was consensus that these protocols increase on-path privacy, a matter of concern with them still remains: who does finally get to see these queries? Internet service providers (ISPs) were concerned that applications could easily run DoH queries to whatever resolvers they like, effectively bypassing them. Such private information would now be available to big tech companies operating browsers or cloud services, which have been involved in the development and deployment of DoH.

The Internet Services Providers Association in the UK even nominated Mozilla as an 'Internet Villain' for planning to roll out DoH in a way that bypassed them and their content filtering mechanisms. The European Telecommunications Network Operators' Association published a position paper noting their concerns for how all DNS traffic may move to a small number of players, and called for more scrutiny of the impact of DoH deployment on regulation and competition in the industry.

New developments at the IETF may have significant policy consequences, given regulators in the EU and around the world becoming increasingly sensitive to both privacy and competition law concerns in the tech industry.

What's happening now?

While Mozilla made DoH the default for users in the US, the fervent backlash that caught the eye of UK regulators meant that they stopped their plans to do the same in the UK. Several developments at the IETF provide an indication for what may happen with how DoH and DoT are rolled out increasingly around the globe.

Instead of going directly to third-party DNS resolvers, there may be two reasons for sticking to ISPs' resolvers (now with DoH/DoT instead of plaintext DNS). First, ISPs can continue to provide parental controls or other filtering services if customers have opted for (or are involuntarily subject to) them. Second, the relationships that ISPs have with local cloud providers may mean that they provide better responses, i.e. the network addresses they provide in response to DNS queries may be closer, and thus such responses can result in more efficient traffic routing.

Earlier this year, the Adaptive DNS Discovery (ADD) working group was set up at the IETF to explore some related questions: How can a user or device discover DNS resolvers that are available to them in their network? How can a user select one if multiple resolvers are available?

Just add a DHCP option!

Traditionally, your device picks a DNS resolver that your access point tells it to using the Dynamic Host Configuration Protocol (DHCP). The access point itself retrieves these details from your ISP. One way then to implement a way for your ISP to instruct your device to use their DoH/DoT resolver is to have a way in DHCP to do that, which is exactly what a group of engineers have proposed with the Internet Draft DHCP and Router Advertisement Options for Encrypted DNS Discovery within Home Networks.

Discovering 'equivalent' resolvers

On the agenda for the Adaptive DNS Discovery working group at IETF109, however, was Discovery of Equivalent Encrypted Resolvers, which approaches the matter differently. Developed by technologists at Apple, Microsoft, Cloudflare and Fastly, the proposal seeks to answer the specific question of what a device can do once it does have a traditional DNS server that it seemingly trusts: how can it discover an equivalent service that uses DoH/DoT instead? In the usual case, the Internet Draft proposes that each device performs an additional DNS query (that uses the service binding and parameter records, being developed separately at the IETF) when it finds out an unencrypted resolver exists: the response to this query will contain information on how to contact related resolvers that support encrypted DNS protocols.

Of course, it would be uncharacteristic of IETF participants to leave potential for pedantry untapped. For around two hours at the IETF109, the discussion focused on what 'equivalent' could mean.

Other rooms, other wonders

A related Internet Draft comes in the context of Mozilla enlisting US telecom giant Comcast in their trusted resolver program. Their Internet Draft, CNAME Discovery of Local DoH Resolvers, proposes that a name 'doh.test' be reserved for a CNAME DNS query for discovering DoH resolvers. An application (like Mozilla's browser Firefox) can perform this query with traditional plaintext DNS: if it receives a response with a resolver that exists in the trusted resolver program, the application will use it instead of using the default (which, for the Firefox is currently Cloudflare in the US).

A shifting mood

If the initial conversations on DoH seemed indifferent about the role of ISPs, the current phase of discussion centres around their involvement (or at least deployment not without their involvement). Two things are becoming increasingly clear however. First, that encrypted DNS is here to stay. Second, with all these proposals moving at the IETF, DoH/DoT deployment globally may be more conservative than originally anticipated: it has not, at least immediately, concentrated power in the hands of web browsers. Simply put, internet service providers may still continue to play an important role in providing DNS services to their users.

III. Standardising an end-to-end encrypted messaging protocol at the IETF

Last month, an Austrian media report kicked up a storm by suggesting that the Council of the European Union was drafting a resolution to prohibit the use of end-to-end encrypted communication. This was quickly corrected: the draft resolution, in fact, affirms the previous position of previous EU policy documents that recognise the importance of end-to-end encryption (E2EE) in providing secure and private communication.

While the European Commission has been considering questions around E2EE and information access to law enforcement agencies since 2016, there have emerged no serious and binding proposals that threaten popular use of E2EE communication. A couple of developments, however, portend some uncertainty about how strongly this position will be held in the future.

Earlier this year, Politico leaked documents that revealed deliberations of a working group of the European Commission on 'technical solutions' for detecting child sexual abuse material in private E2EE communications, such as those provided by Signal and WhatsApp. Civil society organisations fear that these proposals, which include client-side scanning of content and "exceptional access" to encrypted data, undermine the security and privacy guarantees that E2EE messaging provides.

The second threat to E2EE communication comes from counter-terrorism efforts in the EU. While the latest

draft of the proposal on regulation of the dissemination of terrorist content online does not apply to private messaging services, the EU Counter-Terrorism Coordinator has been advancing a different position. In May 2020, they wrote to EU Member States advocating for an encryption "front-door" and increased state intervention in regulating encryption. In October, when the Five Eyes (the US, UK, Australia, Canada and New Zealand), India and Japan issued a joint statement calling for cleartext contents of communication to be available to law enforcement agencies on demand, the EU Counter-Terrorism Coordinator welcomed the proposal.

The policy position of E2EE at the EU-level is thus becoming somewhat polyvocal and/or stuck at a question that has no real answer: when device access is not possible, how can law enforcement agencies access end-to-end encrypted messages without 'breaking' said forms of encryption? Unfortunately, such policy aspirations may as well be in "a laundry list of tortuous ways to achieve the impossible."

Pertinently, one of the ways the EU Counter-Terrorism Coordinator has recommended is to monitor standards development. In their (translated) words:

"Member states and EU Institutions should be encouraged to collectively challenge changes to the encryption landscape in the international standards bodies, particularly the Internet Engineering Task Force (IETF), to ensure they are involved in the development of international standards and technological norms, impacting encryption and wider cyber security for the years to come."

So, what is happening at the IETF?

The Messaging Layer Security (MLS) working group is unperturbed by these policy debates on end-to-end encryption. Set up in 2018, the working group has a clear objective to standardise an architecture and protocol that can facilitate end-to-end encrypted messaging. MLS will have several key security properties, including:

- *Message confidentiality*: messages cannot be read by anyone except the sender and recipient(s)
- *Message integrity*: messages cannot be tampered with
- *Message authenticity*: recipients have an assurance of the sender's identity
- *Forward secrecy*: compromise of a key at an endpoint does not cause all previous communications to be immediately decryptable
- *Post-compromise security*: compromise of a key at an endpoint does not cause all future messages to be revealed, i.e. there a way to recover security properties even after a compromise

All the properties listed here are already guaranteed by some existing solutions, such as the Signal protocol, a version of which WhatsApp also uses. What is new about MLS is its design philosophy: it starts with group messaging as a default, whereas older protocols are designed for one-to-one communication. The intention is for MLS to be much more scalable than current solutions (like Signal, iMessage, WhatsApp, etc.). This performance edge and the open nature of the standard is likely to be incentive enough for lots of platforms and services to adopt MLS as their message encryption protocol of choice. That is why, besides academicians, the working group has active participation from companies, including Google, Mozilla, Facebook, Twitter and Wire.

Will it federate?

Since traditional E2EE protocols were designed keeping one-to-one conversations in mind, the logic of how chat 'groups' operate has been left to individual services and platforms. Coupled with the fact that some organisations may deliberately not want to federate their service (for commercial or non-commercial reasons), true interoperability on a public scale has arguably never been achieved with E2EE messaging.

MLS has the potential to change that. While the working group has not set complete federation/interoperability as an explicit goal, an Internet Draft by authors from Google and Wire clearly lays out that it is technically possible with the existing architecture of MLS. If successfully demonstrated, it is likely that details on how to achieve federation with MLS are incorporated into the proposal in early 2021.

The ecosystem is still moving

Fortunately, the MLS working group is concerned with usability as much as it is with security and privacy. With working group participants actively having accommodated support for multiple devices per user in addition to business use-cases, MLS offers the

promise of a protocol that can be widely deployed across all applications that need a messaging feature.

As the charter for MLS notes, the working group "hope[s] to have several interoperable implementations as well as a thorough security analysis" before standardisation. This was confirmed at IETF109, where the plan for the protocol specification was discussed. The Internet Draft will go on a freeze until developers can get deployment experience with the current version, and academicians can formally analyse the cryptographic properties.

With broad industry buy-in and the likelihood of open source implementations cropping up in the near future, the MLS open standard may just become the backbone of private communication online.

IV. Is this a privacy concern? Reverse search in registry data

Since 2018, The entry into effect of the General Data Protection Regulation (GDPR) has reinvigorated privacy concerns associated with registry data in the EU. The question of whether the traditionally-public nature of registry data is in conflict with European data protection requirements had already sparked conversations at ICANN. If the recent proposals in the Registration Protocol Extensions (regext) working group are any indication, regulatory developments at the EU continue providing context to and/or affecting standard-setting on registry data at the IETF as well.

Since January 2019, the regext working group has adopted a specification that outlines how to add 'reverse search' capabilities to the Registration Data Access Protocol (RDAP). This feature gets its name from the many websites that use public information to provide 'Reverse Whois' capabilities, i.e. they allow anyone to find out what domain names are registered by a particular person (or a particular email address).

Why reverse search at all?

Mario Loffredo, who works at Registro.it and is one of the co-authors of the Internet Draft, presented the proposal at the meeting of the regext group at IETF109. Notably, a significant part of Loffredo's brief presentation focused on regulatory context in the EU that may speak of requiring this sort of capability. Amongst other things, Loffredo cited the European Commission's proposed regulations, E-evidence - cross-border access to electronic evidence, that seek to establish clear principles for law enforcement access to information held by service providers. The matter was also discussed in the 63rd CENTR Legal & Regulatory Workshop.

Apart from a statement in the Draft that says how the feature may allow for "registrars searching for their own domains", the primary motivation of standardising the feature seems to be easing access to information to law enforcement agencies.

Note that the presentation further said that "[a] uthorities should be able to access unpublic [sic] registry data without submitting written requests", a statement which this author could not substantiate or reconcile with the EC's proposed regulation which specifically speaks of judicial orders for information access.

Another consideration for standardisation at the IETF, of course, flows from the standards body's ethos of rough consensus and running code. For the latter, working groups generally prefer to record demonstrable interest in deploying the technical proposal before it is standardised, particularly if multiple implementations do not exist in the wild already. Currently, the Internet Draft only lists the Italian registry as having implemented this feature as a proof of concept.

Dealing with privacy concerns

According to the authors of the Internet Draft, the privacy concerns that apply to Reverse Whois are largely absent from their proposal because RDAP can allow for authentication before data access. Yet, most of the discussion of the Internet Draft now is centered around how to deal with the privacy considerations of denoting such a feature as a standard. Some participants believe that appropriate technical and organisational controls can entirely mitigate the privacy risks: when registries implement the feature, they should have strict control over who can run these 'reverse search' queries, and authenticate their identity each time.

The Internet Draft does have a section on privacy considerations, but it is brief and largely asks registries to follow legal procedures. Alexander Mayrhofer, working at the Austrian registry, pointed out the absurdity of the text considering that "there is no need to say, in a technical document that you 'must follow the law', because that's quite obvious."

Ulrich Wisser, who works at the Swedish national domain registry, added, "How do we know that [the considerations in the Draft are] good privacy consideration[s]?" While the IETF process mandates a section on security considerations for network protocols and standards, there is no similar requirement for listing down privacy risks and their associated mitigation. There is some guidance on this aspect in RFC 6973, Privacy Considerations for Internet Protocols, but it is not clear whether the Draft authors have considered those.

With no other apparent hurdle for the Internet Draft to proceed to the next stages, the forthcoming discussion on the proposal may tell us how the IETF working group will debate privacy concerns, when a clear and primary motivation of a proposed standard is law enforcement access to information.

(Disclosure: The author of this report has previously commented, in their personal capacity, on older versions of the Internet Draft.)

V. Wrench and numbers: Is the DNS centralized?

"Is Internet traffic consolidating, i.e., moving towards a larger fraction of traffic involving a small set of large content providers, social networks, and content delivery platforms? It certainly appears so, though more research on this topic would be welcome."

> -- The Internet Architecture Board on Consolidation in March 2018

Ask and you shall receive… Well, or as academicians may say: ask for research, and two years later, you may be fortunate enough to receive some initial evidence that potentially answers your question.

In 2018, when the deployment of encrypted protocols had sparked concerns around consolidation of DNS queries in the hands of a few large private companies, there was little evidence to show concentrated the market already was. Over the last two years, there has been mounting evidence that favours this hypothesis. In the meeting of the Measurement and Analysis for Protocols Research Group (MAPRG) at IETF109, Sebastian Castro presented such a paper, Clouding up the Internet: how centralized is DNS traffic becoming?, which was published in the proceedings of the ACM Internet Measurement Conference (IMC) 2020.

The authors' approach relies on analysing DNS traffic flowing from resolvers to three authoritative servers: one in Netherlands (.nl), New Zealand (.nz), and B-ROOT (multiple top-level domains). The final dataset looks at queries for a single week across three years, ending up with information on 55 billion DNS queries. Then they identify traffic coming from the big five companies (Google, Amazon, Microsoft, Facebook and Cloudflare) involved in providing hosting services.

Compared to .nz and B-ROOT, the .nl authoritative server saw the most concentration of received traffic: more than a third of the traffic came from just these five companies, with Google leading the pack. Google's large share of the traffic can be partly explained by the fact that only Google and Cloudflare, out of the five, operate public DNS resolvers. The authors also identified queries from public DNS resolvers were also the majority in the dataset.

At the same time, this data alone may not be enough to capture concentration within the DNS market. For instance, it does not answer whether there are a comparable number of internet service providers (ISPs) that are directly or indirectly responsible for similar levels of traffic to authoritative servers. However, given the fact these five companies are involved in providing a host of other services, the results of this paper indicate worrying levels of consolidation in the internet economy at large.

Also keep in mind that the authors' approach relies on measuring resolver-to-authoritative server traffic, i.e. it is not representative of how consolidated the market is on the user end. Considering the fact that Google or Cloudflare resolvers are caching responses (and serving those to users without contacting the authoritative resolvers every single time), the user-to-traffic DNS traffic may be even more concentrated than the paper's findings.

That would be in line with the results of Roxana Radu and Michael Hausding's paper, Consolidation in the DNS resolver market – how much, how fast, how dangerous?, published in the Journal of Cyber Policy in February. From an analysis of 100,000 measurements from the Open Observatory for Network Interference (OONI) database, they conclude that "there is a high concentration of power in the hands of Google and Cloudflare, which control half of the overall market."

Apart from being sour news for those seeking more competition in their digital markets, the security and

privacy ramifications of consolidation in the DNS market are also significant. Large DNS providers can be singular points of failure, as evidenced by the denial of service attack mounted on Dyn in 2016, which led to the unavailability of several prominent services across Europe and North America. The sensitive nature of DNS queries can also be exploited by companies for commercial advantage, either by selling datasets entirely, or to aid their micro-targeting advertising services.

The deployment of encrypted DNS protocols, like DNS over TLS (DoT) and DNS over HTTPS (DoT) is likely to entrench this trend, considering that Cloudflare and Google are influential players in pushing those protocols to end users. While regulators across the world are quickly catching up on competition concerns in the internet economy, this recent evidence is a clarion call for policymakers to pay more attention again to market consolidation in the 'invisible' parts of our networks. If nothing else, they can always ask for more research.

VI. Transparent censors and other extensions of extended error codes

The DNS Working Group of the IETF is continuing to expand the DNS code base with both new features and enhancements to previous features. In the latest session, a proposal on private space in the DNS with two letter codes received mixed comments, while the policy-heavy work on the operational fall-out of DoH is still not welcome.

All those failed DNS queries

Under the current technical specifications for the DNS, receiving an error message in response to a DNS query can mean any number of things. The new RFC 8914 on Extended Error Codes proposes to change this, so that administrators will at least be able to know the specifics of an error.

Among the different problems a query could run into, and that an administrator might need to be aware of in order to take the right countermeasures, are issues with DNSSEC certificates (such as expired certificates, signatures that are not yet valid or even unsupported crypto algorithms), network problems or upstream issues with the authoritative servers of the domain. Queries can also fail due to policy reasons, for instance if a resolver or authoritative server is based in a jurisdiction which places blocking, filtering or prohibition requirements on the resolution of queries. The list of 26 error codes in RFC 8914 carefully differentiates between these cases.

But no sooner was the electronic ink dry for RFC 8914 than a group of editors from McAfee, Open-Xchange, Citrix and Orange asked for additional transparency with regard to the "filtering and blocking category".

Under the current error code list, users do not know why a domain was filtered or blocked, Tirumaleswar Reddy explained during the DNSOP session at IETF 109. Reddy and his co-authors propose an extended DNS (EDNS(0)) option that would return a Uniform Resource Identifier (URI) that explains the reason a DNS query was filtered. Foreseen benefits include an ability for end-users to send timely objections to responsible parties when content that should be available is made unavailable.

However, the proposed solution comes with considerable security issues, notably the malignant injection of an error page by an attacker. Reddy, whose draft already identifies this issue, promised the draft would try solving this by making DNS encryption mandatory and also by forcing a rejection of any displayed URI EDNS(0) options that are provided by unauthenticated origins.

With such limitations, the implementation of transparent filter messages risks becoming pretty restricted, US academic Wes Hardaker noted, and he recommended waiting to see whether the Extended Error Codes from RFC 8914 would see a greater adoption before taking any further steps. He also proposed that a free text field – while limited in length – could be used to signal the URI of an explanatory error page in the meantime.

Signaling errors into the other direction

Two ICANN employees, Roy Arends and Matt Larson, would also like to signal errors towards the authoritative name server experiencing the problem.

A reporting agent for the authoritative domain, specified in the EDNS(0) option received from the authoritative server, could receive indications of the error-related queries from the recursive resolvers,

Arends proposed.

The proposal raises similar security concerns as the one by Reddy et al, but Arends seems intent on going ahead. After discussions in the DNSOP working group, he noted that the IETF document is currently listed as an independent submission which the working group would not need to adopt.

Private zones by name and not only by number

As previously reported in the CENTR Tech Trends Watch Q2/2020, Arends has also proposed – together with Joe Ably – the creation of an IETF-managed list of two-letter private namespaces following existing two-letter codes in ISO 3166-1. This proposal has generated heavy e-mailing list traffic since Q1, and now made it to the DNSOP meeting. Working group members such as Ted Hardie, former IAB Chair, warned that this issue had to be discussed between ICANN and ISO.

Fight over – new edition in DNSOP?

With the DoH WG being closed the authors of a draft on guidelines for operators are desperately looking for a new space to place their work. But the DNSOP Chairs certainly want to keep the policy-heavy discussion out of their WG as good as they can. For the ongoing dispute on DoH, discovery and the related privacy issues – stay tuned for the next CENTR blog post.

VII. Diversity at any price? IETF looking for a new chair

The ongoing search for a new IETF Chair offers the community a possibility to look into diversity issues and choose a candidate sponsored by one of the newer participants in the standardisation process. It is unfortunate that the most plausible candidate from the standpoint of diversity, is sponsored by Chinese vendor Huawei, who is currently locked in a trade war with the US.

Huawei already sends more developers to the IETF than the most long-standing participants in internet standardisation. For IETF 109 Huawei and its subsidiary Futurewei together registered 92 attendees, while Cisco, one of the oldest sponsors of the IETF and employer of reigning Chair Alissa Cooper, this time registered a mere 66. According to Cooper's statistics for IETF 109, Chinese companies and universities additionally stepped up to become the second biggest group of participants after the US participant cohort.

Two candidates for the chair position, Barry Leiba and Alvaro Retana, are employed by the research focused Huawei subsidiary Futurewei and a third candidate, UK based consultant Adrian Farell, is known to have cooperated with Huawei on a number of projects. In the run-up, it seems clear that Huawei is seeking to sponsor their first IETF Chair.

Full time positions

The IETF Chair position is a near full-time job. Job tasks include overseeing IETF work in general and the work of the IESG, the peer body of the IETF, in particular. IETF Chairs serve as director of the so called General Area workstream, which is tasked with things like the recent disentangling from the Internet Society. Plus the IETF Chair has to represent the IETF to the outside world, as well as in various internet governance related bodies.

The IETF LLC, the organisation formally charged with running IETF meetings and intersessional infrastructure, does not remunerate the position, so individuals taking on the role have to be supported financially by their employers or industry partners. Historically, one of the more curious sponsorships was certainly the United States National Security Agency's sponsorship of 2007 -2013 Chair Russ Housley.

Anti-Huawei climate

If US public authorities have previously sponsored chairs directly, and the IETF is on the look-out for improving its representational diversity anyway, why should there not be a Huawei-sponsored IETF boss?

For the IETF Nomination Committee, formally responsible for selecting suitable candidates for important positions, the heavy political bias against Huawei and other Chinese vendors in the US and some of its allies is certainly a complication. This bias is on clear display both by trade sanctions and entity listings in the US, as well as in a number of European Union countries. Another illustration of this bias is the so called Clean Network Initaitive from the US State department.

The IETF is not under US regulation, one US observer notes, so no legal issue would arise. But according to this long-time IETF expert, it could be a problem

politically if some "crusading congress critter" wanted to make an issue of it. And despite the State Department changing hands soon, the anti-China hysteria might very well stay around because the incoming president could be expected to tread carefully if only to push back against early "China puppet" screams.

On the list of candidates are:

- Adrian Farrel, Old Dog Consulting
- Alvaro Retana, Futurwei
- Barry Leiba, Futurwei
- Deborah Brungard, AT&T
- Fred Baker, Consultant, Board Member at ISC, and former IETF Chair
- Lars Eggert, NetApp
- Rich Salz, Akamai

VIII. DNS transport: The race is on!

Not one, not two, but three new protocols are offering internet transport layer options for the Domain Name System (DNS). We must not lose sight of the dernier cri (*last shout) though. Here is a quick look at the catalogue of options and opinions on DNS over TLS (DoT), DNS over HTTPS (DoH) and DNS over Quic (DoQ).*

End-destination better security

Securing DNS transport is becoming quite fashionable. Mozilla pressed the pace when announcing its implementation of browser-based DNS over HTTPs (DoH) in the US in 2019. Microsoft, Google and Apple all followed suit to announce implementations, as did network operators like ComCast, which partnered with Mozilla last summer.

There is also no shortage of European implementers of DoH on the network operator side. Both Deutsche Telekom and British Telecom are in on it. According to Nicolas Leymann, the German network operator will offer experimental DoH for its customers in the first quarter of 2021.

The original front-runner for a privacy-friendly solution was DNS over TLS (DoT). It is still seen as the natural evolution to secure infrastructure DNS and leaves the configuration of service parameters to users and network providers. Compared to DoH, DoT suffers from the fact that DoT traffic is easily discernible because it runs under a special port number.

In a recent column about DNS Trends, APNIC Chief Scientist Geoff Hustonalso points to another issue: DoT does not eliminate the potential for the manipulation of DNS answers, but places trust in the hands of the DNS provider of choice. In Huston's words: "all you really know is who is lying to you".

The candidates

Using HTTPS web transport as the substrate, DNS queries benefit from TLS encryption. They also become part of the vast HTTPS traffic flows and cannot be easily identified by networks. Mozilla engineers never tire of underlining these privacy gains for users. Using DoH DNS becomes part of the application, and it allows applications to bypass local and remote networks as well as platforms.

The most recent development is oblivious DoH (ODoH), just promoted by Cloudflare as the ultimate answer to concerns over the concentration of user information. ODoH adds a proxy between the public resolver and end user, separating DNS information from the user's IP.

During IETF 109 Christian Huitema, an expert in privacy by design, further asked the DNS Privacy (DPRRIVE) working group if he could go ahead with secure DNS protocol number three, DNS over Quic (DoQ).

With Quic, the IETF's new transport protocol, on the finish line, DoQ could be pursued in earnest. Quic is UDP-based and integrates the TLS stack to become the first natively privacy preserving transport protocol. Many believe it will become a big competitor to TCP. What could make DoQ attractive for DNS providers is that the encryption is dealt with at the transport level. Plus, the DNS could benefit from additional Quic features like multiplexing.

One to rule them all?

While Huston does not see a big future for DoT and also calls the half-forgotten UDP-based Datagram TLS (DTLS) – the fourth secure DNS transport - too fragile, other experts see a potential division of labour between the candidates.

DNS privacy expert Sara Dickinson from British-based

consultant company Sinodun believes "we will have multiple protocols which have specialised areas".

She can see that DoH is preferred by applications, while DoT makes more sense for basic stub resolvers. For DoQ, which came late to the game, she does not currently see enough appetite, at least for the path between the user's stub and the provider's recursive resolvers. On the other hand, Dickinson expects that the path between recursive and authoritative resolvers could be encrypted, running either DoT or DoQ. The DPRIVE working group just started to work on securing the upper part of the DNS resolution path. DoH is not being considered for this.

In the end, speed could be the decisive factor. "I happen to think DoQ will need to prove it is more performant in order for it to be chosen in preference to DoT for that role, because DNS folks are now reasonably comfortable with DoT". However other voices are pointing out that DoQ could still beat DoT, even for stub to recursive resolvers, because DoQ might be simpler to use.

Burdened by parallel deployments

For implementers it is hard to decide who to put their money on. There was a certain risk that one of the candidates would become dominant – and efforts to deploy the other protocols would be wasted, Wes Hardaker from the University of Southern California's Information Sciences Institute (USC/ISI) warned during IETF 109. Yet picking a winner upfront has not been the means of choice in the IETF recently.

Furthermore, implementers at Deutsche Telekom are happy to deploy at least DoH and DoT in parallel for now, while waiting for DoQ to arrive. This means that the race is on...

IX. Choosing the right encrypted DNS resolvers: who discovers the options?

The Adaptive DNS Discovery (ADD) working group (WG) at the Internet Engineering Task Force (IETF) has been trying to catch up with the deployment of encrypted DNS and met six times last year. Its goal is to provide standardised means of discovering which encrypted options are available to various network users, and a means for those same users to select the option most appropriate for their intended use. The work entails manoeuvering between technical tasks and policy choices that other WGs, such as the DNS Operations (DNSOP) WG were reluctant to pick up.

DNS queries are invisible to most internet users. Typically, the query for mapping a domain name to a server is sent by the web browser to the resolver as the user tries to visit a web address. In many places, the crucial resolver services have been operated by the network provider unless the user has specifically indicated that they want a different resolver service. Neither network nor DNS providers have made big efforts to educate the users about privacy issues in this arrangement, nor were privacy failures high on the agenda in the underlying protocols until after 2013.

But with a rising tide of privacy and security priorities for the internet's most fundamental infrastructures, service providers have launched a number of encrypted DNS initiatives. Choosing a secure and private DNS solution should be as easy as deciding whether to allow your browser access to your microphone or camera, seems to be the message.

Discovering Equivalent Encrypted Resolvers

A draft proposal by engineers from Apple, Cloudflare and Microsoft is making a first step with "Discovery of Equivalent Encrypted Resolvers" (DEER). Their aim is to provide two mechanisms for upgrading clients to encrypted DNS resolvers.

The first mechanism relies on querying a special domain in the .arpa TLD to look up encrypted DNS resolvers. The second mechanism fits the case when the hostname of an encrypted DNS server is already known to the user application. For the second case a new resource record type (SVCB) will convey information on the encryption protocol and blocked ports.

The proposal is yet to be adopted by the ADD WG, but nothing is easy in encrypted DNS. In a two-hour discussion the WG tried to establish whether the "equivalence" in „equivalent encrypted resolvers" is limited to queries, responses, name-pools, performance requirements or laws.

Harald Alvestrand, former IETF Chair and Google engineer, recommended not to make any equivalence assertions in DEER at all. In the end, he argued, DEER contains mechanisms for providing recommendations to end-users on encrypted DNS services and the end-users are capable of deciding for themselves how similar or different they want their DNS services to be.

Privacy, law and user expectations

Many experts pointed out that the wide-spread use of unencrypted DNS in user home networks implies an a priori lack of privacy expectations. Switching on DNS encryption would be a net benefit for this large user group, who are often completely unaware of the DNS.

In opposition is the view that users have chosen to trust their network providers, including through long-standing society discussions on content management and liability. Sending their queries on to a third party provider would change that equation.

Balancing the commercial and social interests involved in information management remains an issue for the internet standardisation community. While our common networked infrastructures are being made more and more robust against privacy and security threats, power dynamics that have reigned since the beginning of the 1990s are being challenged with the deployment of new technical solutions by new commercial actors. And even as the wild, wild web is again attracting criticism from, among others, Commissioner Thierry Breton, it is also true that as long as tradition is allowed to rule, we all know what we have got.

About our authors:

Gurshabad Grover is a technologist and legal researcher based in Bangalore, India, where he is Senior Researcher at the Centre for Internet and Society. Gurshabad's writing focuses on network security, privacy and censorship.

Monika Ermert has been working as an IT journalist for over 20 years. She has covered the evolving internet governance landscape, EU and worldwide attempts to regulate and the risks and fun of technology. She holds an M.A. in Chinese/Media Studies from the University of Tuebingen and lives and works in Munich, Germany.

CENTR is the association of European country code top-level domain (ccTLD) registries, such as .de for Germany or .si for Slovenia. CENTR currently counts 53 full and 9 associate members – together, they are responsible for over 80% of all registered domain names worldwide. The objectives of CENTR are to promote and participate in the development of high standards and best practices among ccTLD registries. Full membership is open to organisations, corporate bodies or individuals that operate a country code top level domain registry.

Rate this CENTR Report on IETF109

(Thank you for your feedback!)

CENTR vzw/asbl
Belliardstraat 20 (6th floor)
1040 Brussels, Belgium
Tel: +32 2 627 5550
Fax: +32 2 627 5559
secretariat@centr.org
www.centr.org

To keep up-to-date with CENTR activities and reports, follow us on Twitter, Facebook or LinkedIn