



**Council of European National
Top-Level Domain Registries**

Report on IETF102

Montréal
14-20 July 2018

Contents

Highlights **3**

From DNS to DNS over TLS and DNS over HTTPS: on to resolver-less DNS?	3
IASA 2.0 and some IETF soul-searching	4
REGEXT – Talking RDAP search to fill in for suspended WHOIS?	6

Working Groups **8**

DNS – Is “business as usual” good enough?	8
DPRIVE WG: Re-chartered, authoritative not yet reached	9
MAPRG: DNSSEC roll-over issues and other DNS measurements	10
ANRW: DNS, leaky certificates and sea-level rise	10
TLS: Encrypting SNI	11

Highlights

From DNS to DNS over TLS and DNS over HTTPS: on to resolver-less DNS?

The draft RFC for DNS over HTTPS (DoH) is just about to reach the IESG (currently AD draft) for final review and IETF last call. Now there are some interesting questions before the DNS and HTTP communities: will DNS, DNS over TLS and DoH live side by side? More critically, will DNS become to a large extent a web application? Or, as those most pessimistic ask: will browsers and CNDs even get a say over what TLDs should be resolved at all? And will the web win over the internet in the end?

The IETF in Montréal did not see another DoH WG meeting, but a non-WG BoF meeting was held on “DNS resolver identification and use” (DRIU) and two unrecorded Bar BoFs took place on resolver-less DNS, SRV and HTTP that discussed related issues.

The most controversial discussion developed around Mark Nottingham’s DoH Digest draft, which was presented in the DRIU BoF. Nottingham’s idea basically aims at changing the current DoH model with regard to the pre-configured DoH servers. Instead of having only one DoH server configured – this is the model currently implemented by Mozilla/Cloudflare – clients should be able to choose from several servers.

Mozilla describes the Mozilla-Cloudflare contract as a combination of trusted resolver plus DoH. HTTPS-embedded and encrypted queries arrive from Firefox users at Cloudflare, which leverages its large DNS infrastructure to shuffle back the DoH answers to users. Such an outsourcing of DoH provision has been anticipated by some.

DoH Digest

Nottingham reiterated the advantages of DoH from the web community point of view: privacy and reliability by encryption and taking one party (no going out to third party DNS operator). Performance could also be upped through combining HTTP Client and DNS Client and by using the information in the DNS request stream to aggregate all of its traffic into a small number of connections (possibly only one), thereby allowing greater coordination of congestion

control and avoiding connection setup costs.” Nottingham also envisages the potential use of secondary certificates (from httpsBis).

The Digest idea, according to Nottingham, basically means “that each DoH Server that the client configures would send a digest to the browser (through some means TBD), and that digest would be used as a hint as to the requests that that DoH server would like to see. [...] If (a browser-like) Mozilla were to support this, your browser would be configured with some number of DoH servers; when it connects to each, it would get a digest, and use that as input to its decision about which DoH server to use for a given request.”

Nottingham called the proposal a first strawman. He also acknowledged the risk of concentration, which was highlighted during the BoF by former IETF Chair Jari Arkko. Nottingham wrote in an email to this author: “There’s a trade-off here. DoH works best for anti-censorship when it’s co-located with a very popular Web server -- e.g., a large Web site or a CDN. However, we don’t want to give large Web sites or CDNs more advantage over small sites, so it could be that we try to come up with another mechanism that doesn’t require pre-arrangement with the browser. That hasn’t been discussed yet, however.”

In a reaction after the BoF, IAB Chair Ted Hardie called the incentive for selecting from a choice of DoH servers “poor” with the local device already having a good upstream with an established TCP/TLS/HTTP session “and it’s going to want to avoid the latency of establishment of load balancing in a lot of cases.” Hardie, contrary to other Google colleagues, said that to attract query traffic bound for its network a DoH server needed “a big cache, good connections to other DNS services, the lot.”

“I think the end game of this model is that the user has no control over where the queries go and the heuristic system underneath them ends up sending them to the site willing to offer the highest number of names (more specific routes) and the biggest DNS query infrastructure. That’s going to

land everything behind a few CDNs, unless I miss my guess”, Hardie said.

DNS over HTTPS overtaking DNS over TLS

Despite concerns over DoH and the Digest idea, DoH currently seems to be the technology moving quicker than DNS over TLS (DoT). The latter, already standardized, would have been an alternative more in line with the traditional DNS system and would come with at least the option for decentralized DNS operations (even though concentration in 8.8.8.8, 1.1.1.1 and 9.9.9.9 is progressing).

A reasoning behind DoH being preferable from a Web developer perspective has been explained by Patrick Mc. Manus from Google in a blog post in May:

“DoH implementations, by virtue of also being HTTP applications, have easy access to a tremendous amount of commodity infrastructure with which to jump-start deployments. Examples are CDNs, hundreds of programming libraries, authorization libraries, proxies, sophisticated load balancers, super high volume servers, and more than a billion deployed Javascript engines that already have HTTP interfaces (they also come with a reasonable security model [CORS] for accessing resources behind firewalls). DoH natively includes HTTP content negotiation as well – letting new expressions of DNS data (json, xml, etc.) blossom in non-traditional programming environments.”

Sara Dickinson, Sinodon, shared insights on the race between web community and DNS community for an interview [posted on the CENTR blog](#). During IETF, Dickinson gave a short version of her talk given at ICANN’S Global Domain Division Summit in which she gives an overview of the DoH/DoT race.

Dickinson also authored a new document calling for a special profile to use DoH in order to mitigate DoH-related privacy risks. New privacy concerns, she writes, result from the mere fact that a new transport (compared to DNS over UDP, TCP or TLS [RFC7858]) includes client identifiers (e.g. user-agent, accept-language) not present in any existing DNS transport. How best to mitigate this is still under debate. While Dickinson recommends, e.g., that “DoHPE clients should send queries over connections used solely for DoHPE (‘dedicated DoHPE connections’) to avoid mixing with other HTTPS traffic that might contain HTTPS messages with client identifiers”, other engineers point out using dedicated connections might help traffic analysis instead.

DNS and Web communities – to cooperate or to compete?

Additional cooperation between Web and DNS communities was called for at the DNS WG (see below) and at the two Bar BoFs. One was a discussion on SRV and HTTP on service location. CNAME used for using bigbank.example instead of www.bigbank.example was a stretch, according to experts, because CNAME changes direction of DNS lookups and frequently leaves the zone. The idea to use SRV records instead to express which servers will provide a site was decided to be suboptimal as well due to wildcard issues. Now Web and DNS people are looking for other potential candidates to solve the something@apex issue. A dedicated mailing list will be established by the DNSOP WG Chairs. According to Olafur Gudmundson, one candidate might be a “HTTPS RRtype that has the servers and SNI+KEY’s information in the record as that record can be added to”.

The second BoF meeting where web and DNS experts gathered was a meeting on “resolver-less DNS”. According to observers, it touched mainly on “using off-network resolver over random protocol instead of the network-provided DNS resolver”.

According to Dickinson, the DNS people had somehow brought the DoH development upon themselves ([see interview](#)) by not addressing some of the issues the web community wanted to see solved. For the DNS, the major shift toward Port 443 could result in a considerable change of their environment.

IASA 2.0 and some IETF soul-searching

The IETF is advancing its new administrative structure and is becoming a little bit more independent. However, concerns remain about how to continue to attract enough engineers to participate – and thereby to also raise enough money to fund the IETF activities, the secretariat, meeting expenses and RFC series. An IAB-initiated debate about the future of the RFC series under the title RFCplus shed some light on more soul-searching.

IETF LLC

During the plenary meeting in Montréal, IETF Chair Alissa Cooper announced that by the end of August,

the IETF hopes to have officially set up a new Limited Liability Company. The IETF Administration Limited Liability Corporation (“LLC”) will become the “corporate home for the IETF, the IAB and the IRTF”, said Cooper. It will allow the IETF to independently contract with the secretariat operators, meeting hotels and to hire personnel.

With the so-called [IASA 2.0](#) reform, the IETF will become a “disregarded entity” of ISOC, sharing ISOC’s tax-free status (as a 503c non-for-profit organisation according to US law), residing in Delaware, US. For people interested in the details, a comment period on the legal documents will be open for a short period of time in August, according to the IETF Chair.

The links between both IETF and ISOC were said to remain close. ISOC reserves some rights, especially to approve:

- amendments to the LLC Agreement
- fundamental and material changes to the nature of the LLC’s activities
- material change in accounting or tax policies previously agreed
- admission of new members, mergers, sale of all LLC assets, etc.
- conversion of LLC to another form of legal entity

In an effort to allow a swift transition to the new structural set-up, ISOC contributed an annual [extra funding of USD\\$5 million](#) for the coming three years, in addition to the regular USD\$2 million annual contribution. In addition, money from the IETF Endowment will be transferred (USD\$2,6 million in 2018).

	2018	2019	2020
Revenue			
ISOC annual contribution	\$2,692	\$5,000	\$5,000
ISOC in-kind	315	0	0
Meeting revenue	3,909	4,154	4,219
In-kind revenue	113	35	35
Total	\$7,029	\$9,189	\$9,254
Expenses			
Meeting expenses	\$3,089	\$3,172	\$3,215
RFC services	1,238	1,262	1,198
IETF secretariat	1,375	1,404	1,436
Operating costs	670	1,564	1,795
ISOC support	315	0	0
Transition costs	75	0	0
Special projects	50	50	50
Tools	217	221	226
Total	\$7,029	\$7,674	\$7,920
Net surplus		\$1,516	\$1,334

From IASA 1.0 to IASA 2.0 – The new administrative structure

The Montréal meeting saw the last gathering of the IETF Administrative Oversight Committee (IAOC), the very body established during the first IETF reform. Initiated by then IETF Chair Harald Alvestrand (then Cisco, today Google), IASA 1.0 (established in 2005) resulted in the first formalization of the administrative work, the introduction of the position of the IETF Administrative Director (IAD) – and some time later, the IETF Trust.

[With IASA 2.0](#), only the IETF Trust will remain [\(with some minor adaptations\)](#). The LLC is generally tasked with supporting the ongoing operations of the IETF (meeting and non-meeting activities), managing the IETF’s finances and budget, fundraising and compliance (“establishing and enforcing policies to ensure compliance with applicable laws, regulations and rules”).

The LLC Board will take-on the oversight role and hire the LLC executive Director, who in turn will perform the day-to-day operations (including hiring additional staff members for a number of things, including fundraising, outreach and communication).

During the IASA 2.0 BoF session in Montréal, a lengthy discussion was entertained about the role of the LLC Board and how many members the new oversight body should have. Too large a board created a risk for mission creep, some participants warned. In the end, the result was a fixed number of five members including: the IETF Chair, one ISOC Board of Trustees’ member and three appointed by the IETF NomCom (call will go out on 16 August). Furthermore, the LLC Board itself can opt to choose up to two more Board members on their own.

New meeting fees

Dwindling participant numbers remain a concern for the IETF. With 1,020 attendees, the Montréal meeting was smaller than the Prague meeting one year ago. IETF Chair Alissa Cooper announced a rise in fees, beginning with those registering late:

Fee Type	Due Date	Deadline	Amount (USD)
Early Bird	17 Sept 2018	7 weeks prior	\$700
Standard	22 Oct 2018	2 weeks prior	\$875
Late & On-Site	9 Nov 2018	≤ 2 weeks	\$1000

Another experiment planned for the Bangkok meeting is the reduction of the meeting days to four (Monday to Thursday, or six for those attending the highly-popular Hackathon meeting the weekend before).

IETF evolution and soul-searching – The debate about RFCplus

More soul-searching came at another BoF initiated by the IAB over potential changes to the RFC publication series. In an attempt to avoid confusion about the RFC “brand”, the IAB proposed an experiment to only call IETF standards RFCs in the future and find other labels for documents of the IAB, the IRTF and the so-called independent submission stream.

With this change, outside observers and users of the RFC series should be prevented from mixing up IETF standards that went through the regular IETF standards peer review procedures with documents by individuals. The proposal was flatly rejected by an overwhelming majority, with many participants calling the BoF unwarranted and the lack of inclusion of the current RFC editor, Heather Flanagan, “shameful”.

There were some voices arguing that the inability of outside users to differentiate between the RFC and an internet standard caused problems. For example, the RIPE NCC has recently been debating draft documents with regard to IPv6 numbering, promoted by the respective authors as IETF RFCs while not having made it through the official standards process. While not urgent, this was an issue, said Marco Hogewoning (RIPE NCC). Misrepresentation of individual submissions or informational documents could lead to problems, for example when implementation was called for by some outside body or industry.

Nevertheless, the majority warned against closing the other streams and hand control of RFCs over to the IESG exclusively. The age-old problem of “not every RFC is an internet standard and not every internet standard is an RFC” could either be addressed as an educational problem or perhaps through innovative formatting. Changes were complicated, some pointed out, including with regard to the IPR statements, disclosures and licensing.

A deeper-rooted problem, several participants warned, was losing quality in the documents. Some parties would just push to get an RFC in any of the streams, and then capitalize on the public confusion

over the status of the documents – standard or just informational submission not vetted by IETF WGs, said Area Director Mirja Kühlewind (ETH Zürich). She called on the IETF community to think about this kind of abuse of the system and do something about it. After the highly contentious debate, the BoF closed with only one next step, which was to request data on the “confusion problem” from Flanagan.

REGEXT – Talking RDAP search to fill in for suspended WHOIS?

The Registration Protocols Extensions Working Group may well be one of the groups that could be questioned with regard of their processes to hammer out RFC standard documents in numbers without much involvement with the IETF community. Being basically dominated by one large registry (VeriSign, with a small number of contributions from three or four other registries), one sole registrar (GoDaddy) and ICANN, the WG is always in short supply of reviewers for the continuous stream of special EPP extensions for the ICANN-Registry-Registrar industry. The set of materials posted for WG members is massive, containing 346 pages of draft documents (of which not all are new, yet most remain work-in-progress). The volume of proposals alone sometimes allows, it seems, to push through items favoured by a tight-knit group of experts without much involvement of a broader IETF community.

During the second meeting, WG Chair Jim Galvin, being the de-facto sole chair at the meetings (with Co-Chair Antoine Verschueren usually participating only remotely) acknowledged that the RegEXT community was a very small community and document “shepherds” (who cannot be authors themselves) are in high demand. At the same time, the WG is about to re-charter, giving the group leeway to take on broader work items. Concerns about the potential overly broad scope resulted in the addition of one half sentence in the new WG charter to ensure a check against potential mission creep: “The working group may also, in consultation with its responsible area director, take on work related to the operation of Internet identifier registries, beyond the EPP and RDAP protocols”. A pretty large number of WG documents have actually reached IETF last call:

- [Allocation Token Extension for the Extensible Provisioning Protocol \(EPP\)](#),
- [Registration Data Access Protocol \(RDAP\) Object](#)

[Tagging](#),

- [Extensible Provisioning Protocol \(EPP\) Organization Mapping](#),
- [Organization Extension for the Extensible Provisioning Protocol \(EPP\)](#)

Some are post WG last call and close to IESG review:

- [Extensible Provisioning Protocol \(EPP\) Domain Name Mapping Extension for Strict Bundling Registration](#)
- [Change Poll Extension for the Extensible Provisioning Protocol \(EPP\)](#)
- [Registry Fee Extension for the Extensible Provisioning Protocol \(EPP\)](#)

With these completed, and the WG rechartering expected to be accepted by the IESG, the WG could be looking for new work items, Galvin said. Several candidates presented their draft proposals in Montréal.

Galvin said that one of the candidates that could move up right away was a proposal presented by Roger Carney from GoDaddy which seeks automation (as much as possible) for registrars in their effort to onboard new registries. Processing questionnaires for the new registries about their handling of shared registry system aspects and EPP plus extensions took his company six weeks, Carney said. With the Registry Mapping proposal, a reduction for the checks shall be formalized, brought down to two pages and allow to deal with 80 percent of the questions in five minutes, before addressing the left-over 20 percent that is not automatized as easily. The Registry Mapping draft proposal is [here](#). Scott Hollenbeck announced his company had filed an IPR claim on technology in the draft.

While attendees pushed back against another proposal by Verisign's Jim Gould on the issue of servers and clients not being in sync about a common set of supported EPP features, Gould nevertheless announced he would come back with a draft proposal for the WG to consider.

ICANN GDPR temporary specification as base for making RDAP search obligatory?

For ICANN, Francisco Arias proposed a number of work items for the WG to take on. During the longer working session on Monday, Arias [presented](#) ICANN's

request for extended RDAP search functions, namely the following options that are currently not yet supported by RDAP:

- Partial match supporting leading "wildcard"
- Support for multiple occurrences of the "wildcard"
- Support for logical operators "AND", "OR", "NOT" to join a set of search criteria at client request
- Explicitly specify the search-pattern parameters to be used with each object type search
- Internationalization improvements

Interestingly, Arias pointed to the recent temporary specification for the GDPR as the document requesting that RDAP search was offered by registries and also registrars (besides webwhois). Scott Hollenbeck offered existing work on using regular expressions as a pathway to mightier search routines in RDAP.

One participant questioned the link between ICANN's temporary specification and the change from Websearch to RDAP search, while Arias pointed to an annex in the temporary specification, which he said will enforce the switch for the ICANN-contracted parties. Concerns on the potential abuse of searches ("mainly used for data mining") were raised by Andy Newton (ARIN).

Several participants also questioned a second set of proposals from ICANN's experts, covering data escrow.

- [Internet Domain Registry Data Escrow specification \(draft-arias-registry-data-escrow\)](#)
- [Registry Data Escrow Specification \(draft-arias-noguchi-registry-data-escrow\)](#)
- [Domain Name Registration Data \(DNRD\) Objects Mapping \(draft-arias-noguchi-dnrd-objects-mapping\)](#)

With the implementation of privacy regulation (GDPR) just underway, the attempt to move the technical parts forward might be unwise, according to Roger Carney (GoDaddy) and Richard Wilhelm (Network Solutions). Scott Hollenbeck called for insights on how much the data escrow standard proposed was also in use by ccTLDs and thereby "global", otherwise he proposed for the potential document to be informational instead of following the standards track.

With regard to the choice between the informational or standards track, the proponents usually all call for a standards track in the RegExt meeting – just neatly highlighting the issue discussed during the RFCplus

BoF. For more details on the RDAP session, see the [Chair's slides](#) and the minutes posted for [one of the sessions](#) so far.

Working Groups

DNS – Is “business as usual” good enough?

With DoH advancing and pressure on the DNS community coming from their Web colleagues, one might wonder if the DNS WG will end up having a fundamental discussion on the future of the DNS in more general terms.

Something(CNAME)@APEX?

According to some, joint gatherings of DNS and Web communities are in urgent demand due to developments like DNS over HTTPS, DoH Digest, but also SRV at HTTP (which met for a Bar BoF – meeting documented [here](#)). DNSOP Co-Chair Suzanne Woolf made it clear that the Chairs considered this work out of scope of the DNSOP WG.

Yet the pressure on DNS from the Web was illustrated during the DNSOP meeting in the debate on something@Apex. WG Co-Chair Tim Wicinski appealed to the DNSOP community to come up with some sort of solution to allow a CNAME-like resolution (like the redirection of example.com to www.example.com). Wicinski pointed out that a number of cloud providers had it, even if proprietary solutions were used at the moment. For example, Amazon is offering it, but since it is breaching existing standards, it only works when they know the target (because it is in their own customer data base).

While Wicinski just asked for a solution that would also allow smaller operators to do it in a standard honouring way (and not only leave it as a trick to the big ones), there was considerable push-back about what problem would be solved. As discussed at the side-meeting, SRV would be a cleaner solution, despite having some problems (corner cases, for example), said Stephane Bortzmayer. It was a little like solving a problem another problem had, Wes Hardacker and Joel Jaeggli said.

Experiments putting either CNAME or CNAME plus DNAME together in the Apex were held during the Hackathon by Willem Torop (NLnet Labs) and Ondrej Sury (ISC). Currently, according to RFC 1034, CNAME@Apex is not allowed: it disallows to have CNAME RR alongside other data and it works badly. However, the CNAME plus DNAME version worked better (see results in the graph). Sury's old draft on this version is [here](#).

CNAME + DNAME @ PARENT Results

DNSSEC Validation Enabled	No QNAME Minimization	Relaxed QNAME Minimization	Strict QNAME Minimization
BIND 9.11.4	OK!	N/A	N/A
BIND 9.12.2	OK!	N/A	N/A
BIND 9.13.2	OK!	OK!	OK!
PDNS Recursor 4.1.3	DNAME fails [2]	N/A	N/A
Unbound 1.7.3	OK!	OK!	OK!
Knot Resolver 2.4.0	N/A	Mixed [1]	N/A
Google Public DNS	OK!	N/A	N/A
Verisign Public DNS	OK!	N/A	N/A
Quad 9	DNAME fails [2]	N/A	N/A
Cloudflare 1.1.1.1	N/A	Mixed [1]	N/A

1. DNAME returns SERVFAIL "AND" Correct Resource Records
2. PowerDNS 4.2 has some DNAME fixes in the roadmap

Some were complaining that the DNS people had not moved for 20 years and needed to be pushed. Another meeting in Bangkok might follow an interim meeting that Wicinski said he would organize.

Eat those (DNS) cookies!

The other controversial subject discussed in DNSOP was on the DNS cookies. Ondrej Sury and Willem Torop proposed to harmonize cookie implementation through an RFC document standardizing cookie crypto and other operational details. Currently, multi-provider as well as same-server discrepancies would result in problems with different versions of DNS cookies being used. Even large servers like K-Root would switch between various cookie implementations, also due to anycast situations. The

proposal by Sury and Torop is to make a number of features mandatory to allow for harmonized cookie production and consumption. This instigated calls to not use cookies anymore, and more drastically to “kill the cookies, use TCP (for stateful transport)!” (Olafur Gudmundson, Cloudflare). However, a number of experts underlined that as long as UDP was an option, it had to be secured and cookies were an option standardized in [RFC 7873](#).

For the time being, the DNS WG still has quite a number of documents on its agenda. Relatively uncontroversial is the update of the crypto for DNSSEC keying and validation (see graph).

DNSKEY Algorithms

Mnemonics	DNSSEC Signing	DNSSEC Validation
RSAMD5	MUST NOT	MUST NOT
DSA	MUST NOT	MUST NOT
RSASHA1	NOT RECOMMENDED	MUST
DSA-NSEC3-SHA1	MUST NOT	MUST NOT
RSASHA1-NSEC3-SHA1	NOT RECOMMENDED	MUST
RSASHA256	MUST	MUST
RSASHA512	NOT RECOMMENDED	MUST
ECC-GOST	MUST NOT	MAY
ECDSAP256SHA256	MUST	MUST
ECDSAP384SHA384	MAY	RECOMMENDED
ED25519	RECOMMENDED	RECOMMENDED
ED448	MAY	RECOMMENDED

A number of drafts that have arrived at the IESG table (past WG last call) include:

- draft-ietf-dnsop-kskroll-sentinel
- draft-ietf-dnsop-terminology-bis
- draft-ietf-dnsop-attrleaf
- draft-ietf-dnsop-attrleaf-fix
- draft-ietf-dnsop-isp-ip6rdns

One of the documents that was discussed for quite some time, the 5011 key-roll-over security considerations, is still waiting for expressions of support or objection from the WG to proceed. During the meeting in Montréal, Mike St. Johns from the IESG said that while making sense in its current form, the document was more harmful than useful. Author Wes Hardacker said that he did not intend to go back to make further changes: he asked the WG to decide on how to proceed.

New work put on the table of the DNS WG in Montréal that looks like it will be going ahead is a [multi-provider solution for DNSSEC](#) (WG adopted the document) and potentially a way to flag that a

parent will only delegate and not do signing for its child. According to Author Paul Wouters (Apache), this would allow for transparency on potential attempts of parents to DNSSEC re-sign child zones.

Wouters also presented a [proposal to avoid potential DNSSEC chain extension downgrade attacks](#). With keys residing in the DNS (DANE Authentication of a TLS server), there was a need for additional securing against downgrade attacks. Wouters said the TLS WG and DNS WG could consider the following answers:

- Do nothing
- Fix everything in new TLS extension
- Two zero bytes in this RFC, specify non-zero semantics in a separate update RFC
- Two byte TLS extension pin TTL (in hours)
- Variable-length (0..255) reserved field (default empty) in this RFC, syntax and semantics in - separate update RFC
- Nested extension block (just like new TLS extension, but even more complicated)

For additional notes, see the extensive [minutes](#) here. The DNSOP WG has a new, third Co-Chair in Benno Overeinder, NLNet Labs to shoulder the workload.

DPRIVE WG: Re-chartered, authoritative not yet reached

The DPRIVE WG has re-chartered since IETF101 and according to the new charter, has the following scope:

1. providing confidentiality to DNS transactions between Iterative Resolvers and Authoritative Servers,
2. measuring the efficacy in preserving privacy in the face pervasive monitoring attacks
3. defining operational, policy, and security considerations for DNS operators offering DNS privacy services. Some of the results of this working group may be experimental.

The extension of DNS over TLS from the stub-resolver to the resolver-authoritative viewed as the immediate next step after securing the stub to resolver part was not discussed as planned in Montréal. Instead, due to time constraints in Montréal, Co-Chair Tim Wicinski announced an interim meeting later in the summer which will discuss on how to proceed with securing the lower branches of the DNS tree. The document by

Stephane Bortzmeyer has been around for some time but has not garnered much discussion.

The parallel development of DoH might certainly have resulted in a slow-down of DNS over TLS developments. Deployment figures of DNS over TLS after four years still seem to grow very, very slowly – despite the availability of the option in most of the DNS open source software.

Brian Haberman, Co-Chair of the DPRIVE WG, reported from a RIPE Atlas measurement campaign that out of 3,659 unique DNS servers queried (40,841 total queries), 61 DNS servers responded over TLS. A mere 1.67 percent of servers supports DNS over TLS. “Measurements from some known DNS-over-TLS-capable servers failed due to TLS capability mismatch”, Haberman also added.

Addressing the third point on the new charter (“defining operational, policy and security considerations for DNS operators providing DNS privacy services”) is Sara Dickinson’s co-authored draft for a best practice document (BCP) on operational guidance for DNS privacy services. The document gives an overview on the various options now available for encrypted transport (DNS over TLS, DNS over HTTPS), discusses the various features and provides assistance to operators in writing up a document on their operational practices with regard to privacy, a DNS Privacy Policy and Practice Statement (DPPS). The DPPS can be published to allow clients to evaluate the DNS operator’s privacy policies. The opinion in the room favoured keeping the two parts (privacy features and DPPS) in one document.


The BCP also lists privacy policies of major DNS operators (Quad9, Cloudflare, Google and OpenDNS) and points to a [full table including small operators here](#).

MAPRG: DNSSEC roll-over issues and other DNS measurements

DNS was also a key topic in the Montréal session of the Measurement and Analysis for Protocols Research Group. With the trust-anchor singling-out RFC 8145, experts are trying to get to the bottom of the slow uptake rate of the 2017 KSK. Last fall, the slow uptake and lack of insight into the situation drove ICANN to


stop the planned roll to the new key. Wes Hardacker presented a [case study](#) illustrating the problem of one VPN operator.

RFC8145 Measurements of DNSSEC KSK Trust



When looking at the number of sources signalling they only had the old KSK at the beginning of 2018, Hardacker and his co-authors found that instead of falling, the number had been rising at the beginning of the year at the B-Root server. Checking on relevant sources, Hardacker found that 63 percent of sources sent only very few (one or two) queries per month. A quarter of the checked queries went to a VPN provider. By contacting the respective provider, Android/OS Software updates could be triggered to mitigate. Hardacker concluded that flag-day changes were hard (the KSK roll-over is now set for 11 October 2018) and that software should include automatic DNSSEC key updates in the first place.

Impact of This Effort



DMAPI – Easier, unified measurements of DNS

More DNS-related work included the presentation of a new tool to map DNS properties. DMAPI, the domain name ecosystem mapper, automates measurements of five protocols: HTTP, HTTPS, DNS, TLS and SMTP. It allows a unified view via an SQL interface. The SIDN-supported tool can be tested and used by everybody [here](#).

Caching and retry effects during DDoS

A study on caching and retries quantified the number of queries answered during DDoS attacks under different TTLs. The study concludes that together, caching and retries allow up to half of the clients to tolerate DDoS attacks that result in 90% query loss. Almost all clients can tolerate attacks resulting in 50% packet loss. Tail-latency increases for clients during attacks. According to the results for servers, retries increase normal traffic up to 8 times.

ANRW: DNS, leaky certificates and sea-level rise

For the first time, the [Applied Network Research Workshop](#) (ANRW) was held alongside the regular IETF meeting instead of the week before. Both scientists and IETF participants welcomed the scheduling, as it allowed for the two communities to be brought together.

ANRW Co-Chair Sharon Goldberg called on the academic community not to expect their ideas and proposals to simply be taken up by engineers and become RFCs. Goldberg gave a little how-to-guide for those interested in succeeding to have their drafts passed by the IETF. The substantial sessions covered the topics of TLS, routing, infrastructure and anonymous communication.

DNS was the topic of one of the invited talks in which Mark Allman (International Computer Science Institute at the University of Berkely) presented figures on the progressing concentration in DNS. Addressing robustness, defined in RFC 1034 as having two authoritative nameservers per SLD and having them geographically distributed according to RFC 2184, Allman found that a growing number of SLDs did in fact fulfil the minimum standards set in the RFCs. Digging deeper, though, he also could measure that 20 percent of all SLDs are served from

only 19 networks (numbers 1 and 4 in the graph are Cloudflare, for example).

Allman called this at least an “unhealthy habit”. He agreed that there was a need for additional research on concentration of DNS infrastructure in addition to the trend towards using just a number of large cloud providers and big hosters. Anycast usage, he said, would regionalize, but not necessarily completely solve the robustness issue (being able to always find an authoritative server for a queried name).

Allman was invited by Ondry Sury (ISC) to the next DNSOARC meeting in Amsterdam in October (alongside the RIPE meeting).

TLS 1.2 and leaky client certificates

Other highly interesting talks touched on the data leakage through TLS 1.2 in combination with client certificates. Certificates are not encrypted in TLS 1.2. With considerable information put into the client certificates, users could be individually tracked, the researchers from the TU Munich found. In their [paper](#), they showed how Apple’s push service allowed for this kind of tracking with every authentication for a new TLS 1.2 session providing a marker of the user’s whereabouts. By the time it was contacted by the researchers, Apple had closed the leak. However, the client certificates are also used in other places, for example VPNS and mobile communication. As long as it is using TLS 1.2, the use of client certificate authentication should be avoided, said author Quirin Scheitle.

Another talk presented predictions about the impact of sea-level rise due to climate change on cable infrastructure in the US. Based on sea-level incursion projections from the National Oceanic and Atmospheric Administration (NOAA) and data from Internet infrastructure deployment in the Internet Atlas, the researchers from the University of Oregon found that 4,067 miles of fibre conduit will be under water and 1,101 nodes (e.g., points of presence and colocation centres) will be surrounded by water in the next 15 years. Regions with especially high risk are New York, Miami, and Seattle metropolitan areas.

Rank	Full SLDs	Partial SLDs	/24s	Same Last Hop
1	71,472	3,066	2	✓
2	69,637	328	2	
3	15,421	17	2	✓
4	13,044	3,727	2	✓
5	8,347	3	2	
6	6,111	631	2	✓
7	5,568	375	3	✗
8	5,076	69	2	
9	4,788	648	2	
10	4,611	4,820	4	
	204,075	–	23	–

Table 2: Information about the top ten SLD groups based on /24 address prefix.




Figure 4: Overlap of Internet infrastructure and seawater in New York (left) and Miami (right) with average sea level rise of 6 feet.

From the individual service provider infrastructures, Level3, Inteliquent, and AT&T are at highest risk. The researchers called for the development of mitigation strategies, e.g. Alternative infrastructure deployments. They did not study other parts of the world.

TLS: Encrypting SNI

The TLS WG adopted a document aiming at [encrypting the Server Name Identification \(SNI\)](#) during the meeting in Montréal. The lack of opposition came as a surprise, since a number of experts had expressed their objections on the mailing list before the IETF meeting. The SNI is one of the remaining points of meta data still available with most of the handshake being encrypted in TLS 1.3.

TLS 1.3 encrypts many data points, including the certificate. SNI was originally created to identify the recipient of a packet, as due to Cloud and CDNs, IP addresses are regularly shared. Practically, SNIs offered a flag allowing censoring and quality of service differentiation. The WG had for some time [discussed the leakage through SNIs and possible mitigation](#). A standalone solution for encrypting SNIs so far had not been pursued, because experts were afraid that due to the complexity, there would be few implementors, making those “stick out”. The authors from Apple and Mozilla now hope that private sources will hide in larger networks and app servers. If those would switch on ESNI then this would only point to the provider of the App, Cloudserver or CDN.

Technically, the provider will publish the public part of a key, possibly in the DNS (as txt or resource record). DNSSEC can provide authentication. The provider finally decrypts with the private key and sends the packets to the intended recipient. According to Eric Rescorla, this would be rather straight forward. A potential source for failed connections remains home routers that block the encrypted traffic.

A number of company representatives rejected the proposal on the mailing list once it was published, some warning that the balance between privacy and manageability of networks had shifted too much in the direction of the former. Former Security Area Director Kathleen Moriarty (Dell) complained that so far, the WG had expressed it would not go all the way to also encrypt the SNI. Moriarty’s statement was rejected by Rescorla and others who pointed to the earlier draft.

It is not clear if the ESNI concerns originate from the same groups already calling for static keys in TLS 1.3 to allow for better data centre manageability. Interestingly, though, a representative of the British National Cybersecurity Center called for [some place in the IETF to discuss the side effects of encrypted protocols](#) for cyber defence and law enforcement. During the London IETF, the National Cybersecurity Center had hummed loudly for a static key in TLS 1.3.

Another concern is related to ongoing centralization. With the ESNI solution building on the large platform model for obscurity, the trend to centralization and concentration continues.

IETF103 will be held on 3-9 November 2018 in Bangkok, Thailand.



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.

Rate this CENTR Report on IETF102

(Thank you for your feedback!)

