Council of European National Top-Level Domain Registries

Report on IETF104

Prague
22-29 March 2019

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

Highlights 3

Debating DoH 3

DoH Discovery 5

Running DNS in a more privacy-friendly way – special service or for all the DNS? 7

Fallout of the DNSPionage attacks:
The debate about standardising a registry lock 9

WGs and BoFs 11

DPRIVE WG: Recursive to authoritative is still a topic 11

RDAP at regext IETF – Policy/privacy-related or not? 11

DNSOP – DNSSEC, DNS Server Cookies and “mopping up” the special-TLD mess 13

SMART RG: “Encrypted data” removed from target list 13

Weirdest BoF: Validated brand logos in email 14

IETF News 15

Highlights

Debating DoH

The implementation of DNS over HTTPS (DoH) by Mozilla and Google, who both recently made announcements on their use of DoH in their respective browsers, fuelled the ongoing discussion on DoH during the IETF 104 week. DNS over HTTPS (RFC8484) was discussed by both the DoH and the DPRIVE Working Group. The DoH WG meeting aimed to focus on the future discovery of DoH servers. The DPRIVE WG meeting allotted a slot to Vittorio Bertola from Open-Xchange to present a draft on potential implementation guidelines for DoH Clients. Bertola's proposal somewhat mirrors Sara Dickinson's ongoing work for a BCP on privacy implementations for DNS servers, which originally started with DNS over TLS in mind. In addition to these meetings a dedicated side-meeting organized by Stéphane Bortzmeyer (Afnic) provided the two (or three?) camps with additional time to let off steam over what some DNS and network operators see as a coup d'Etat against their business models.

Laundry Lists?

By now network operators have become fully aware of the potential massive change DoH will bring by sucking up DNS queries from their customers. Instead of the DNS resolver being controlled by the respective network operator DNS, queries sent via browsers will be answered by external resolvers. So far these resolvers have been chosen by the browser companies. Cloudflare's global resolver network is currently the sole DoH operator contracted by Mozilla and Google. For the Chrome browser, queries will be resolved through Google's own resolver network. As long as a broader discovery mechanism for different DoH servers is not in place, the change from DNS to DoH will result in a considerable concentration of DNS traffic.

Market concentration has therefore been listed prominently as a major concern in two documents authored/co-authored by major telecom operators who are weighing in on the DoH debate. Comcast and British Telecom have partnered with Sky and the Georgia Institute of Technology to write "Centralized DNS over HTTPS (DoH) Implementation Issues and Risks" and with Deutsche Telekom, Open-Xchange and well-known DNS expert Jim Reid for "DNS over HTTPS (DoH) Considerations for Operator Networks". In the decentralisation document the operators express some shock about the switch-on manner for DoH via the browser implementations: "It appears to be unprecedented that a new protocol could be so rapidly deployed and thus displace an existing, long-standing, highly distributed protocol". They also list problems with the potential centralisation of DNS traffic:

- *change of the Internet ecosystem by the shifting of traffic to a few platforms*
- *decreased stability through fewer points of failure (while acknowledging that the concentration of DNS traffic is not new)*
- *possible security issues through fewer points of failure allowing the attacker to go for few sites only (including targeting of individual DNS administrators)*
- *loss of a more widespread visibility of security threats*
- *loss of parental control and other content control*
- *issues for split DNS and potential leaks of internally used domains through DoH requests*
- *reduced software diversity due to fewer players*
- *more commercial use of DNS data*
- *potential negative issues for localisation (a raison d'être for content delivery network traffic shaping, possible latency effects)*
- *DoH as a source or malware command and control (HTTPS://github.com/SpiderLabs/DoHC2/blob/master/Mitre_Attackcon_Playing_Devils_Advocate_With_Attack_1.0.pdf)*
- *issues with legally mandated DNS blocking (and disruption of walled garden or captive portals)*
- *increased complexity, problem for troubleshooting due to additional providers unknown to end-users*
- *business risks following concentration (smaller DNS software market, fewer public DNS operator choice, smaller market for CDNs, smaller DNS labor market)*

The centralisation document also makes a number of recommendations, including pushing for DoH discovery standardisation and the need for conventional DNS operators to start testing DoH, while slowing down the implementation through browsers through a mixture of technical and political/administrative steps. This means more testing and measurements, not allowing DoH to be the default, an ICANN review, community assessment, a push for DNSSEC implementation,

the development of centralised DoH Data privacy guidelines).

Whilst the "Considerations" document lists most of the above issues, it takes on the changes resulting from DoH from an operational perspective. This includes aspects like potential problems with IPv6-IPv4 NAT translation, failures in recovery/fault reporting/user support, and the question of "meaningful consent" from users. With DoH the provision of DNS and Internet connectivity might be decoupled, and users might unknowingly become customers of parties unbeknownst to them. For different applications, in the future DNS might be provided by different parties and via different protocols.

A basic question according to the network providers is who will decide which DNS servers (and protocols) are to be used in the future – connectivity providers, app/browser providers, users?

Fighting Camps: Network Gang vs Browser Gang – split DNS community

There was considerable backlash against the presentations of the network providers (Jason Livingood, Comcast, presented during the DoH session, Jim Reid during the dedicated session on concentration). In essence, three camps have developed in the DoH fight. They can be described as the browser/HTTPS/web camp (with Mozilla and its provider Cloudflare and Google in the vanguard) and the opposing network/connectivity providers (large telcos already afraid of losing out to the GAFFA). The third camp is those DNS providers that seek to use DoH as well as they can.

The DNS community seems to be split. A number of DNS operators (for example PowerDNS), and CDNs like Akamai (Ralf Weber made their case in Prague) reject the switch to a small number of trusted resolvers outside of their service networks. However, there are also those DNS operators who point to the potential positive effects with regard to privacy and anti-filtering effects. They recommend reconsidering the implementation in a more decentralised way. Representatives from both Afnic and CZ.NIC made comments in that line.

Privacy proponents from civil society organisations and open source DNS software developers question the arguments of telcos against DoH. They point to similarities with previously-attempted push-backs by telcos against better encryption of traffic (for example in QUIC and TLS 1.3). Together with DNS experts and the browser group, they speak of a laundry list of concerns in the draft documents.

Daniel Kahn Gillmor, a technical expert at the American Civil Liberties Union (ACLU) stated clearly that the call by telcos for informed consent comes far too late: "DoH has forced us to grapple with the idea that we are leaking data. We have never informed the user before. Now we may change who this data gets to, and this upsets people." With many users (in the US) not having a choice when it comes to connectivity providers and no GDPR-like legislation in place, the current default DNS resolution is "not necessarily more respectful of user choice." Yet the activist said he certainly did "agree with folks who are terrified of DNS over Cloudflare. But that is no DoH problem."

Mozilla's CTO Eric Rescorla, who prominently led the discussion in favour of DoH in Prague, also reinforced the argument calling out network operators who can and more often than not meddle with DNS resolution attackers: "Someone who controls the network but does not control your computer is an attacker." Rescorla did acknowledge the need "to make sure that the web level was not leaking information", because "since it is multiplexing there is an opportunity for leakage". Nevertheless, he demanded more focus on the positive aspects of DoH and on technical merits and policies ensuring protection against potential data leaking. Rescorla argued that while the GDPR was mentioned a lot in the debate as a policy that protected users against data leaking from DNS services (in the EU), it did not address filtering and blocking by local DNS providers.

From these discussions one fundamental difference becomes clear: telcos try to make the case for "good blocking" (parental control, malware filtering by the network provider, blocking of sites deemed illegal in a given jurisdiction) and offer the idea that DoH's encryption might only be necessary for "dissidents" (in non-rule of law-countries). Yet the browser community views interference by a network provider ("somebody with full or partial control of the network", Rescorla) as some kind of "attack" anywhere in the world.

Mozilla announces next steps on DoH

During the IETF week, Rescorla announced the next steps (without a clear timeline) concerning its DoH strategy, and in this announcement he acknowledged concerns with regard to concentration. Rescorla reiterated that there was "ample evidence of monitoring/manipulation of user traffic via this vector".

According to its CTO, Mozilla "would like to deploy DoH by default for our users" and "select a set of trusted recursive resolvers (TRRs) that we will use for DoH resolution." To address the privacy issues, the future TRRs would have to adhere to a privacy policy set by Mozilla that would "roughly match" the one Mozilla has put in place for the Cloudflare resolvers currently used.

The privacy policy would still be refined, the statement reads, but would be based on the following points:

1. Copies of Firefox will be configured with a set of TRRs. Different regions may have different TRR sets or different defaults. In addition we may have DoH/TRR on by default in some regions and not others, especially initially.
2. Users will be informed that we have enabled the use of a TRR and have the opportunity to turn it off at that time, but will not be required to opt-in to get DoH with a TRR.
3. Any given client will automatically select a resolver out of that set and use that for all resolutions [with the two exceptions noted below*]
4. At any time, the user will have the option to select a different resolver out of the list, specify their own resolver, or disable DoH entirely.

* The exceptions are: cases where the network also controls the client (e.g., they are able to remotely manage it via MDM); in this case the respective user/network should be able to select a resolver and/or disable DoH. Also where a system has a preferred resolver that is on Mozilla's TRR list, a choice should be possible (perhaps, Rescorla wrote, via Paul Hoffman's DoH discovery draft).

In the short term, the need for resolvers to be on Mozilla's list "creates some challenges for resolver operators. We would be open to discussing how to adapt our security constraints to suit the needs of multiple applications, so that as more systems deploy DoH/TRR, they can share a list of resolvers vetted to a common standard."

Just allowing network operators (or users, who Rescorla said should not be asked to decide for their resolver themselves) "to dictate the DoH resolver would obviate the security objective" intended.

Mozilla has meanwhile posted the privacy policy (including a transparency policy) that lists a limitation to retain data (24 hours only as long as data is not anonymised), a ban to market/sell/transfer the data (except for transferrals required by law), a ban to combine/aggregate this data with data from other sources, a ban to sell/grant access to it. Interestingly Mozilla obliges the TRR candidates to support query minimisation, but not to implement DNSSEC. While the company would welcome DNSSEC validation by the DoH resolver, they did not think it should be made mandatory, company representatives wrote in an ongoing discussion on the mailing list.

A list of DoH servers, browsers and tools can be found here.

DoH Discovery

All the "camps" do essentially agree on one point. A mechanism to allow for the choice of resolver is needed. It was referenced in Mozilla's next steps announcements, underlined as indispensable to fight further consolidation or market concentration on the DNS market by various speakers at the DoH meeting in Prague. "Without a discovery mechanism we will be stuck with Cloudflare and Mozilla", said Petr Spacek, CZ.NIC, during the discussion. The DoH WG discussed the respective discovery draft edited by Paul Hoffman, ICANN.

In short, the draft acknowledges that clients might want to use an internal or preferred external server for DNS resolution. The draft therefore proposes "protocols to get the list of URI templates [RFC6570] or addresses for the DoH servers associated with at least one of the resolvers being used by the operating system on which the application is being run." The two mechanisms envisaged are "DoH servers from HTTPS" to use "a well-known URI [I-D.nottingham-rfc5785bis] that can be resolved to return the URI templates in an HTTP response" and "DoH servers from DNS" that put resolver addresses into a new special use domain name (SUDN) [RFC6761] "that can be queried to return the URI templates as a TXT Rrset" (or allow to query the resolvers from a SUDN for A and AAAA Rrsets). Browsers need to have a special entry in their configuration interface in which the allowed DoH servers for the

respective traditional DNS (Do53) or DNS over TLS Servers (DoT) can be found ("DoH server associated with my current resolver"). Some preliminary thoughts were presented by Ted Hardie (IAB Chair) on the issue of the nature of a special domain name and its relation to the ICANN root.

Nevertheless, a major topic of controversy revolved around ensuring that when making the choice for resolvers, clients (end users) are not led astray and right into the hands of malicious actors who would then not even have to perform cache poisoning or other DNS re-routing in order to 'own' end users with regard to their DNS query and the ability to sell a biased or even faked DNS view.

Such security issues are clearly mentioned in the draft (alongside the potential privacy issues caused by both TLS and HTTPS, allowing for "user identification in ways that plain Do53 does not"):

"If DNS queries sent from stub resolvers to recursive resolvers are not sent over transports that assure data integrity and server authentication, the "DoH servers from DNS" and "Resolver addresses from DNS" protocols are susceptible to on-path attackers directing a user to a DoH server that is not actually associated with their resolver. Do53 is not a secure transport, and neither is DoT using the opportunistic profile."

It was impossible to authenticate unauthenticated sources, said both Rescorla and Patrick McManus from Mozilla. This explains the reluctance by Mozilla to open up resolver choice too much, including perhaps the idea of end user choice.

In the DPRIVE WG a draft on how to organise DoH (and DoT) server discovery while leaving Split DNS or security monitoring of the provider intact was presented. The concept essentially wants to use an "Enrolment of Secure Transport" (EST) server in the provider network as a control point for DoH and DoT discovery for the client (end user). This would allow the provider to allow secure transport (not blocking DoT and DoH) and at the same time enable them to continue security monitoring, Tiru Reddy (McAffee) explained during the DPRIVE WG. By inserting this control point at the network edge, split DNS would also be possible again.

DoH Debate – What next?

Another point was made on the efforts to form a "truce" between the various DoH groups, namely the fact that discussions should not question the protocol of DNS over HTTPS (RFC 8484). Instead the focus should be on how DoH will be implemented. Therefore the DoH discovery draft should be given special attention. Addressing implementation and operational concerns is also the topic of two drafts which are currently being discussed in the DPRIVE WG (see below).

One additional idea raised on the DoH mailing list is the idea of a special port for DoH. To ease setting DoH by default, Tomas Krizek (CZ.NIC suggests using Knot Resolver, which is intended to "use (port) 44353 as the default port for DoH". Using the classical HTTPS port 443 for DoH resulted in clashes, Krizek wrote. There was considerable opposition against the idea, with people complaining that using a new port would complicate DoH implementation, but also, that the very idea of hiding DNS traffic inside HTTPS traffic would be demolished with a dedicated port – which would stick out like the DoT port 853. CZ.NIC developers on the other hand argue that they do not expect quick consensus on the discovery draft and see the extra port solution as a way to ease deployment.

What complicates the whole DoH debate is spread over various places. After passing the basic DoH standard, the DoH WG is now working on DoH discovery. The operational issues and BCP documents for implementers are currently covered in the DPRIVE WG. None of the groups are interested in taking on the new drafts which focus on the concentration and operational issues. The DNSOP WG Co-Chair Suzanne Woolf rushed to underline that these were not issues for DNSOP.

Additional work on potential mechanisms to push additional DNS responses when answering queries to the client might also be revived in the future, McManus thinks. It would potentially be better for such work to be moved to HTTPBis.

To complicate the ongoing debate, the IESG felt it necessary to set up yet another mailing list (DoH is already discussed in the DoH and DPRIVE mailing lists at minimum). According to the new Area Director (ART), Barry Leiba, the Applications Doing DNS (ADD) mailing list will be dedicated to "DNS over HTTPS, DNS over TLS, implementation choices for those, application usage, operational concerns, privacy concerns, performance concerns, and any other such." Leiba encouraged engineers "to take all that related discussion to the new list and please stop discussing it on DOH, DPRIVE, DNSOP, and any other lists." While Leiba said that the

motivation for ADD (which might also become a WG) was to "avoid fragmentation", for now ADD only seems to add to the fragmentation, especially given the fact that Barry Leiba acknowledged that some work certainly cam under the scope of other WGs.

Leiba envisaged that a possible BoF in Montreal (IETF 105) "aimed at forming an 'ADD' working group, most likely in the ART Area but with significant crossover expected and desired from Ops, Sec, Int, and probably the rest of the -solar system- IETF community."

According to some observers the DoH-DoT debate could be bigger and should potentially be led by the IETF community at large. During the dedicated DoH side meeting, the idea of a Human Rights/Privacy WG was mentioned by one participant. Sara Dickinson (Sinedun), who last year called on the DNS community to take a deep look into the changes DNS over HTTPS might bring, asked the DPRIVE WG to consider a possible new WG that looked more generally into policy and deployment issues.

Dickinson currently co-chairs the newly-established Privacy Enhancements and Assessments Research Group (in the IRTF) and edits both a bis-version of the DNS Privacy Considerations (RFC 7626) document as well as a Best Current Practice Document for Operators of DNS privacy services.

Running DNS in a more privacy-friendly way – special service or for all the DNS?

During the DPRIVE session in Prague, Dickinson presented both the bis-version of the DNS Privacy Considerations document as well as the Best Current Practice document "Recommendations for DNS Privacy Service Operators". The follow-up (bis-) document on the DNS privacy considerations' RFC had become necessary given the considerable changes that have taken place over the last three years. The adoption of both the DoH and DoT RFCs marked major steps for DNS service operators. Besides considerations of new threats, for example the threats DNS services inherit from taking on HTTPS transport (potential for tracking) and attacks on encrypted transport, the document also added sections on the blocking of encrypted services as well as existing issues of personally identifiable data in the DNS payload (ECS, DNS Cookies).

The BCP mirrors the privacy concern document by providing a set of minimum standards (and recommendations or optimisations) that operators should adhere to if they want to call their service a DNS privacy service. The document addresses best practices for data on the wire (stub to recursive resolver), data at rest (data minimisation) as well as upstream traffic. It also includes a dedicated chapter on a "DNS privacy policy and practice statement". Such a statement, if standardised, would allow users (and monitoring parties) to compare different options they might have from the different operators. Policies to be covered are the handling and potential logging of IP addresses (PII or not?), the nature and condition (anonymisation?) of data aggregation and transfers, data retention times, data sharing or selling policies, declaration of partners in the loop, data correlation practices, filtering policies (legal or other filtering ongoing?). A practice statement should declare the current operational practices and deviations from it, the jurisdiction, agreements with law enforcement agencies/agencies, the mechanisms for users to contact the operator and enforce the policies and user consent policies.

Tables have been created by Dickinson and the DNSprivacy project which illustrate what a comparison of relevant policies and actual practices could look like. Dickinson said that she had to read 7000 lines of fine print in order to come up with the tables. A standardized DPPPS would allow easier comparison and monitoring. There was a brief discussion in the DPRIVE WG about whether the DPPPS framework should be put into a dedicated document, and some participants (like Dan York, ISOC) favoured this, as the audience would be different for the operators and the privacy policy declaration document.

An additional request to include recommendations with regard to how CDNs should work with the privacy-enhancing DNS services was made in Prague by Rich Salz (Akamai).

With regard to Best Current Practices, the WG also talked about the potential to add another BCP recommendations document that would focus on clients (instead of servers). Dickinson said that the time might be right to add this kind of client-side DNS Privacy recommendation's BCP. She suggested that putting client-oriented recommendations from her draft into a new document would be possible, albeit overlapping partly with a draft presented during the DPRIVE WG by Vittorio Bertola (OpenX-Change). Bertola's "Recommendations for DNS Privacy Client Applications" were prepared as a contribution to the DoH debate, but were better received than the Telco/ISP documents.

Policy

List Item	1	2	3				4		5	6	7	
Redirect NXDOMAIN	IP address are PII	IP address logging	Clear list of what data stored and for how long	Share anonymized data with partners	Share identifiable data with partners	Share or sell data to third parties	Exceptions to collection for attack analysis	non-profit	Partners	Combine DNS data with other data sources	Redirect NXDOMAIN	Block domains
Quad9 Secure	Y	N	Y	Y	N	N	Y	Y	IBM PCH GCA	N	N	Y
Quad9 Unsecured	Y	N	Y	Y	N	N	Y	Y		N	N	N
Cloudflare	Y	N	Y	Y	N	N	N	N	APNIC	N	N	?
Cloudflare DoH	Y	N	Y	Y	N	N	N	N	Mozilla/ Firefox	N	N	?
Google	N	Y(1)	Y	?	?	?	N	N	?	N	N	N(1)
OpenDNS	Y	Y	N	?	Y	Y	?	N	?	Y	N	?

(1) Only in temporary logs

Configuration Matrix		TLS	TLS 443	Strict Name	Strict SPKI	Cert 0	Cert 14	QNAME min	RTT 250	DNSSEC	Keepalive	Padding	TLS 1.3	OOOR
dnsovertls.sinodun.com	v4	✓	✓	✓	✓	✓	✓	!	✓	✓	✓	✓	!	✓
	v6	✓	✓	✓	✓	✓	✓	!	✓	✓	✓	✓	!	✓
dnsovertls1.sinodun.com	v4	✓	✓	✓	✓	✓	✓	!	✓	✓	✓	✓	!	✓
	v6	✓	✓	✓	✓	✓	✓	!	✓	✓	✓	✓	!	✓
getdnsapi.net	v4	✓	!	✓	✓	✓	✓	✓	✓	✓	!	!	!	!
	v6	✓	!	✓	✓	✓	✓	✓	✓	✓	!	!	!	!
dns.quad9.net	v4	✓	!	✓	●	✓	✓	!	✓	✓	!	!	✓	!
	v6	✓	!	✓	●	✓	✓	!	✓	✓	!	!	✓	!
dns.google	v4	✓	!	✓	●	✓	✓	!	✓	✓	!	!	✓	✓
	v6	✓	!	✓	●	✓	✓	!	✓	✓	!	!	✓	✓
1dot1dot1dot1.cloudflare-dns.com	v4	✓	!	✓	●	✓	✓	✓	✓	✓	!	✓	✓	✓
	v6	✓	!	✓	●	✓	✓	✓	✓	✓	!	✓	✓	✓
security-filter-dns.cleanbrowsing.org	v4	✓	!	✓	●	✓	✓	!	!	✓	!	!	!	!
	v6	●	●	●	●	●	●	●	●	●	●	●	●	●
unicast.censurfridns.dk	v4	✓	!	✓	✓	✓	✓	!	✓	✓	!	!	!	✓
	v6	✓	!	✓	✓	✓	✓	!	✓	✓	!	!	!	✓
kaitain.restena.lu	v4	✓	!	✓	✓	✓	✓	✓	✓	✓	!	!	!	!
	v6	✓	!	✓	✓	✓	✓	✓	✓	✓	!	!	!	!
dnsovertls3.sinodun.com	v4	✓	!	!	✓	!	!	!	✓	✓	✓	✓	✓	✓
	v6	!	!	!	!	!	!	!	!	!	!	!	!	!
dnsovertls2.sinodun.com	v4	✓	!	✓	✓	✓	✓	✓	✓	✓	✓	✓	!	!

Bertola underlined that the aim of the document was not to stop DoH, but rather to elucidate issues and possible mitigations from the client/application point of view. During the DPRIVE meeting, Bertola clarified the two basic concepts possible; the current one where DNS resolvers are chosen by the network a user sits in by default (while he retains the option to configure other DNS servers) and the potential future one, where applications come with their own choice of DNS (trusted) resolver.

Network-level	Application-level
All applications use the same resolver (the operating system one)	Each application uses its own resolver
The default is usually the resolver automatically suggested by the network	The default is usually supplied by the application (no local resolver discovery)
The user is in charge, either accepting the default or changing it in a single place	The application is at least partly in charge, choosing the default and/or constraining the choice to its own «trusted resolvers»

Issues that need to be addressed by DoH clients according to Bertola are:

1. *Trust model and user choice*
2. *Consolidation*
3. *Namespace fragmentation*
4. *Privacy*
5. *Content access control*
6. *Security and network management*
7. *Jurisdiction*
8. *Disaster recovery*
9. *User support*

With the DoH debate just gearing up, there seems to be quite a struggle ahead for the DNS community. Several DNS operators have already announced that they will get into DoH as well. For the DNS to remain "a little" decentralised (or get back to more decentralisation), a certain investment of time, energy and funding in DNS evolution seems to be necessary.

Fallout of the DNSPionage attacks: The debate about standardising a registry lock

The recent DNSpionage attacks that targeted a number of public authorities in Middle East Countries (in particular Lebanon, Iraq and Egypt) resulted in intensified talks about further security mechanisms for domain registrations. A dedicated side-meeting was set up by Alex Mayrhofer to evaluate the possibility of using a standardised registry lock as one counter-measure. While many registries offer some kind of registry lock, participants at the meeting in Prague warned that the lock could itself be "weaponised". After changing a compromised name server, the attacker could set up the lock, thereby complicating countermeasures by the legitimate owner.

Servers attacked according to Brian Krebs

nsa.gov.iq: the National Security Advisory of Iraq
webmail.mofa.gov.ae: email for the United Arab Emirates' Ministry of Foreign Affairs
shish.gov.al: the State Intelligence Service of Albania
mail.mfa.gov.eg: mail server for Egypt's Ministry of Foreign Affairs
mod.gov.eg: Egyptian Ministry of Defense
embassy.ly: Embassy of Libya
owa.e-albania.al: the Outlook Web Access portal for the e-government portal of Albania
mail.dgca.gov.kw: email server for Kuwait's Civil Aviation Bureau
gid.gov.jo: Jordan's General Intelligence Directorate
adpvpn.adpolice.gov.ae: VPN service for the Abu Dhabi Police
mail.asp.gov.al: email for Albanian State Police
owa.gov.cy: Microsoft Outlook Web Access for Government of Cyprus
webmail.finance.gov.lb: email for Lebanon Ministry of Finance
mail.petroleum.gov.eg: Egyptian Ministry of Petroleum
mail.cyta.com.cy: Cyta telecommunications and Internet provider, Cyprus
mail.mea.com.lb: email access for Middle East Airlines

DNSpionage: targeted attacks on DNS infrastructure as door opener to victims

The DNSpionage attacks combined several well-known attack vectors to produce what experts called a whole new type of attack. According to Patrick Fältström (Frobbit), attackers used stolen credentials, for example to change Netnod's operational servers in phases of an hour in order to be able to make changes to DNS entries to route traffic from the company's mail servers to their own servers and, while using quickly obtained certificates, they used their servers as proxies to phish for the targets' account and password information. A targeted attack on DNS infrastructure allows for one entrance ticket to the target victim's traffic and is hidden as they only use it for a short time.

According to Fältström, while the attacker's name servers were visible to Netnod for some time in the Whois, monitoring did not help, due to the fact that monitoring software only checks this information once every four hours. What became visible in one out of three attacks on Netnod were DNSSEC failures, but only when the attackers forgot (or consciously) did not remove DNSSEC on the domain for a third stage of the attack, meaning that the validation failed. The DNSSEC aspect illustrates that whilst it is one possible countermeasure, DNSSEC does not protect against attacks once the attacker gains access to registrar credentials and can change domain information. It will be interesting to see how such attacks change in DoH settings.

Registry/Security Lock

Another potential countermeasure is to lock-down registration data at the registry and make changes dependant on more or less manual interventions and is now being discussed by the DNS community. During the side meeting in Prague, participants from registries noted that they offer registry locks (with a few exceptions: .de, .ch, .br, .ua). In most cases, the registry locks are turned on and off by registrars, who are also the ones that can process change through some sort of manual processes (fax, phone call, note-sharing and passwords). Registries that work in this way include VeriSign (.com, .net, .name, .cc, .tv.), .fr, .jp, .ca and .se. A number of registries (.at, .cz), rely on the action of registrants before they process changes. There is also the example of a VIP domain status under .dk, which requires confirmation for every EPP request to change data in the registration, making all changes "asynchronous".

The intent of the meeting was to consider a potential standardisation of registry lock processes as currently, the procedures which have been put in place by registries vary considerably from one to the other, making it difficult for registrars to implement them all. Prices for the service can range from zero to 500€.

One attempt at standardising was presented by Ulrich Wisser (.se) during the session. The draft RFC puts out an EPP extension that inserts a manual authorisation step inside the EPP to protect changes to an object by the sponsoring client or its customer. The draft RFC, which is now on the table of the RegEXT WG requires "additional authorisation for transform commands", using in-band EPP options available through EPP Standards [RFC5730], [RFC5731], [RFC5732], [RFC5733].

With a registry object locked, transform commands can only be executed if proper authorisation is provided (or the object was unlocked out-of band). There are a number of open questions to be discussed in the WG.

At the same time, a number of participants at the side meeting called for a clear statement of the motivation for standardisation ("what are we trying to achieve?") and also clear terminology ("what does lock mean?"). One counter argument against standardisation was that diversity might be a feature, not a bug, as diversity could make attacks more difficult as well.

An important discussion also took place around the need to include two-factor-authentication in the registration and domain managing EPP processes.

Conclusion of a DNSpionage victim

A registry lock is an "extremely heavy tool", and is perhaps "too heavy for normal business", as it made quick changes difficult and tedious, concluded Fältström, talking to this reporter. While acknowledging that Netnod and Frobbit did "not have the horses in the barn" and two-factor authentification (as well as registry lock for Netnod) was in the planning, he suggested that the community should consider something "between registry lock and nothing".

One option Fältström mentioned was a push notification which registrants could subscribe to with the registry, that would alert them to actual changes. Another option was to find a simpler (more light-weight) registry lock solution. Registries might also consider installing monitoring systems, comparable to credit card companies that check "abnormal behaviour". Other general hygiene recommendations which are already out there (also in SSAC advice) was not to use cleartext inside one's network, and to control one's own nameserver (instead of outsourcing it to an external provider).

The discussion about registry lock will continue and Wisser's draft is open for comments.

WGs and BoFs

DPRIVE WG: Recursive to authoritative is still a topic

Besides taking some time to talk about DoH, the DPRIVE WG discussed the way forward for the long-standing question of whether queries travelling between recursive resolvers and authoritative servers should be protected. Alex Mayrhofer (nic.at) and Benno Overeinder (NLnet.labs) have taken up the topic after several attempts to get the discussion going.

Users signalling their privacy wishes?

Mayrhofer and Overeinder laid out the issues to discuss in a future draft (see their Github document) and asked for more comments in Prague. They asked the group if it should become more of a prescriptive document (operators must do privacy-friendly recursive-to-authoritative), or if it should present other options. A related question was how to deal with the differing interests of operators, users and developers. Decisions on how to organise the signalling of what the different parties offered/wanted could differ depending on the answers.

Another question that was addressed was whether DoT would be the protocol of choice for protecting queries travelling between the recursive and the authoritative resolvers. The WG discussed this briefly and seemed to reach a consensus that DoT is the way to go, at least for now. According to Mayrhofer, the functional aspects of the future draft could include privacy protection mechanisms, the authentication of servers (how to deal with non-authenticated authoritative servers), performance, the detection of availability (by zone, by identified nameserver or by IP-address), as well as end-user policy propagation.

The discussion erupted over what signalling might be needed and how end-users' wishes should be reflected. During the scoping discussion, Mayrhofer pointed out that the interests of the different parties might not be aligned. From a user's point of view, the transitive trust established when queries go up to an authoritative server could be problematic, as the user has had no "chance to identify which data was exposed to which authoritative party (via which path)". Users might potentially want "to be informed about the status of the connections which were made on their behalf", the authors reminded the WG, also sparking a debate about potential options to allow end-users to receive signals about the choices made.

Most participants clearly favoured signalling to applications only. Signalling to users was too hard, said Daniel Kahn Gillmor (ACLU), and the IETF was not good at it. Eric Rescorla, CTO of Mozilla, said it was unclear to him what a signalling of choices for users would result in. Others argued that while it is nice to have, signalling to end users should only be considered at a later stage in order to avoid further delays to the production of the recursive to authoritative document. Sara Dickinson (Sinedun) argued that another option would be for users to trigger signalling only when they want resolution, as long as it would not expose their private data.

The discussion on the document will continue on the list. Interestingly, Mayrhofer also openly asked if the discussion should be taken on by the DNSOP WG, since it may ask for all DNS operators to use the privacy-preserving mode. As with the DoH debate though, DNSOP Chairs seemed happy to keep the privacy discussion outside of DNSOP for the time being.

More workarounds for DNS privacy

The DPRIVE WG briefly discussed the possibility of easing the implementation of DoT. Manu Bretelle (Facebook) presented the idea of using the combination of a delegated simple public key infrastructure and DNSSEC at parent level to allow insecure sites to participate in DoT without themselves being forced to introduce DNSSEC. Since the signature for the PKIX comes from the parent DNS servers, servers which are lower down in the hierarchy would be able to introduce DoT without making the effort to deploy DNSSEC for authentication. Bretelle's draft wants to introduce a "Delegation SPKI (DSPKI) resource record" for that purpose. Reactions at the WG have not been conclusive so far.

RDAP at regext IETF – Policy/privacy-related or not?

Since IETF103, RFCs 8521, 8495, 8543 and 8544 have been published. With two more documents on their way to the IESG review (Registration fee extension and strict bundling registration), the WG is looking ahead and has to make a decision about how many new milestones it will take up under its renewed charter. The question is: should RDAP get its own WG?

As four out of five documents were chosen to become new milestones related to the Registration Data Access Protocol (RDAP) during an interim meeting, the WG discussed if the whole RDAP effort merited a dedicated Working Group to allow the work on EPP extensions to continue. George Michaelson (APNIC) argued that RDAP had long been slightly neglected as the solution to a problem the community had, and with a lot of work ahead on the Whois follow-up protocol, a special WG could make sense. According to several experts, RDAP will impact a much larger community than EPP, which is only of interest to about 20 back-end providers and their 20,000 registrars.

Several participants in Prague were clearly against the idea of splitting the work, especially given the fact that so far, the WG has always been short on experts to review the draft. By splitting the work, expert review might become even more elusive. Those following the work of regext are the same people who would follow RDAP standard suite development. Through rough consensus, the group also recommended that the incoming new area director Barry Leiba should be more flexible with the documents (including the number of documents) taken on.

Everything RDAP – and some policy questions

The WG will be working on and trying to standardise four RDAP-related documents, namely:

- Federated authentication for RDAP
- RDAP Query Parameters for Result Sorting
- RDAP Partial Response
- RDAP Reverse Search
- Login Security Extension for EPP

Federated authentication is an older topic that has been presented by Scott Hollenbeck several times over recent years. Scott Hollenbeck's draft summarises how RDAP will perform authentication of a browser-based client. The RDAP client (OpenID user) queries RDAP servers, which check with an OpenID Provider if the RDAP client is authentic. A match of the client ID token and access token (received from the authorisation server) authenticates the client vis-à-vis the RDAP server and allows for (differentiated) access (depending on policy).

Three-level test implementations are being run by VeriSign Labs. They offer basic answers for unauthenticated users, a larger set of information for those identifying via Google mail and Microsoft Hotmail. In addition to this, for those fully authenticated ("using more restrictive identity providers", namely https://testprovider.rdap.verisignlabs.com/ and CZ.NICs https://www.mojeid.cz/) all information has been made available.

Mario Loffredo of .it Registry presented three other RDAP proposals the WG will be working on under its new milestones:

- "RDAP Query Parameters for Result Sorting" (allowing to organise and limit query results for access data, including registration metadata),
- "RDAP Partial Response" (allowing to receive subsets of possible query results to save bandwidth and time) and
- "RDAP Reverse Search" (allowing to search for all domains related to an entity, registrant, email, address).

Loffredo asked the WG members if they felt that the privacy issues related to reverse search were aptly addressed and received mainly negative answers. Stephane Bortzmeyer concluded that the "Privacy Consideration Section" of the draft only went as far as to confirm that local laws had to be complied with. Instead of confirming the obvious ("follow the law"), the section at least had to describe the risk associated with the reverse search. Loffredo argued that he wanted to focus on the technology in the draft instead of dealing with potential risks and rules outside the scope of the draft. Sensitive registration data MUST be protected and accessible for permissible purposes only. The section mainly underlines that "RDAP servers MUST provide reverse search only to those requestors who are authorized according to a lawful basis" and also mentions "performing a specific task in the public interest that is set out in law" as a legitimate reason or the "permitting reverse searches, which take into account only those entities that have previously given the explicit consent for publishing and processing their personal data". Discussions about the privacy issue related to reverse search will certainly continue. In fact, the notion that policy has no place in RDAP documents seems to be fallacious, given part of the stated motivation for reverse search in the draft:

The first objection has been caused by the potential risks of privacy violation. However, TLDs community is considering a new generation of Registration Directory Services ([ICANN-RDS1], [ICANN-RDS2]), which provide access to sensitive data under some permissible purposes and according to adequate policies to enforce the

requestor accreditation, authentication, authorization, and terms and conditions of data use. It is well known that such security policies are not implemented in Whois ([RFC3912]), while they are in RDAP ([RFC7481]). Therefore, RDAP permits a reverse search implementation complying with privacy protection principles.

Other participants, including Peter Koch (DENIC), reiterated the need to consider privacy issues more thoroughly in regext given that RDAP had developed into some kind of "passenger name records" for governments. The question about what was supposed to come first, requirements developed by ICANN or the technical implementation at the IETF, was also raised. Koch warned against the danger of "policy laundering" through a technical WG at the IETF.

Ideas about a possible privacy draft on RDAP (one central document) were briefly mentioned, but once again might be rejected due to the "let's keep it technical"-mantra.

DNSOP – DNSSEC, DNS Server Cookies and "mopping up" the special-TLD mess

So far, the DNS Working Group has avoided taking on the DoH or DNS Privacy discussions on their agendas, happy to have these discussed at DPRIVE or elsewhere. It will be interesting to see if that might change, given the calls for the DNS Privacy BCP to become operational practice (or even called for by local regulation) for all DNS providers.

On the other hand, keeping the DoH controversy at bay might be the result of the now three chairs' reluctance to overload their agenda, which is already pretty packed with drafts on:

- multiprovider DNSSEC (offering several models of how keys could be either shared or several key sets by the DNS providers of one customer be used),
- running local instances of root zone (aka hyperlocal root zone development),
- recommendations against switching servers in case the DNSSEC validating servers fail, guidelines on TCP as transport protocol for DNS.

A new draft that is being discussed is an attempt to standardise DNS server cookies, which so far have been constructed in highly diverse ways by programmers.

A rather policy-leaning discussion that cannot be further delayed by the WG is the one concerning special TLDs. WG Co-Chair Suzanne Woolf argued that the current specification to allocate special top-level domains for non-DNS services (such as the Tor domain .localhost, .onion, RFC 6761) had to be revised or clarified in order to avoid more people coming to the IETF for TLDs and thereby opening a potential avenue for people trying to circumvent the onerous and expensive ICANN new TLD process. Even within the WG, there is no consensus yet on how to deal with the Special Use TLD RFC.

In her proposed document, Woolf makes an attempt to give further guidelines on what could be considered a special name. Another option considered by the WG would be to put the RFC to rest as "historic". Several participants pointed to the necessary cooperation with ICANN for clarifying potential processes. Peter Koch (DENIC) suggested that the debate might also need additional audiences within the IETF as a whole.

One former applicant for a special use domain, researcher Christian Grothoff, declared that after being rejected in receiving .gns, GNU was cleared by creating .gns as an encrypted name resolving system available in parallel to the DNS.

SMART RG: "Encrypted data" removed from target list

Another group getting organised in the IETF is the Stopping Malware and Researching Threats (SMART) Research Group. Held as an Internet Architecture Board (IAB) meeting, the gathering was packed, perhaps thanks to a rather high-level guest, Ian Levy, Technical Director of the National Cyber Security Center (NCSC), the defensive/cyber security body of the British intelligence service General Communication Headquarter (GCHQ).

The NCSC has been a main force behind the initiative for the research group, which in the original draft charter declares that it "will investigate how cyber attack defence requirements can be met in a world of encrypted data". According to the new version of the charter, the SMART RG declares it "will research the effects, both positive and negative, of existing, proposed and newly published protocols and Internet standards on attack defence." According to Kirsty Paine (NCSC), the main goal is for designers, implementers and users of new protocols to be better informed and for the SMART RG to become "the authority" for attack defence in the IETF/IRTF to be consulted by developers.

In his presentation (which was the last in a pretty packed SMART agenda), Levy promoted the work of his agency (recommendations, annual reports to make security better; it tries to make it easier for users to "use cybersecurity"; it also develops a red-yellow-green label for IoT products; and pushes for the adoption of DMARC in the UK administration). Another project is to build a national BGP peering platform for British ISPs to avoid BGP hacks, as BGP was even worse than its reputation let on. The agency finally blocked huge numbers of queries from UK public agencies (450,000 WannaCry, thousands of Conficker).

Security incompatible with resilience

Levy said that security was more and more baked into protocols and warned that encryption was not the same as security. "Encrypting something does not make it secure", he said. TLS for example and initiatives like "let's encrypt" are all very well, but "always remember, the bad guys use the shiny too". Developers therefore need good information when they make their decisions not to enable new attack modes. Security, privacy and resilience are different. If done badly, security and resilience are incompatible. SMART is therefore important from the NCSC's point of view.

Daniel Kahn-Gilmore from the ACLU acknowledged that the IETF needed many more discussions on user interface failures and user interface people needed to come to the IETF to say what signals they needed. At the same time, obliging providers to cooperate with intelligence services and law enforcement could be "used by the bad guys, too", much in the same way that others used the "shiny" security protocols. In relation to this, Kahn-Gilmore also asked Levy's intentions with the so-called "Ghost proposal", a proposal in which Levy and his colleague, the GCHQ Technical Director, propose to allow intelligence agencies to become a "silent" party in encrypted conversations with specific targets.

Levy said to this reporter that he quickly agrees with the ACLU that it would not be good to have a centralised key-escrow for encryption. However, the basic problem of how intelligence services could get to encrypted communication needed to be solved.

First draft document in SMART

The WG also briefly discussed its first draft document, a lengthy draft on endpoint security capabilities and limitations. According to Arnaud Taddei from Symantec, out of 275 types of attacks, only 32 could be detected at the endpoint. The argument that control by network operators is indispensable has been made in several recent discussions on new protocols (e.g. TLS 1.3 or QUIC). The draft document is intended to become a first reference on attack vectors for protocol developers.

Weirdest BoF: Validated brand logos in email

The idea was not well received at the IETF, but a two-hour BoF was still spent on a proposal from several US companies, including Valimail, Agari and network provider Comcast, to allow large brand owners to publish brand indicators for domains and use them for authentication, based on existing standards like Domain-based Message Authentication, Reporting, and Conformance (DMARC). "If both the email and the logo authenticate, then the receiver adds a header to the message, which can be used by the MUA (mail user agent) to determine the domain owner's preferred brand indicator."

Assertion for the graphical logo of the brand owner is made through the publication of a text record in the DNS ("default._bimi.example.com"). Authentication for the record is performed via a check using a Certificate Authority (in the way TLS-certificates are checked).

The proponents argued in Prague that BIMI might push for the adoption of e-mail authentication standards, Sender Policy Framework (SPF), DomainKeys Identified Mail (DKIM) and DMARC, which provide mechanisms for domain-level authentication for email messages. The adoption of these standards has been slow so far, and BIMI making use of the mechanisms might change that. During the discussion, Seth Blank said the BIMI draft intended to "provide mechanisms to prevent attempts by malicious domain owners to fraudulently represent messages from their domains as originating with other entities".

The very idea that the mechanism used by large brand owners could be marketed as an anti-phishing tool was rejected at the BoF session because the mechanism would not deliver, as domain and CA-based authentication only allowed for a party to have control over a certain domain at best. The fact that no central, acknowledged database for the relevant

intellectual property rights was not available – and IP rights were globally disputed in many aspects – was another objection raised during the session.

The authors distanced themselves from earlier announcements that anti-phishing was a goal. However, they did acknowledge that the proposal has several problems, namely that the graphical logo concept was only for large brand owners (who own such logos and being able to make the necessary investment to propagate their logos via the BIMI structure). The authors also listed a number of rather grave security concerns (see also the long security section in the overview draft). The logo can easily be abused as a web bug to track users, malware can be hidden in the payload and copycat logos (similar to the ones of large brands) could be used.

There was overwhelming consensus that users would be misled into thinking that with the logos displayed, their email was more secure. Several developers, including David Schinazi (Google), called on the IETF community to never standardise such a mechanism. Interestingly, according to earlier press releases, Google was one of the supporters of the project ("BIMI is an Initiative of the three largest mailbox providers Microsoft, Google and Oath [Verizon, AOL, Yahoo] as well as Comcast, Agari, RP, Valimail and PayPal"). The BIMI proponents said they were considering next steps and would possibly ask for another BoF.

IETF News

The IETF Administrative Oversight Committee (IAOC) is history. At IETF104, the IETF community had its first opportunity to meet the new LLC board members. Following the IETF stepping up to become a legally independent organisation (responsible for hiring, contracting and fundraising outside of ISOC), LLC members met alongside the Prague IETF meeting. The members are:

- Maja Andjelkovic
- Alissa Cooper
- Jason Livingood, Chair
- Sean Turner, Treasurer
- Peter Van Roste

Agendas and minutes of the LLC Board can be reviewed here. Interesting points on the LLC's current IOC agenda include the search for an executive director and budget planning.

IETF105 will be held on 20-26 July 2019 in Montréal, Canada.

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 54 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

Rate this CENTR Report on IETF104

(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