Council of European National Top-Level Domain Registries

Report on IETF105

Montréal
20-26 July 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

Expert panel: Privacy considerations a MUST for standardization 3

Entrenched Positions: The ongoing discussion on DoH 3

A new threat model: Attacks on end-user devices and systems 7

IETF in crisis mode 8

WGs and BoFs 11

DNSOP Working Group: More straw for the camel, or how to avoid policy-sensitive issues 11

REGEXT: More privacy discussions are needed, EPP-RDAP dual stream 12

DPRIVE: Privacy work, DoT/DoH discovery questions, private zone transfers 14

BoFs on LAKE, MOPS and barriers to newcomer integration 15

Highlights

Expert panel: Privacy considerations a MUST for standardization

In a new version of the technical plenary, the IETF community listened to two well-known US privacy experts. Arvind Narayanan, Assistant Professor at Princeton University and Steve Bellovin, former long-time IETF developer, former IAB Chair and Professor at Columbia University, called on the IETF community to consider privacy when developing standards.

Narayanan and Bellovin both said that the privacy reviews which have already been completed and the privacy consideration sections in some RFCs were good, but Narayanan also underlined that they could be complemented with implementation audits once the standards were adopted. Narayanan, who also leads the Princeton Web Transparency & Accountability Project (WebTAP), emphasised the potential for cooperation between developers and academics on the topic. Academics could be lured with prices, Narayanan said.

The Princeton WebTAP is an example of audit-like work. It is a software platform that is curated by a group of researchers and used to perform regular screenings of a million sites for privacy violations, for example fingerprinting or leakage of data. Contrary to earlier studies stating that 90 percent of users were identifiable via their browser through fingerprinting, more recent studies found that the numbers were lower, according to Narayanan. "The ship has not sailed," he said, referring to fingerprinting. In addition, he said that developers were not alone in the fight for privacy, pointing to the decision by the US Federal Trade Commission to fine Google for deliberately circumventing Safari's blocking of third-party-cookies in 2012.

In recent years, Steven Bellovin has worked in the political sphere, including as an advisor in the Obama administration. He called on the techies to take into consideration the political environment, as it was necessary to share technical expertise with legislators and regulators. Regarding what the IETF could do, Bellovin also added that developers should avoid allowing too much flexibility in the implementation of their standards. "Leaving it to implementation should stop, as this is what becomes finger-printable afterwards", he warned.

Bellovin's main argument in his privacy talk was that the current privacy paradigm is broken. "Informed consent is dead", he said. The concept, which stems from legal research in the '60s, has not kept up with new developments. Privacy statements to which users consented were too long and complex. The business models of data brokers, which do not even have a contractual relationship with the users, were stated as an example. According to Bellovin, it is not clear what could replace informed consent. A declaration of use, which would allow users to declare what can be done with their data, was also impractical. The provision of a user-friendly interface to make such a declaration was tricky and enforcement was an issue.

Bellovin said that what was necessary first and foremost was a new privacy paradigm. "This is a task for research", he said.

Bellovin also made a brief comment on the controversial topic of DNS over HTTPS (DoH), underlining his scepticism with regard to the aggregation of traffic that would be favoured by DoH implementation.

Entrenched Positions: The ongoing discussion on DoH

In Montréal, both Mozilla and Google presented their plans for DoH to the IETF community during the Applications Doing DNS (ADD) BoF – and they differed considerably.

Mozilla: encryption hygiene, the end of the DNS as a control point

Martin Thomson (Mozilla), who is also on the IAB, announced that the company was gearing up to implement the technology with an opt-out for places/users that have "controls" in place. In short, if users/networks do not want to "be DoH-ed", they must signal it to the browser. According to current plans, the signals will be un-authenticated, opening them up to abuse through downgrade attacks. "It is a stopgap", Thomson acknowledged.

Nevertheless, despite considerable criticism from network providers, Mozilla wants to stay on course. If the DNS' name space were not fragmented, the company would already have shipped DoH, Thomson said. However, a late start also allowed for the issue of service discovery to progress.

During the well-attended BoF session, Thomson explained that the main reason for this decision was the attempts to use DNS as a control point. He stated that issues such as content filtering, malware detection and blocking, captive portals, enterprise- and network-specific service access, routing policies and regulatory mandates are the reasons why applications did not even want to use the DNS in the first place. Many of the "control techniques" were not functional, and filtering based on DNS resulted in over- or under-blocking.

Switching on encryption undermines the DNS used as a control point, Thomson said, adding that alternative naming was already happening, for example with RFC 7838 "HTTP Alternative Services":

"Alternative services do not replace or change the origin for any given resource; in general, they are not visible to the software "above" the access mechanism. The alternative service is essentially alternative routing information that can also be used to reach the origin in the same way that DNS CNAME or SRV records define routing information at the name resolution level. Each origin maps to a set of these routes -- the default route is derived from the origin itself and the other routes are introduced based on alternative- service information."

Thomson predicted that in the long run, applications (not only browsers) would want to encrypt everything and choose who they trust with the data. In that case, the only way to exercise control over traffic would be to engage with the endpoints directly. The topics of filtering, monitoring and "control at the endpoint" came up several times in the following discussion. Meanwhile, a draft resulting from the TLS 1.3 discussions has been published, detailing the operators' concerns regarding the change to monitoring and securing endpoints only (instead of monitoring at the network gates).

The call from operators to allow end-users to decide for themselves who answers their queries in the DNS resolution process is not implementable. The DNS is part of the "plumbing" and nobody would expect users to care about the pipes in their house. Thomson also reiterated that Mozilla's core motivation was that the company saw encryption as basic hygiene, and privacy and security should not be optional.

Just after the IETF, Mozilla published a blog post on their plans for DoH, which sound a little more careful than expressed during IETF Montreal. The post also announced more testing "to understand how often users of Firefox are subject to these network configurations (parental controls). To do that, we are performing a study within Firefox for United States-based users to collect metrics that will help answer this question. These metrics are based on common approaches to implementing filters and enterprise DNS resolvers."

Google: Treading carefully

Kenji Baheux, Chrome and Web Development Manager at Google, announced a much more restricted approach. A unilateral decision about a switch to the new resolution could confuse users who had certain expectations about their resolution. "We will not force a change of DNS providers on users", Baheux said. Google also thought that businesses should be "in charge of their users' experience".

Regarding the roll-out-plan, Google's DoH document notes:

"We don't have any current plans to rollout the experiment described above past 1% stable. If the experiment has performed well, then we will consider a full launch, which would likely include a formal settings UI and improved support for secure mode, especially around captive portal resolution. The rollout would require a new Finch flag to control whether the UI is visible.

For the moment, Google wants to stick to an experiment with 5 to 10 DoH service providers who fulfilled a set of privacy criteria. If a user is already using a server from that list, "we will upgrade to DoH with that provider. A user interface for choosing providers could be added later"."

The published "Testing plan" reads:

"For the automatic mode experiment, we should verify that the system DNS config is discovered properly on an actual Windows device, both with and without a VPN. We should also manually confirm that the Android DoT config is properly read and upgraded. Before we formally launch secure mode, we should verify that captive portal resolution succeeds on all platforms for a range of live captive portal implementations."

Contrary to Mozilla, Google seems to be shifting down a gear with its DoH development. It is unclear whether this is the result of calls from network operators and/or politicians, but during the open mic discussion, Baheux reported that in talks with ISPs, Google learned that simply moving forward with DoH was untenable. "If the choice is between a world where we can achieve

99 percent security against cyber threats, while still providing law enforcement 80 percent of what they seek and a world where we have boosted cyber security to 99.5 percent but given law enforcement zero access, the choice for society is clear."

Despite this careful treading, David Schinazi, a former Apple engineer who has since moved to Google, presented a draft for an additional element for a future DoH environment. For performance reasons, he recommended that web servers themselves indicate to clients which DoH server could best resolve their addresses. Therefore, an HTTPS origin would indicate its preference regarding the DoH server to be queried by a client with a "DoH-Preference header field":

*DoH-Preference = doh-uri *(OWS ";" OWS parameter)*

doh-uri = quoted-string

parameter = token "=" (token / quoted-string)

The (so far) rough draft by Schinazi, Nick Sullivan and Jesse Kipp (both Cloudflare) is available here. During the ADD BoF session, Schinazi explained that the concept could be used to mitigate response times when DoH requests end up at servers that are not optimal to resolve names, for example one CDN instead of another. Browsers would ship with a list of vetted DoH servers, would use preferred DoH providers where a preference was conveyed and where a proposed server was on the list, and content providers themselves could provide a vetted DoH server themselves.

One might wonder what will happen if the parallel development of origin servers meant that local networks announce their preferred DoH server, which would result in complicated negotiations about where the DNS resolution finally has to resolve.

Operators' concerns

This time, network operators/ISPs were represented by British Telecom. IETF newcomer Chris Box laid out a list of issues and concerns which were already well-known and communicated, namely the impact on their NATs, proxies, captive portals, load balancers and the effects for CDNs.

It is interesting to note that while the HTTP-DoH camp (such as Google's Schinazi) is making an attempt to clear some agreed-upon barriers such as CDN performance and service discovery to avoid monolithic resolution, network operators/Telcos/ISPs point out that they are interested in locally implementing DoH themselves. Cox underlined that he was not opposed to DoH but wanted the IETF to establish a Working Group to develop best practices for operators to run DoH servers. He not only asked how best operators should run DoH servers locally, but also if operators should put DoH in home routers.

Jim Reid, also speaking from the operators' camp, put up a few additional barriers to DoH, asking questions about potential conflicts resulting from the marriage of the two protocol worlds (HTTP and DNS). According to Reid, from the HTTP side, HTTP Push is an issue. From the DNS side, he listed standards such as DNSSEC, negative caching, the handling of TSIG signed requests, dynamic updates and other DNS additions, asking if DoH servers would have to implement them.

He also underlined the "Gretchen-Frage" of the traditional DNS – how will DoH servers treat the IANA root, will it be possible to have different responses depending on where resolution will take place or will DoH answers differ from "vanilla DNS" (now also labelled Do53).

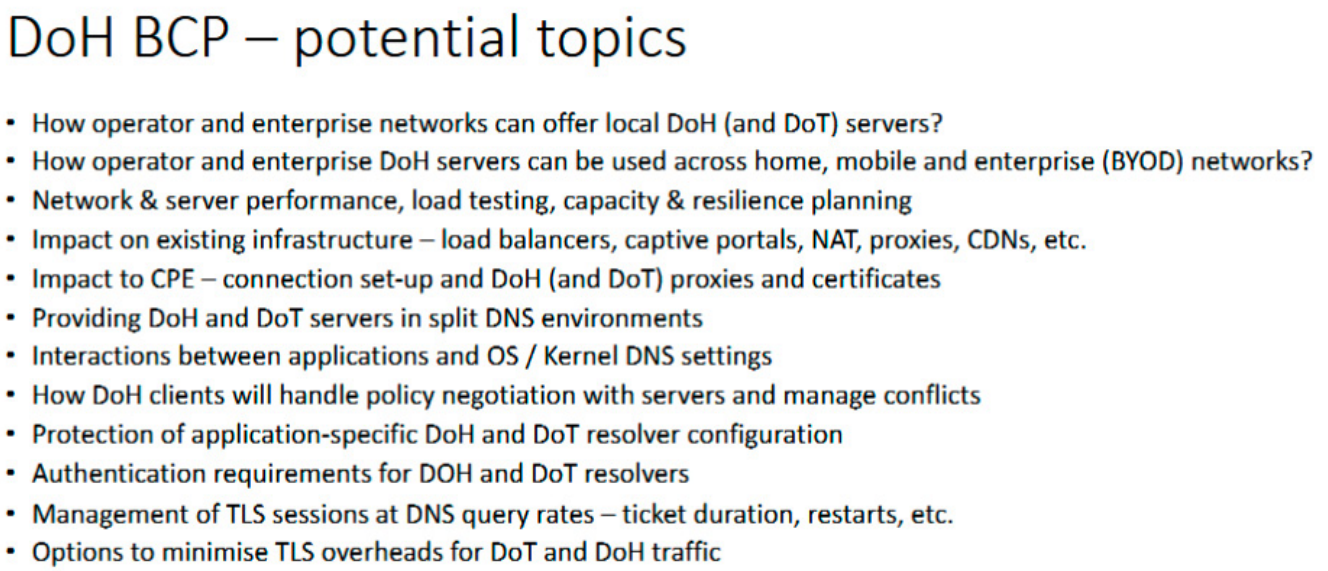

Asking DoH providers to be "nice DNS citizens" by implementing DNSSEC and other standards which are not mandatory (and not implemented) by traditional DNS providers was hypocritical, said Stéphane Bortzmeyer (Afnic). talking to this reporter.

On the other hand, by presenting privacy as a core motivation, Thomson said they had to face enquiries over their privacy "cover story", as DoH only secured privacy on the path. Regarding the "Trusted Resolver" implementation, privacy policies and possible leaks from HTTP headers have to be checked.

End-game for the DNS? Next steps

During Thomson's talk, it became clear that there was a challenge to the idea of the DNS as a central part of the infrastructure. Thomson argued that the developer community should consider whether the functions provided by the DNS could be performed in other ways. Alternative naming options are already out there, he said, hinting for example to Mark Nottingham's RFC on HTTP naming.

Wes Hardaker said that one of the pending questions was: "should the IETF standardize the concept of name resolution done per application, making a choice or leaving this to an end-game?" Leslie Daigle challenged the BoF, calling for more consideration on how the DNS could evolve and clearer analysis of the architectural consequences, also by defining what an application is, what the services are and what a naming system is. Lorenzo Colliti (Google) also called it an architectural issue. "The problem arises when clients use resolvers that are different from what the network configures". It was not a DoH issue, plus there were not technical issues with DoH. For some of the proponents of the DoH suite, the ongoing discussion with operators is a reminder of the fights around backdoors for operators and network enterprises in TLS 1.3. In both cases, the loss of "control points" resulted in lengthy discussions.

The concentration of DNS traffic streams – depending on implementation – was a major concern raised by several speakers during the ADD BoF, including by former IETF Chair Jari Arkko, as well as Roland Rijswik (NLnet Labs). In a draft presented for the emerging discussion on a new threat model for internet protocol design (see below), Arkko wrote that the DNS is showing some signs of ageing yet was changing very slowly, motivating the change to DoH. Nevertheless, Arkko warned that "while the security of the actual protocol exchanges improves with the introduction of this new technology, at the same time this implies a move from using a worldwide distributed set of DNS resolvers into more centralised global resolvers", creating welcome targets for pervasive monitoring.

The ADD BoF Chairs and the IESG will have to make up their minds after the controversial discussions, which were rather "civilized" compared to the IETF104 side-meeting exchange of views. They will have to consider the following practical questions: should a working group be chartered besides DoH, DNSOP and DPRIVE, and what should its scope be?

During the session, operators called on the IETF to establish an "Operational Concerns" Working Group which, according to Barbara Stark (AT&T) should have a Co-Chair from the operator community. Stark warned that DoH had risen to one of the top concerns with regards to the necessary changes of networks and rising troubleshooting costs.

Many in the HTTP camp questioned the need for an additional working group, underlining that there were already several working groups dealing with DNS. In fact, Stephen Farrell said that the split of DoH and DPRIVE – both preparing their own version of DNS privacy – had been a mistake.

Ben Schwartz (Google) pointed out that there was already a DNS Operator Working Group, the DNSOP Working Group, although this group was prioritising the production of new RFCs instead of focussing on operational issues. IAB Chair Ted Hardie said that a number of BCPs could be developed regarding the process of choosing a resolver and transport.

Since the IETF meeting, the ADD mailing list has gone wild, with a high number of postings.

So far, it remains unclear whether the IETF will be chartering an ADD Working Group or reconsidering how (in how many working groups) DNS work in the IETF should be performed.

How fast is DoH?

With DoH grabbing more and more attention, the first studies on DoH performance are underway. One comparison about the latency of DoH compared to DoT and Do53 (as normal DNS is now called from time to time) was presented during the Applied Network Research Workshop, held alongside the IETF. Depending on the network's capacity the DoH provider, encrypted DNS can be faster than normal

DNS, interestingly with DoT beating DoH in several settings (see graph).

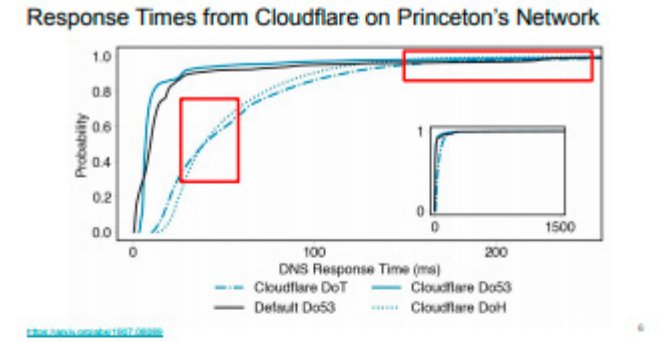

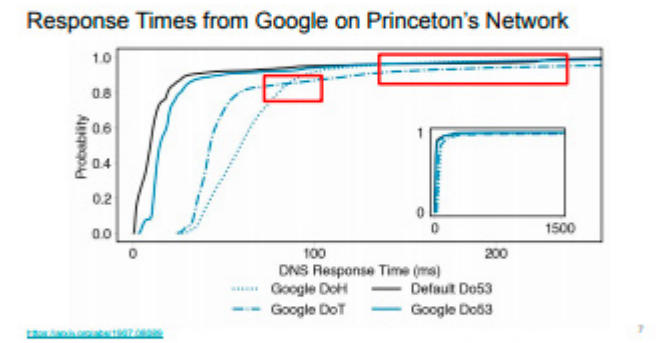

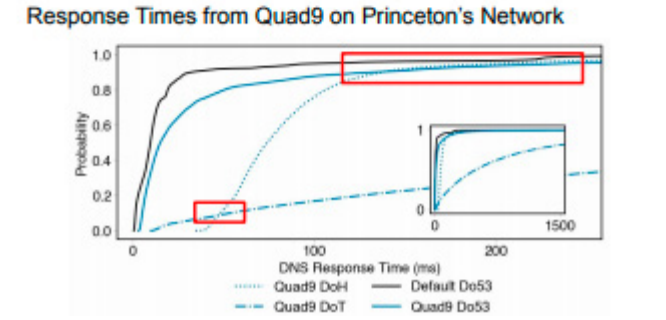

Observers noted that the figures given had to be treated with caution due to the fact that the researchers had done the testing from their University lab at Princeton. A recent check using RIPE Atlas probes and measuring DNS response times from Africa at the Africa Summit found that on UDP, local resolvers returned responded 27 times faster than the Cloud resolvers (Cloudflare, Google, Quad9). Measuring DoT, they found that its response time was three times longer than TCP on the cloud DNS providers.

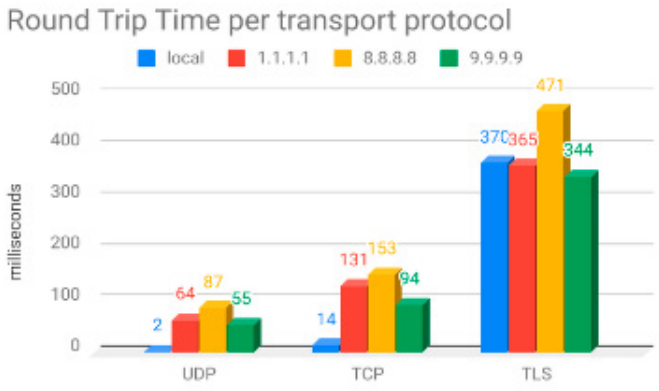

A new threat model: Attacks on end-user devices and systems

Four IAB members, although not speaking for the IAB, called on the community to reconsider the threat model that IETF protocols are developed for. During the Security Area Open Meeting in Montréal, former Security Area AD Stephen Farrell (Trintiy College) and former IETF Chair Jari Arkko (Ericsson) presented some of the ideas.

According to Farrell, for many years, the threat model developers had in mind was that attackers had complete control over the network (path), while devices were expected not to be compromised. Over several years, and especially since the Snowden revelations in 2013, IETF developers have taken communication security up a notch, for example by encrypting transport (packets on the path) through TLS 1.3 or Quic, the new UDP-based transport protocol nearing completion, by encrypting DNS queries (DoT, DoH) and by encrypting the Server Domain Names (ESNI).

Partly triggered by this closing-down and to circumvent encryption, attackers turned to providers to get data – helped by surveillance capitalism – or got their hands-on end-user devices. In his draft document on "Changes to the Internet Threat Model", Jari Arkko warns against continuing to focus on communications security only, as this might "lead to accidental or increased impact of security issues elsewhere".

For example, allowing the collection of data through protocol design even by unsuspicious parties meant that "overtime, unnecessary information could get used with all the associated downsides, regardless of what deployment expectations there were during protocol design." Arkko lists aggregation, consolidation and concentration of traffic (as in DoH implementations) as

some of the issues that require attention: "it is difficult to imagine that DNS resolvers wouldn't be a target in many future attacks or pervasive monitoring projects", Arkko writes in the draft, noting that "there is little that even large service providers can do to refuse authority sanctioned pervasive monitoring."

Arkko envisages the following potential guidelines:

- consider principles in protecting information and systems
- minimize information passed to others
- perform end-to-end protection via other parties
- minimize passing off control functions to others
- avoid centralized resources
- have explicit agreements
- be suspicious toward parties your device connects to
- encrypt everything to everyone

Farrell, who is the author of a second document, has another set of quite interesting recommendations:

- developing a BCP for privacy considerations
- consider not only use-cases, but abuse-cases when developing
- re-consider allowing protocol "extensions"
- think about isolation (to protect against linkability)
- mainstream the transparency concept taken by certificate transparency, pushed through GDPR
- if applications don't see data, it's harder for them to misbehave (minimize)
- same origin policy
- generalize the threat model of OAuth (two of three parties might collude against the third)
- update threat model from end point from not compromised to in general not compromised
- re-visit transport/communication security
- attack recovery must be part of protocol designing
- protocol endpoints might not be "hosts" in the classical sense

A start to these discussions was made at an IAB workshop on Design Expectations and Deployment Realities, which featured quite a number of interesting topics including some related to DNS, namely Andrew Sullivan's "Three kind of concentrations in open protocols".

Farrell and Arkko's call received mixed reviews in Montréal. Several participants underlined the need to focus on implementing the many security and encryption standards passed by the IETF in recent years instead of producing another set of papers. Many pointed to ongoing work, for example the newly-established Working Group on "Remote ATtestation ProcedureS" (RATS), which aims at providing a mechanism to "assess the trustworthiness of the peer", a kind of trustworthiness check-up for remote systems.

However, others welcomed Farrell and Arkko's drafts, namely those working in and around the SMART group (which has not been chartered yet). The SMART people are mainly concerned with how to solve the problems caused by encryption for (security) monitoring. Arnaud Taddei from Symantec advertised his draft on endpoint limitations with regard to security monitoring.

Farrell tried to make clear that the security model considerations must not result in questioning communication security. Instead, encryption of communication protocols had to be further developed.

It is also interesting to note that looking to the security of end-devices is a shift from the usual IETF practice. End-devices (and end-users) are not in the realm of IETF standardization, as Daniel Kahn Gilmore reminded the Security Area meeting participants. It will be fascinating to see if another IETF mantra is unravelled, as recommended by IAB member Mark Nottingham with an informative draft on why the IETF should keep its eyes on the users and decide in their interest in times of conflicting considerations.

IETF in crisis mode

The Administrative Plenary in Montréal was privy to lengthy and very tense discussions over three controversial issues: IETF conduct, LLC budget decisions and the unexpected resignation of the RFC Series Editor.

Once more, IETF Chair Alissa Cooper talked about the IETF code of conduct that has been sharpened over recent years, in part due to rude exchanges on the IETF mailing lists. In several meetings, rough language and aggressive behaviour in working groups and on the mailing lists have been brought up. The

IETF Chair pointed to several recent initiatives aimed at mitigating possible violations of "good conduct", including discussions on Ads and working group chairs, assistance from professional trainers, and the addition of more rotating sergeant-at-arms to step in in case of aggressive behaviour. Following a recent incident involving a sergeant-at-arms' intervention, Cooper also said that a clarification on escalation procedures might be necessary. Some mentioned the possibility of adding mechanisms to treat sexual harassment (proposed by Kristy Paine, National Cyber Security Center). In a rather unusual act, two IAB members, Martin Thomson and Mark Nottingham, apologized for their own violations of good conduct or failure to step in against aggressive behaviour.

An apology was also sent by the RFC Series Oversight Committee (RSOC) just hours before the plenary meeting. It was addressed to Heather Flanagan and concerns her decision not to renew her contract as RFC Series Editor (RSE).

The RFCs are the most important output of the IETF and are published in different streams: Internet standards and Proposed Standards, IRTF and IAB RFCs, and finally independent submissions (see RFC editor site). Since taking the job in 2012, Flanagan had started working on the formatting of the RFCs (see several steps of the process here).

Flanagan announced her resignation after being confronted with the RSOC's announcement to call for bids for the RSE function before the next possible extension of Flanagan's current contract (renewable in 2021 for another two years). According to some opinions (including the RSOC's), she decided to resign because she interpreted the early start of a bidding process as a reflection of her performance. In the apology letter, RSOC Chair Sarah Banks explained that the low number of candidates to earlier bids for IETF functions was the motivation behind that decision and in retrospect, it had been a bad move and obviously badly communicated.

Flanagan herself said to this reporter that her decision was routed in "too many managerial burdens" in doing her work.

Flanagan has overseen the production of the RFC series since 2012. The production is performed by the RFC production center, a group of editors at AMS, which are bound by contracts that are independent from the RSE contracts. AMS is also the operator of the IETF secretariat. In addition to the contractual split, there is also an oversight split, since besides the RSOC (which is nominated by the IAB), there is also the RFC Advisory Group, which according to the site, is nominated by the RSE. In her statement to the mailing list Flanagan pointed out several issues:

"Over the last year, I've seen the rfcplusplus BoF happen, against my recommendation. My oversight committee, which is a group that I must work with most closely, was almost completely replaced without any input from me. I have what essentially acts as a design team, the RFC Series Advisory Group. They generally aren't consulted either. The RSOC/IAB is pushing hard on the missed SLA, not acknowledging that that statements were made (with full support and understanding of earlier leadership cohorts) on plenary stage and in meetings that the SLA would be missed as the format tool testing and transition ramped up."

Due to the ongoing reform and upcoming switch to the new RFC format, the SLA has not been met recently, also leading to a decision by the LLC to add staff to the AMS RFC production center.

IAB Chair Ted Hardie had outlined the IAB's view and steps for the now urgently-needed bidding process here. During the discussion in the plenary, a majority favoured the suggestion to look for a successor urgently, while at the same time revisiting the current construct and sources of tensions.

One of the major sources of tension at the meeting were the widely differing views of IETF participants regarding the RSE. For long-time IETF contributors, the Editor is a member of the community and party to the process – with a somehow independent standing. For the younger IETF participants, many of whom have moved up in the IAB, the RSE is more like an employee. Several members of that group made "get on with it"-comments, while the older participants saw the RSE stepping down as a fundamental failure by the current IETF leadership. Discussions on how to proceed are ongoing, with the format switch and the lag in RFC production pressing.

The LLC was also challenged and got its first "baptism of fire" over budget questions. According to LLC Chair Jason Livingood, in the first four months the LLC Board had prepared to hire an executive director (receiving 134 applications!), had prepared its own procedural policies (which are under consultation here), had established insurance coverages, had to look at the

S export regulations and extended contracts with Secretariat Service and the current Interim Executive Director. It also took over the IETF Endowment from ISOC (around $3 million). Leslie Daigle (as well as Harald Alvestrand) warned against the LLC's micromanagement by revisiting exiting contracts in order to simplify them. The most critical question for the RFC came from Bob Hinden, who called out the LLC for a rise in expenses by $1.6 million in 2019 compared to 2018. Cooper rejected the notion that the rise had not been communicated to the community and said that the budget had risen every year, albeit not by that much each time.

These discussions seem to reveal that the IETF is at a number of inflection points with regard to how the community will govern itself (including considerations about cutting back meetings to meet remotely for ecological reasons). Besides tensions within the community over procedures, there are tensions between established IETF contributors and new groups attracted by the ongoing work or trying to bring new, related work for getting the RFC stamp (see the Media Operations, Mops BoF, the end-user oriented Medup and the law enforcement community oriented Smart side-meetings). New groups often feel unwelcome due to established contributors pushing back – hence the conduct issues. One pessimistic participant reacting to the RSE issue said to this reporter that the IETF had a lot of work to do to ensure remaining in business.

WGs and BoFs

DNSOP Working Group: More straw for the camel, or how to avoid policy-sensitive issues

The DNSOP Working Group took a small hit during the ADD BoF when Ben Schwartz (Jigsaw/Google) pointed out that some aspects of DoH, such as DoH operational concerns, would in fact fit the mandate of the "DNS Operations" Working Group. So far, the Working Group has managed to dodge most of the DoH debate. Now it has one DoH-related draft up for adoption. It is a draft on the self-publication of what kind of encrypted resolution a resolver wants to offer. Another highly-political issue, alternative names, was also briefly discussed in Montréal, with DNSOP Co-Chair Suzanne Woolf considering asking the IESG for a final decision on what to do with the document on alternative names. Once more, the DNS Working Group had plenty of additional straws for the camel during two working group sessions, with a long list of currently-active working group documents and more up for adoption.

Routing loop with alternative naming

In its first session in Montréal, the DNS Operations Working Group briefly dipped into one of the highly controversial issues that had been postponed several times, namely the alternative names issue. It was evoked by a number of applications which were brought to the IETF based on RFC 6761 (Special Use Domains). While some applications, namely .tor, was granted (with Apple's .local being the original use case), a fierce discussion took place on the possibility of the IETF granting rights to more special domains. Many participants felt this was an invitation for parties to circumvent ICANN's gTLD application policies and were afraid the IETF could get into a fight with ICANN over the issue.

A document that tries to "corral" alternative use cases such as .tor (or .home, which recently became home.arpa) into a .alt TLD is now in its 11th version. During the meeting, DNSOP Co-Chair Tim Wiczinski proposed sending the document off to the Internet Area or General Area Ads so they could "beat" on the document some more. But several participants, including co-author Andrew Sullivan (now ISOC CEO), called on the Working Group Chairs to make a decision instead of creating what could ironically be called a "perfect routing loop". Paul Vixie, ISC founder and BIND developer, recommended that the Working Group should document a decision and, if they were to chose not to allow such alternative uses, to list the reasons why.

A little DoH in DNS

A bit of DoH-related work has arrived in the DNS Working Group with a draft that specifies DNS resolver information self-publication, including the necessary "DoH" name-value pair to allow servers to express what they provide to be standardized elsewhere. Instead of limiting this to the information about existing DoH service by the respective (local) resolver, the authors (Paul Hoffman and Roy Arends, ICANN and Puneet Sood, Google) opted to generalize the method, allowing resolvers to announce additional information.

Format options to address the resolver were discussed, with two now left on the table: the in-addr/IPv6.arpa or https:/.well-known/info URI. The idea to use a special use domain, which some participants said would be preferable, would not be pursued, according to the authors. The draft establishes a special new resource record type "RESINFO". The chosen format is I-JSON, as the authors believe that it would be better than JSON for interoperability reasons.

During the discussion, some participants questioned the need for standardizing this. For example, Stéphane Bortzmeyer (Afnic) asked why one should not rely on normal domain provisioning mechanisms that were currently developed in the Internet Area. Others warned against potential gigantic record sizes that might evolve over time, given that the self-publication was extensible. The Working Group still must decide if it wants to adopt this work.

Additional work on DoH was done during the Hackathon. Petr Špaček (CZ.NIC) wrote a DoH proxy implementation on fastcgi, and Witold Krecicki is preparing BIND9 for DoT and DoH.

More decisions to make: HTTPSSVC instead of ANAME?

The long discussion over ANAME or an alternative solution to simplify lookup options to connect to HTTPS URIs took a turn in Montréal with the presentation of a proposal for a new DNS record type, HTTPSSVC. HTTPSSVC records would allow HTTPS

origin hostnames to be served from multiple network services. According to Erik Vyncke (Akamai), co-author with his colleague Mike Bishop and Ben Schwartz (Google), each can be enriched with information about the transport protocol and keying material for encrypting TLS SNI (ESNI). The HTTPSSVC would provide a solution to the inability of the DNS to allow a CNAME to be placed at the apex of a domain name. At the same time, the information in the record would allow for the omission of http bootstrapping for all domains that offer HTTPS and offer clients the opportunity to learn about all alternative services available at the origin before the first contact. According to the author, the record would boost performance and privacy.

For the time being, DNSOP Co-Chair Benno Overeinder said that the ANAME draft will be pursued, but with interest from the browser group (both Erik Rescorla, Mozilla, and Eric Orth, Google, expressed clear interest during the meeting), HTTPSSVC already looks like a "winner". Nevertheless, Orth said that he saw its implementation tied to DoH, as he was worried about sending out an additional query for each request. ANAME could still be an option for the classic DNS.

Some concerns which were briefly discussed again included the potential size DNS answers could be brought up to, which could potentially be used for DDoS attacks (Warren Kumari, Google). Olafur Gudmundsson recommended that subtyping, which is currently foreseen in the draft, should be avoided, and that a format with an extensible string at the end should be used instead.

It is yet to be decided if the draft will be taken on by the DNS, the TLS or by another Working Group.

Close to last call: Terminology and running a root server locally

Drafts that are close to Working Group last calls include the local running of a root server (which can help mitigating privacy risks) and the DNS Terminology document prepared by Paul Hofmann (ICANN). On the former, a final version was promised for the end of the month before the Working Group last call. On the latter, several participants asked to "ship" the "ter"-Version, as it would subject to updates down the road anyway.

Additional drafts

Further work aimed at making DNS cookies interoperable is underway, with the merging of two drafts on DNS cookies. As reported during the meeting by Willem Toorop (NLnet Labs), one of the authors of the upcoming combined new cookie drafts, problems with interoperability from different DNS server software were also addressed during the Hackathon.

Two more drafts that are up for adoption include the draft to avoid IP fragmentation by the DNS, by Kazunori Fujiwara, as well as recommendations for authoritative server behaviour, developed from a study on that issue by Giovanne Moura (SIDN). Two additional drafts that were briefly discussed include the "Related Domains By DNS" draft by Brotman and Farrell, as well as a draft on resource record for transferring covert information from a primary to a secondary DNS server from ISC.

REGEXT: More privacy discussions are needed, EPP-RDAP dual stream

The REGEXT Working Group discussed if it should advance additional RDAP standards or if it should wait for registries and registrars to gain experience with the implementation of the "Whois"-successor protocol. It also continues to struggle with how to address privacy in registration-related standards, and briefly revisited the questions for RDAP and the continued EPP work.

From 26 August, ICANN registrars and registries will have to use RDAP for registration (Whois) data, a fact that made Richard Wilhelm (Verisign) question the need for quick standardization of additional features in the IETF process and a push for quick deployment.

During the REGEXT Working Group meeting, four additional drafts (functions) were presented. The first two were escrow-related. Data escrow is necessary and for ICANN registrars and registries, subject to contractual obligations – it is used by ICANN, its TLD registries and EBERO operators. The two documents are split into data escrow and the special format for domain name registration data (Registry Data Escrow Specification and Domain Name Registration Data Objects Mapping, respectively). According to Francisco Arias (ICANN), both are extensible. The split in escrow and escrow format would allow the generalization of the escrow part and its use for other data sets.

The format document defines the structure for data to be escrowed for the domain name case, including the following objects:

Domain

Host

Contact

Registrar

NNDN (user for reserved domain names, withheld IDN variants, etc.)

EPP Parameters

IDN Table Reference

Header

Policy

The other two documents, authored and presented by Manuel Loffredo (Registro.IT), are related to sorting and paging, as well as the controversial reverse search. The sorting and paging draft intends to help reduce bandwidth needs and response times. Reverse search would allow for the extraction of information about users to find domain names owned by an individual or a company, starting with the details of the owner, such as a name or email address.

According to Arias, while the escrow-related drafts have been around for some time, and an implementation period was recommended by James Gould from Verisign, sorting and paging as well as reverse search are more recent.

On both topics, the issue of privacy was raised. Wilhelm said that since the escrow-related drafts were working with PII, he saw a need for a more detailed discussion on privacy issues. For the reverse search, several people, including Stéphane Bortzmeyer (Afnic), noted that comments in the privacy consideration section had barely been taken on board by the authors. Instead of pointing out that registries/registrars "should" follow their local laws with regard to privacy, they "must" now do this in the new version. Following the law was a given. Additions that were disregarded included explanations about potential "abuse cases" (something that was proposed by Stephen Farrell in the new threat model-discussions – see highlights).

Loffredo argued that reverse search was already possible in the Whois and was commercially offered. He expected that it would soon also be the case for RDAP. However, this would bring the respective providers (those allowing that use, and those selling the reverse search results) in conflict with the EU's GDPR.

REGEXT Chair Jim Galvin (Afilias) agreed that there was a need for additional privacy discussions in the Working Group and privacy considerations in the texts. The Working Group was split on the question of whether it would make sense to push new features without waiting for ICANN policies to shape up, or if standardizing new features should be paused until experience has been gained with RDAP after 26 August.

What could go wrong with RDAP (and other standard implementations)?

Marc Blanchet, Viagenie, traced mistakes in implementing RDAP, presenting a long list of errors found in the field by ICANN gTLD registries/registrars as well as IP registries. In the list, Blanchet found RDAP non-existent or wrong values in the data base, like object truncated due to server policy instead of object truncated due to server authorization. Other mistakes included self-referencing links that created a loop for requests or RDAP servers not accepting percent encoding.

Besides listing such implementation mistakes, Blanchet also made recommendations for updating the RDAP standard to add, for example, a role object for registries and to switch cross-origin resource (CORS) sharing from a recommendation to an obligation (from should to must).

Blanchet asked the Working Group how to proceed, if for example he should name the culprits that had faulty implementations, and if the document should be worked into a BCP document for implementers. Talking to this author just before the IETF, one of the ICANN registrars who had to implement RDAP by 26 August said that he expected some messy weeks after the protocol switch from Whois to RDAP. More work for RDAP standardization lies ahead, not only with regard to the planned new features, but also with possible updates.

EPP – RDAP: one or two streams?

George Michaelson (APNIC) reiterated the question first brought up at IETF104 on whether the community was considering a split of the work for the REGEXT Working Group. Some think that given the upcoming wave or RDAP drafts, an additional Working Group might be handy, instead of keeping RDAP and EPP together.

Ulrich Wisser (Swedish Internet Foundation) reminded participants in Montréal that there were reasons for keeping the two streams together, namely the fact that the group was already small and that a second Working Group would gather the same people.

The only EPP document presented during IETF105 was the Secure Authorization Information for Transfers (presented by James Gould, Verisign). This draft document "defines an operational practice, using the EPP RFCs, that leverages the use of strong random authorization information values that are short-lived, that are not stored by the client, and that are stored using a cryptographic hash by the server to provide for secure authorization information use for transfers." Several participants hinted at their own implementations of such mechanisms.

RFCs published by REGEXT since IETF105 include the Change poll extensions for EPP, EPP organizational mapping (rfc8543/) and Organization Extension for EPP (rfc8544/). Considered by the IESG are Registry Fee Extension for the Extensible Provisioning Protocol (EPP) and Extensible Provisioning Protocol (EPP) Domain Name Mapping Extension for Strict Bundling Registration (Informational). In Working Group last call: Login Security Extension for EPP.

DPRIVE: Privacy work, DoT/DoH discovery questions, private zone transfers

The DPRIVE Working Group continues to consider securing the next resolution step – from DNS resolver to authoritative name server. At the same time, additional ideas for securing zone transfers resulted in criticism regarding over-complicated proposals.

The DNS Privacy Exchange (DPRIVE) Working Group covered some familiar ground, with the document on "Recommendations for DNS Privacy Service Operators" being close to last call. The document lists privacy considerations and threats to the different encrypted DNS specifications (DoT, DoH). According to Roland Rijswijk-Deij, the authors have decided against splitting the guidance for DNS privacy operators and DNS privacy policy and practice statement (DPPPS). The latter will assist operators in preparing their own DPPPS. The Working Group last call is just around the corner. Instead of waiting for additional policies, for example for DoH to be set, the authors intend to keep the document alive with additions and changes to be considered on a regular basis. The authors believe that waiting for DoH providers to come up with policies like Mozilla's Trusted Server Policy would delay the production of the draft unnecessarily.

The top issue for DPRIVE (after the publication of DNS over TLS, DoT) has been to secure the next step, from resolver to authoritative server. Instead of potential next steps in that regard, Tim April (Akamai) very briefly presented a draft (which is co-authored by April, Jason Livingood, Comcast, and Karl Henderson, Verisign) on operational issues with DoT for authoritative server (AdoT). The concerns include performance issues, such as the "unintended side effect" of erasing the signalling of EDSNO Client Submet (ECS) or the need for a resolver to "test" if an authoritative server will be prepared to do DoT. While Tim April underlined that the document was not on the merits of using AdoT, but rather on concerns to be weighed up by those who implement it, the draft makes the following statement: "At the higher levels, techniques such as QNAME minimization and Aggressive Use of DNSSEC-Validated Cache [RFC8198] arguably provide an alternate path toward mitigating the risk of disclosure of sensitive information without the operational risk of DNS encryption." The draft also points to monitoring problems and originally referred to the highly controversial draft by Matthew Green on static TLS keys to allow for network boundary interception.

Four more documents were considered in DPRIVE, two of which are related to DoT and DoH implementation issues, and two more aimed at addressing the issue of cleartext zone transfers.

A document prepared by Alessandro Ghedini from Cloudflare addresses the potential leakage of data when sending so-called early data in DoT implementations. TLS 1.3 allows the sending of data before the TLS handshake is completed. While saving a round trip for completing the handshake, this allows a potential attacker to extract information about these queries. The document proposes mitigation techniques, but still has to be developed further.

Another document that was presented in Montréal by Michael Richardson (Sandelman Software Works), together with Tirumaleswar Reddy (McAfee), Dan Wing (Citrix Systems) and Mohamed Boucadair (Orange), provides mechanisms to bootstrap endpoints to discover and authenticate local DoT and DoH servers.

It intends to allow employees bringing their devices to discover and authenticate the server for their DNS resolution. According to Richardson, another use case would be for IoT environments.

A snapshot of the mechanism can be seen below.

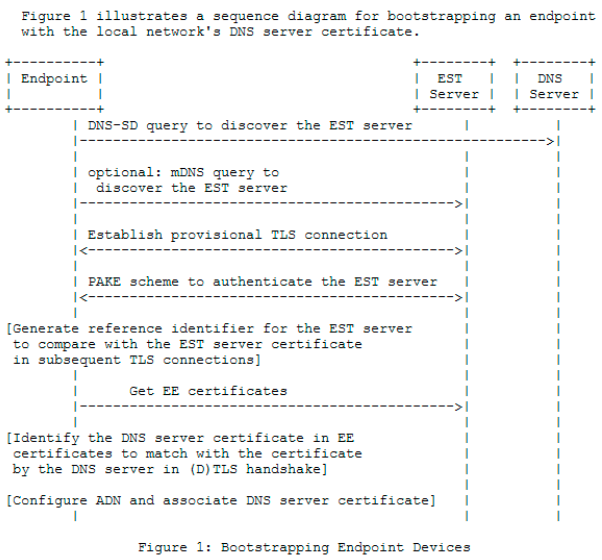

Figure 1: Bootstrapping Endpoint Devices

A highly controversial debate during DPRIVE was spared for one of the two proposals on securing zone transfers.

According to the first of the two drafts (DNS Zone Transfers over TLS), AXFR zone transfers are regularly carried out over TCP; encrypting AXFR using DNS over TLS was therefore straightforward. Besides AXFR (using TCP), the draft also discusses IXFR, which allows both UDP or TCP. Implementations already exist in Unbound Release 1.9.2, so according to some participants, the draft documents current practice. During the IETF105 Hackathon, XOT support for the DNSjava library was coded. While there are several questions, for example regarding the need to make TLS 1.3 mandatory (instead of earlier TLS versions) or if padding was necessary, Working Group participants seemed to welcome the draft.

The second draft was not as well-received. It introduces DNS stateful operations to allow for secure zone transfers without the need for Notify and State of Authority (SOA) interactions upfront. The mechanism looks rather complicated, and not only for the outsider. One participant at the meeting called it a "weird mix of things". Petr Špaček from CZ.NIC was vehemently opposed to the concept. Some also thought the DPRIVE Working Group was not the right place for this topic, so the Working Group Chairs will make a decision on what to do with the proposal, which could include sending it to DNSOP.

BoFs on LAKE, MOPS and barriers to newcomer integration

Several BoFs once more illustrated the interest of the IETF in integrating both new technology and new groups of participants, while at the same time being conservative in how to go about it – and showing a preference for IETF standard suites over alternative ones.

For example, in the Lightweight Authenticated Key Exchange (LAKE BoF), a discussion started on whether a new TLS-like key exchange protocol could better suited to the growing number of IoT environments or if a TLS 1.3 derivative would do the job. Göran Zelander (Ericsson) proposed Ephemeral Diffie-Helman over Cos (EDHEC), arguing that EDHEC could add perfect forward secrecy to the newly-standardized Object Security for Constrained Devices (OSCORE).

During the session, former Security AD Erik Rescorla (Mozilla) offered a slim version of TLS 1.3 instead, arguing that the IETF had put a lot of effort and time into developing TLS 1.3 and that its security properties were well tested. At the same time, he said that there was a lot of overhead that could be cut out from the full version, especially the setting default and avoiding version negotiations for TLS 1.3.

In the end, the BoF was split into two camps over which path to take. Some participants like Elliot Lear warned against not choosing one candidate, having experienced delayed decisions in Working Groups like TCPINC some time ago. Others pushed for allowing both solutions to be developed, arguing that CTLS and EDHEC had different characteristics, with CTLS possibly being too resource-demanding for IoT nodes.

Another example of a clash of worlds could be observed in the Media Operations BoF (MOPS). Glen Deen (ComcastNBC Universal) noted that the video experts sometimes had operational issues with IETF RFCs, on which standard suites such as SMTPE 2110 were highly dependent. SMTPE 2110 defines the use of IP networks for professional video production. The

MOPS proponents therefore asked for a place in the IETF where they could bring operational questions and concerns. While the BoF proponents brought a Taxonomy draft to illustrate their need for better understanding and cooperation between the two spheres, IETF old-timers said that the idea of simply creating such a place for a specific community was not in line with usual IETF Working Group modalities. In what could be seen as an attempt to be IETF-like, the proponents presented a document on taxonomy.

Yet another group is currently struggling with IETF/IRTF polices. Once more, the Stopping Malware and Researching Threats group (SMART) met at the sidelines. The group wants to be chartered as an Internet Research Task Force Group, yet their proposals so far seem to include as much protocol as research work, said Colin Perkins, Chair of the IRTF. An original interest of the group is to work on the consequences of encryption for monitoring, even if that particular topic has been erased from the Charter. If the group manages to get chartered, it will be open. Documents on the agenda of the group, on Capabilities and limitations of an endpoint only security solution (CLESS) and – as in the MOPS group – on a related taxonomy document were discussed in a second side-meeting. Other avenues could be considered to feed these documents into the IETF/IRTF process.

There seems to be a fundamental dilemma for the IETF: should it welcome new groups, especially such as the two latter ones, in an effort to present itself as open (and therefore fostering additional attendance), or could such efforts blur the original IETF mandate?

IETF106 will be held in Singapore on 16-22 November 2019.

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 a country code top level domain registry.

Rate this CENTR Report on IETF105

(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