Council of European National
Top-Level Domain Registries

Report on IETF101

London
18-23 March 2018

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

Struggle over encryption in TLS

The power struggle over encryption continues in the IETF. A noteworthy debate took place in the TLS Working Group in London who fought over yet another proposal presented by Russ Housely, former IETF Chair, for a group of US businesses, namely banks.

Changes from 1.2 to 1.3

The basic TLS 1.3 specification is on its way to the RFC editor, meaning the Internet Engineering Steering Committee (IESG) has approved the draft. Major changes to TLS 1.2 is the choice of modern crypto algorithms settling on the faster, while more secure elliptic curve cryptography (ed25519 and ed448). Static RSA and Diffie-Helman cipher suites will no longer be used. The draft RFC promises perfect forward secrecy. Authentication and key exchange mechanisms have also been separated from record protection algorithm to further defend against active attacks. To speed up the protocol, a 0-RTT mode was added, allowing to start exchanging data with the first packet, but at some security costs.

One of the most important features of the TLS 1.3 certainly is that most parts of the header are now encrypted. After the ServerHello, all handshake messages are encrypted. The encrypted extensions message also allows extensions sent in the clear before to be encrypted.

Concerns of Going Dark by some industry, state actors

The very point of allowing for a higher degree of end-to-end-confidentiality is the focus of heated discussions at the IETF. Various data centre operators, the US banking sector and also the British National Cyber Security Center were very active during the TLS 1.3 session in London. They warned against the negative effects of end-to-end encryption as realized in the TLS 1.3.

The NCSC had put out a warning just a week before the IETF London meeting which says that individual security would win, but enterprise security would lose as monitoring, filtering and troubleshooting would become much more difficult.

A draft individual submission from the newly established Enterprise Data Center Operator organisation explains the thinking of those pushing for an inspection solution:

“Today there are enterprises with extensive packet broker networks who are doing out-of-band TLS decryption to feed network sniffers, intrusion detection devices, fraud detection, malware detection, application performance monitoring tools, customer experience monitoring tools, and other solutions. The capability to do out-of-band decryption has been available for twenty years, and for the first time in history it will be gone with the move to TLS1.3 [TLS13].”

The list of issues network operators had to address includes DDoS attacks, fraud monitoring, intrusion detection monitoring, threat detection and incident response. Alternatives including man in the middle decryption inside the network, use of TCP or UDP headers, staying with TLS 1.2, logging, securing and troubleshooting at the endpoint, encrypted traffic inspection or Ipsec are declared as either too risky, too expensive or not granular enough. Implementing proxies instead would cost millions, former IETF Chair Russ Housley said to this reporter. Regulatory requirements cited so far are very limited to US regulation against insider trade.

An opt-in solution for “visibility”

As an original proposal to allow for a static Diffie-Helman key for decryption by data centre administrators failed to get consensus last year, the US banking community in London came back with a proposal to allow for an opt-in mechanism for letting a smaller number of points in the data centre in on the traffic.

Presented by former IETF Chair Russ Housley, the proposed “extension to TLS1.3” restricts inspection to cases in which a client would signal acknowledgment to be inspected in the ClientHello and would then get ephemeral keys for the session. “No private keys will be shared”, Housley argued. A second set of keys distributed by the key manager in the data centre beforehand would restrict the destinations with which the decrypted packets will be shared.

Housley underlined the advantages – transparency for the client and better security in the data centre against idle use, or attacks against the decrypted traffic. However, he admitted that the mechanism was not limited to the data centre. While clients would normally be load balancers at the edge of the data centre, clients could also be payment terminals outside.

IETF rejects opt-in solutions in tech-thriller like session

The IETF session deciding on Housley's draft was nothing less than a little tech thriller. Decided opponents, namely former security area director Stephen Farrell, questioned the procedure of putting yet another proposal on the WG agenda. IAB Chair Ted Hardie reminded the group that there was a number of state actors that might oblige operators in their region to provide the respective key information, also for later inspection of the traffic.

Therefore, a "voluntary" approach still constituted an architectural weakness. History, and Snowden, had shown that data centre operators and state actors not always agreed in their definition of privacy, Hardie said hinting at the surveillance of data centre traffic shared between Google data centres unencrypted. Keeping the keys from mighty third parties was therefore improbable. Opponents also pointed to the technical alternatives available.

Several attempts of outgoing security area director, Kathleen Moriarty, to build bridges with additional limitations to the solution presented by Russ Housley, were rejected by the latter as insufficient.

In the end, the hum which the WG Chairs decided to hold revealed that the data centre/state security camp had done their best to fill the ranks. The NCSC had come with four, US Bank alone with 14 people. The hum favouring adoption of the visibility draft was clearly as strong as the one against. One observer sitting close to the data centre "camp" joked that breathing exercises obviously had been made before the session.

In the end, the WG Chairs decided that there was no consensus to take on the proposal.

Reactions and next steps

Housley said to this reporter after the session that he did not expect the data centre community to give up on their proposals, but they would not come back to the IETF. Instead, he said, ETSI, the European Telecom Standards Institute, would be happy to step in and he expected data centre representatives going there. The same had happened for the standard for legal interception over a decade ago. The disadvantage from the privacy advocates' point of view was that instead of the more transparent opt-in solution, ETSI might standardize the original static Diffie-Helman key solution (draft-Green).

Quic: Struggle over a Spin Bit

More fight over traffic inspection was delivered during the meeting of the Quic WG, albeit on a different level. While for TLS the data centre camp is asking for the clear text of the packets, in Quic, network operators are looking for meta data information. They hope for at least one bit, the "spin bit", for measuring round trips and do trouble shooting.

The nascent Google lab -originated transport protocol which is based on UDP is much tougher with metadata encryption. "In contrast to TCP, Quic's wire image exposes much less information about transport protocol state than TCP's wire image", the Spin bit draft explains. Especially losing the sequence and acknowledgment numbers as well as time stamps (available in TCP) makes passive measurements on the path impossible.

As in TLS, there was considerable debate by the 200 participants over 90 minutes. In the end, the inclusion of the spin bit solution for passive RTT measurements was not fully rejected, but put on the side-lines to allow to finalize Quic version 1 first.

One bit only

The spin bit draft proposal was presented in London by Brian Trammell, academic at the ETH Zurich and member of the Internet Architecture Board. The list of co-authors (Huawai, Telecom Italia, Nokia, Ericsson and AT&T Labs) illustrates the interest of network operators, mobile operators and network vendors.

The mechanism basically consists of a single bit added to the cleartext part of the header. It will be set to zero by the Client and flipped to one when the answer arrives at the client (one round trip, see graphic). The "flip" allows an outside observer passive roundtrip time measurement, while, according to Trammell, it is light-weight and also easy to implement.

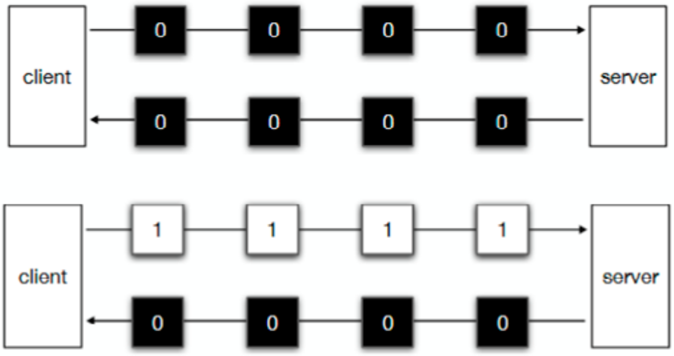

Trammell acknowledged during his presentation that RTT could be slightly overestimated by the mechanisms with imperfect network flows, but filtering out certain effects could help. With regard to privacy, a design group, for which Ted Hardie reported back to the WG, had not found an issue with the spin-bit solution, at least as long as the spin bit in fact was only one bit. A two-bit solution would enhance reliability, according to Trammell, but has not been made part of the draft proposal.

Spin bit solutions had been implemented in MINC and Quic-GO during the IETF Hackathon, Trammell said.

Spin bit put on a side track

While no candidate for the "invariants" of the Quic protocol according to the consensus in the Quic WG, the spin bit proposal came rather close to being put into Quic version 1. Miriam Kuehne, Transport Area Director, and researcher at the ETH, argued very much for including it in order to collect experience with the spin bit.

Even some privacy watchdogs like Daniel Kahn Gillmore, American Civil Liberty Union, seemed to give in, due to the positive evaluation of privacy friendliness.

In the end, the concerns resulting in postponing a final call stemmed from technical and time-to-market concerns for the standard document.

Spin bit: effects and indispensability not clear enough

Quic editor Jana Iyengar (who just moved from Google to Fastly) warned that including the spin bit into version one could delay the finalization of version 1 of the spec which has been postponed slightly to November 2018. He underlined that effects of adding the spin bit to the open header were not fully clear. On the privacy question, Iyengar said, while the spin bit looked inconspicuous, privacy issues sometimes were found later. Iyengar criticized the network operators, though. They missed to make clear why and how the spin bit was indispensable for network management.

Will Quic see similar discussions over its encrypted meta data down the road as TLS? Possibly, says Iyengar. Yet, for content people would be sent back to TLS. Quic could only become a target for a meta data discussion.

Status Quo of Quic

Meanwhile, measurements put Quic traffic to 9 or 10 percent. Practically all is Google traffic, and more than 40 percent of Google traffic now runs over Quic, according to Iyengar. Akamai also implemented Quic, but Quic traffic on Akamai was still "negligible", potentially due to the fact that Akamai customers have to "opt-in" to use Quic. More detailed figures on Quic traffic were presented in the session of the Measurement and Analysis for Protocols Research Group (MAPRG).

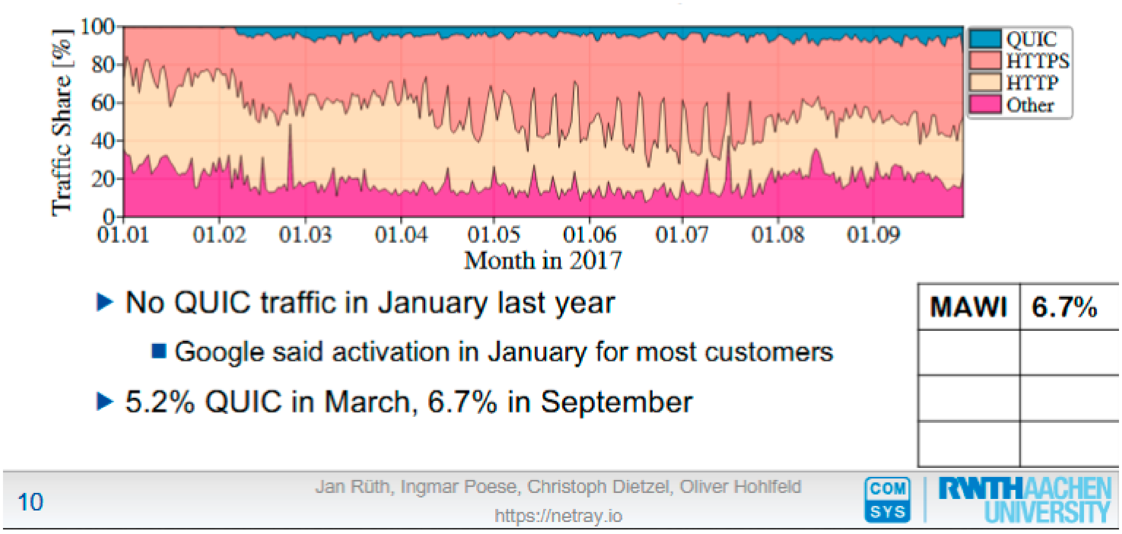

- No QUIC traffic in January last year
  - Google said activation in January for most customers
- 5.2% QUIC in March, 6.7% in September

MAWI	6.7%

10

Jan Rüth, Ingmar Poese, Christoph Dietzel, Oliver Hohlfeld

https://netray.io

On the question of whether Quic will be ready to be finalized by the IETF Bangkok meeting, Iyengar said it was an ambitious plan. The big issue to be solved over the coming months (another Interim meeting is taking place in Stockholm) is related to the hand shake.

The struggle on encryption throughout the stack

The discussions in TLS and Quic over the effects of encryption clearly are expressions of an ongoing power struggle. With encryption, one engineer told this reporter, power is put in on place. The changes in TLS and Quic currently result in a change of where the keys are put. Those who before had access to clear text traffic or meta data are shut out with the added TLS traffic, while the application providers remain in business with the user at the end-point. According to the engineer, the power shift explains the fierceness of debates.

Moreover, state actors see the shift as adversarial to their goals. One NCSC official, talking to this reporter after the hum, said that the office did expect to find their way around where necessary. Interestingly, Sujit Raman, Associate Deputy Attorney General Department of Justice during the Global Privacy Summit March 27 in Washington pointed to the added encryption in protocols and warned that it was wrong to leave decisions over encryption to technologists only.

In the IETF, the discussion take place in a number of places. Beside the discussions in TLS and Quic, there were:

- A presentation of a group of Cisco engineers (presented by Nancy Cam-Winget) that found its way in the OPSEC WG on “TLS 1.3 Impact on Network-Based Security” explaining: *“TLS 1.3 states that the client SHOULD include a “key_share” extension to enable the server to decline resumption and fall back to a full handshake, however it is not an absolute requirement. Example scenarios that are impacted by this are middleboxes that were not part of the initial handshake, and hence do not know the PSK. If the client does not include the “key_share” extension, the middlebox cannot force a fallback to the full handshake. If the middlebox policy requires it to inspect the session, it will have to fail the connection instead.”*
- There is an individual draft by Gory Fairhurst (University of Aberdeen) and Charlie Perkins University of Glasgow) on “The Impact of Transport Header Encryption on Operation and Evolution of the Internet” which was criticized on the Opsec Mailing list for including sentences like: *“Pervasive use of transport header encryption can impact the ways that protocols are designed, standardized, deployed, and operated. The choice of whether future transport protocols encrypt their protocol headers therefore needs to be taken based not solely on security and privacy considerations, but also taking into account the impact on operations, standards, and research.”*
- During the plenary, Stephane Bortzmeyer, Afnic, once more objected against proceeding the document on “Effects of pervasive encryption on Operators” from Security Area Director

Kathleen Moriarty to the RFC Editor. Bortzmeyer argued during the discussion with IESG members who had ack-ed the document that the IETF had to take position on the side of privacy. The proposition of "neutrality" in what has been called the "tussle" between individual privacy and operator security was false.

DNS expert rings alarm bell: do not overload the "Camel"

185 RFCs, 2,781 pages and 888,233 words – it is all DNS and it is just too much for the old resilient, hard-working protocol, Bert Hubert, PowerDNS said in a presentation on the DNS "camel". Highly entertaining, Hubert wanted to address what he sees as a critical issue: rising complexity and over-engineering of the protocol. The DNS session in London was exemplary for the issue, some DNS experts said. Tim Wicinski said the WG currently had 14 documents in various stages of the RFC process (see WG report below), and more on the way – and the two Co-Chairs calling for a third Co-Chair to share the workload.

Complexity: Risk of failure and market consolidation

According to Hubert, DNSSEC was the watershed moment, where complexity introduced into the DNS started to become prohibitive for smaller players. Two smaller, but well established DNS providers, MyDNS and DNS Mara, went out of business because they could not keep up with the pace to implement DNS technology when DNSSEC was introduced.

After DNSSEC, NSEC, NSEC3, Qname Minimization, CDNSKey, DNS over TLS were standardized and the list does not end there.

The growing number of standards led to more complexity of the DNS. This made operations more failure-prone, especially as DNS operations often were under-staffed. A company like Comcast had 21 DNS experts. "That is as much as all my customers combined have", Hubert said about PowerDNS customers. A trend to rely on the large DNS companies was another result. These large providers also brought new standard work, tailor-made for their needs to the IETF, adding to the list.

According to Hubert, one peculiarity of the DNS was that contrary to other protocol areas, there was a lot of open source software which was very good and even free (Bind, NLnetlabs, Knot). Yet the programmers and standardizers were too smart and were always tuning the standards, once more driving complexity. Moreover, there was no push-back against the standardization frenzy. DNS implementers outside of the tight-knit standardizer/implementer group gathered at the IETF would not dare to say new RFCs were too complicated for them to implement. Operators would not participate in the IETF to act as a corrective.

Reactions from the DNSOP WG were mixed. From outright defence for allowing the DNS to be extended and used in new ways (Alain Durand, ICANN) or the mere acknowledgment that the DNS suffering "wild success" was used "in unforeseen ways" (Andrew Sullivan, Oracle) to considerations of writing more documents to explain why the extensions were made (Matthijs Mekking). CZ,NIC representatives pointed out that they indeed had started to address the issue of workarounds to EDNS (RFC6891), see their press release. There were also more fundamental questions raised. A potential temporary stop for new standards

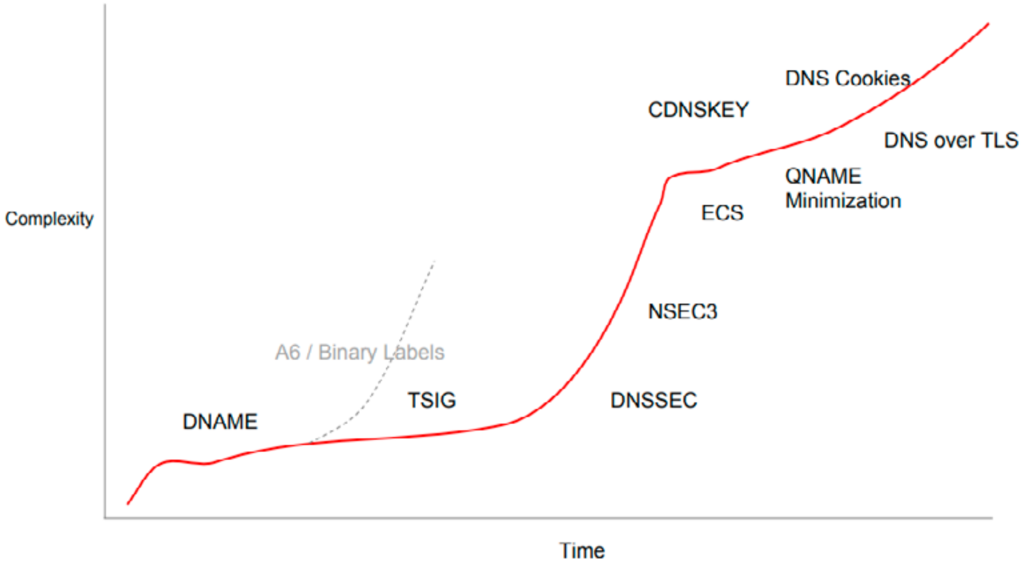

to allow time to prepare an oversight document for guidance was proposed by John Klensin. In a recently published informational RFC, Klensin had called for a discussion on the need for the DNS to diet, reform or develop v2.0 (RFC 8324).

RegEXT – Extension for ICANN's contracting policies

A discussion developed at the Registry Extension Working Group at IETF101 following the announcement that new registry extensions shall be prepared at ICANN. RegExt Chair Jim Galvin, Afilias, noted that registrars were not well represented in the IETF and the RegExt group was a rather small WG. Therefore, proposals developed by registrars at ICANN might come to the IETF, he expects. Both the registrar TechOps Subcommittee, as well as a TechOps group at the registries and also a joint group could be sources. Galvin said there were several documents on their way from the Registrar TechOps Subcommittee (on unavailable names, on a file format for reports between registries and registrars and a file format on billing transactions) to the standardization process.

Alexander Mayrhofer, nic.at, who has been one of the reviewers of pending extensions, warned against allowing standards-setting activities be performed outside of the IETF, with the IETF itself just being used as a rubber stamp mechanism. He was concerned about the fact that the ICANN TechOps groups were membership-only, so standards work was done in a closed space.

In addition to the procedural aspect, some of the mechanisms referred to by Galvin resembled more organizational (even contractual?) aspects, instead of technical issues. An RFC, especially a standard track specification, might be overblown.

Galvin argued, it would be better to bring the respective work to the one place in the IETF RegExt WG, instead of having it brought elsewhere. Those more sceptic envisage a strategic use of the process, with potentially even contractual mechanisms being moulded into "standards" and referred to as obligatory for that reason. One administrative document for example is the currently stalled draft from ICANN's office about the Trade Mark Clearing House functional specification.

Much to digest, lack of reviewers

It has been an issue of concern for some time that the WG has a very limited membership and documents are not scrutinized to the extent they are in other IETF WGs. So far, it has been mainly VeriSign and to a much smaller extent ccTLD registries like SIDN and CNNIC bringing proposals. The only registrar active in the group so far has been GoDaddy.

The three RFCs finalized by the WG are:

- RFC 8056: Extensible Provisioning Protocol (EPP) and Registration Data Access Protocol Status Mapping (VeriSign)
- RFC 8063: Key Relay Mapping for the Extensible Provisioning Protocol (SIDN) (there is a seemingly unproblematic IPR statement from VeriSign on this mechanism)
- RFC 8334: Launch Phase Mapping for the Extensible Provisioning Protocol (VeriSign, CentralNic, Cloud Registry), and the older RFC 7848 on the sunrise trademark procedures Mark and Signed Mark Objects Mapping (ICANN)

The WG is about to re-charter to take on new work, with quite a number of documents still on their way through the process, and also still needing reviews. The list includes:

- Verification Code Extension for the Extensible Provisioning Protocol (VeriSign)
- Validate Mapping for the Extensible Provisioning Protocol (GoDaddy)
- Registration Data Access Protocol (RDAP) Object Tagging (VeriSign, Arin)
- Organization Extension for the Extensible Provisioning Protocol, former reseller draft (CNNIC)
- Extensible Provisioning Protocol (EPP) Organization Mapping (CNNIC)
- Registry Fee Extension for the Extensible Provisioning Protocol (GoDaddy, CentralNic)
- Third Party DNS operator to Registrars/Registries Protocol (CIRA, Red Hat)
  Change Poll Extension for the Extensible Provisioning Protocol (EPP) (GoDaddy)
- Bundling Domains (CNNIC)
- Allocation Token Extension for the Extensible Provisioning Protocol (VeriSign)

The complete list of documents (milestones) is here. New work includes a considerable number of RDAP related extensions for search, reverse lookup, federated access and so on.

Decentralized Whois – another proposal from CentralNic

With GDPR being a top issue for registries and registrars, Gavin Brown, CentralNic took a stab at "decentralizing" Whois. He presented Whoiam as a decentralized alternative that would consist of a thin Whois combined with a domain owner's published v-card data publication (which also could, he proposed, be put directly into the DNS). Registrants could take control over the publication of their data (or outsource this to a third party, if they prefer, including to privacy proxies). They would also be able to check who had accessed the information, giving them more transparency.

Registrars and registries would shift responsibility to the registrants and benefit by not being the party publishing and controlling access to third parties. While Brown said that differentiated access could be realized, he also underlined that data mining still would be possible.

Scott Hollenbeck, VeriSign, viewed it as impossible to enforce the publication obligations against end users.

The document was not yet adopted by the WG.

Working Groups and BoFs

DPRIVE WG

The DPRIVE WG is about to re-charter, as their milestones with regard to DNS privacy mechanisms for the exchange of DNS data between stub resolvers and recursive resolvers are complete (DNS over TLS, DNS over DTLS, Profiles, Qname Minimization). One more draft on padding (against traffic analysis attacks) passed last call, and security review and is now on its way to the IESG.

The next issue to be addressed is the way from recursive to authoritative name servers. Contrary to the stub-recursive rather stable relation, recursives talk to many authoritative servers, making this the more difficult problem. There are IETF participants that hint at DNS over HTTPS as the better solution for privacy-friendly DNS. Stephane Bortzmeyer, Afnic, presented the short draft "encryption and authentication of the DNS resolver-to-authoritative communication".

For secure transport, once more TLS is proposed, but for authentication Dane is proposed. The authoritative server would need to add a TLSA record, the client then would open a TLS connection and authenticate via DANE (the DANE authentication could to speed the process, according to the draft be sent in the TLS session using chain-extension). The authoritative server could separate queries from the recursives could depending on their requesting TLS or not and send them to different servers, according to Bortzmeyer.

The WG still has to officially adopt the document. While Bortzmeyer said that the next step was envisioned in the original charter, re-chartering might still be necessary. The draft charter proposal included measurements on DNSpriv adoption (beside the recursive to resolver path protection). But this was ruled out by several participants calling it a research questions that might be in scope rather for the IRTF.

Privacy Practice

In an effort to document the evolving options for privacy-friendlier DNS services – and also push for adoption – Sara Dickinson (Sinodun) and several co-authors are preparing "recommendations for DNS Privacy Service Operators". According to co-author Roland van Rijswijk-Deij (surfnet.nl), the draft aims at presenting operational, policy and security considerations for practitioners and also help them with writing up their DNS privacy policy statements.

Apart from giving an overview over the new privacy enhancing capabilities for the DNS, the draft also tackles the issue of how operational practices, for example logging of DNS queries at the resolver, can be designed in more or less privacy-friendly ways. Logging and monitoring (and also data retention) could be minimized as much as possible and anonymized, access to stored data also minimized. Privacy DNS services should, according to the document, not track users, not provide data to third parties or aggregate and market query data.

As in the transport and security area discussions (TLS and Quic), some capabilities for troubleshooting could be retained by using pseudo-anonymization (i-cipher, bloom filters). Van Rijswijk-Deij presented experiments with a bloom filter solution (additional research is here) currently underway at surfnet. Bloom filters, originally designed in the 70s to index large data bases. They can be explained according to van Rijswijk-Deji as "a statistical way to test for set membership. Items that are added to a Bloom filter are run through a set of hash functions, and the output of these functions are used to set bits in a bit array. The contents of this bit array are then used to test set membership."

Bloom filters "do not store original query names" (but results of a set of hash functions) and are not-enumerable. Lookups are only possible when knowing what one is looking for. By mixing queries from multiple users in a single filter, tracking users is more difficult.

DNSOP WG

The "Camel" discussion left a lasting impression on the DNSOP WG, at least for IETF101. Entertaining two DNSOP sessions, there was a rapid succession of drafts, which after the Camel-talk, were put in the buckets: Camel – no-Camel. Even before the talks, Area Director Warren Kumari (Google) said that the WG had been good at adopting documents, but not as good in getting them done. WG Co-Chair Suzanne Woolf said in a reaction to Hubert's that not everything interesting coming through the door would become an RFC. Drafts mentioned, or briefly discussed included the below.

Co-Chair Suzanne Woolf asked the WG to comment on a possible .alt TLD draft which had been pushed aside for some time after heated discussion on it. Now Woolf wants to close the issue.

Paul Hofmann, editor of the terminology draft, announced he was getting close to working group last call (mid-April).

Stuart Cheshire (Apple) presented the session-signal draft, in the making since 2015 and a precondition, Cheshire noted for a number of drafts in DNS Service Discovery (DNSSD) WG.

Joe Abley (Afilias) came to re-animate the Refuse any draft (also going back to 2015 and dormant for some time). As ANY queries in DNS were used e.g. for amplification or mining of resource records, there could be a need for small responses. The draft proposes several more minimal answers to ANY questions. The draft will be put in WG last call soon.

The DNS capture format draft, presented by Jim Hague, still has IPR issues. The document shall provide for efficient storage and transmissions of large packet captures of DNS traffic.

WG Co-Chair Tim Wicinski recommended to solve this and go to last call soon. The concept has been deployed at some root servers.

Another WG last call document candidate is the KSK-roll sentinel that shall allow for better monitoring of the KSK roll preparedness of resolvers. It will allow an end user to determine the trusted key state of the resolver that he uses for his DNS queries.

The WG discussed a little longer on the ANAME resource record type, which is similar to CNAME but is limited to type A or AAAA queries. It shall be an alternative to CNAME (where the use of CNAME is prevented). The main discussion was about splitting the document for authoritative and resolver side, something that was rejected by many WG members.

Strong support from ICANN representatives (David Conrad) and others was provided for the trust bootstrapping mechanism document presented by Joe Abley (Afilias). The document provides guidance on how validating resolvers can determine an appropriate trust anchor for the root zone to use at start-up, or when other mechanisms intended to allow key rollover (5011) are not available. There is a lot of talk that 5011 in general should be supposed by a better mechanism.

More DNSSEC related work tries to solve the issue of companies that use different DNS providers for their authoritative DNS service. The draft, presented in London by Shumon Huque (Salesforce) lays out several models of how to deploy DNSSEC in that case.

More straw for the camel's back some thought was the XPF document presented by Peter van Dijk with XPF (ISC). The draft proposes a new "option within the EDNS(0) Extension Mechanism for DNS [RFC6891] that allows a DNS server to receive the original client source IP address when supplied by trusted proxies". It shall solve the issues that front end proxies (for load balancing, e.g.) are hiding the original client's source address from the DNS server, making it more difficult to use ACLs, DNS, Response Rate Limiting and other server side technologies. The draft acknowledges that incorrectly used XPF could expose internal network information. As it was intended for the server-side proxy (under the same administrative control as the DNS servers, there was no change of what private data could be shared. Many WG members criticized the draft warning that it would not be good for the WG to rely on good behaviour and well-meaning actors. The WG was also split on another ISC originated document, a proposal for automatic zone provisioning. Andrew Sullivan called the proposal a "camel farm".

For a longer summary on the DNS WG session, see Paul Hofmann's minutes.

Homenet WG

The homenet WG does not seem to make substantial progress. The only milestone document discussed during the session was the simple homenet architecture document. The document describes the publication, resolving of names and discovery in homenets.

The document got a substantial re-write to clarify, according to Teld Lemon, that it was no full architecture, that no full-service resolver was required to serve homenet queries ("a proxy will do as long as it splits out queries for local zones"). Using a discovery proxy, the following locally served zones would be supported:

home.arpa

fc.ip6.arpa

10.in-addr.arpa

168.192.in-addr.arpa

16.172.in-addr.arpa

queries for all other locally served zones are answered

Lemon said he would implement in OpenWRT and Turris and come back with code.

The home.arpa delegation, chosen after a .home special name delegation was ruled out, meanwhile is stuck in the RFC editor queue, due to the fact that IANA has work through the delegation. In a discussion on the DNSOP mailing list, Kim Davies, IANA, explained that delegating home.arpa to AS112 was chosen as "the best short-term approach". While not without "its own difficulties" it was preferred to have dname records in the root servers for the necessary insecure delegation of home.arpa. DNSSEC insecure delegation is necessary in order to not have the validating resolvers/home routers block home.arpa resolution.

Perimeter security for the homenet was only briefly discussed, the topic, while a milestone, has been dormant and the WG Chair said if nothing was forthcoming, the issue would be closed. Ted Lemon said he would be able to work on the issue between IETF 101 -102. Security for babel (either through hmac or DTLS) was briefly discussed.

In an effort to link homenet to possible bricks it could reuse from Anima, Michael Richardson presented the Anima protocol suite. The problem to be solved by Anima was the secure joining of new network devices to an enterprise network. One component from the Anima suite which Richardson offered was "Bootstrapping Remote Secure Key Infrastructure" (BRSKI) (the other basic Anima components – unnecessary for homenet - are "a secure and dedicated channel (VPN) for management/control (aka. ACP)" and "a generic signalling protocol (aka. GRASP))".

Richardson said a BRSKI profile for homenet might be an option. A challenge was that contrary to the Anima-target enterprise network, home networks were unmanaged (or unprofessionally managed). Not being able to connect to Wifi in the first place (before bootstrapping a device) could result in calls to services provider or vendor. Richardson also acknowledged that the Anima concept of assigning a candidate device starting up automatically in a network to an owner via a voucher to be checked by the "Manufacturer Authorized Signing Authority" (MASA) could be seen as the IETF supporting vendors keeping control. At the same time, Brski was a good compromise between security and usability, Richardson underlined.

DNS over HTTPS, standards done, WG could be closing

The DNS over HTTPS (DoH) has just started, yet still expects to bring its specification to WG last call in April 2018, just seven months after being set up. If no further document is being brought to the WG, it will shut down after finalizing the document. DNS over HTTPS could move one more step towards moving DNS away from users/developers as it will become moulded into http.

DoH maps each DNS query-response pair into a HTTP request-response pair. The approach, according to the draft by Paul Hoffmann (ICANN) and Patrick McManus (Mozilla), establishes default media formatting types for request and responses, but "uses normal HTTP content negotiation mechanisms for selecting alternatives that endpoints may prefer in anticipation of serving new use cases. In addition to this media type negotiation, it aligns itself with HTTP features such as caching, redirection, proxying, authentication, and compression."

Two issues discussed during the WG group was if the get and post mechanism should equally be made mandatory to implement for the server, and there was consensus in the room that it should. Clients on the other hand would then be able to choose. The other issue the authors wanted to see a compromise on was if udpwireformat should be made mandatory. The draft opted for yes. This was supported during the WG meeting by several speakers, including Stewart Cheshire, Apple, who noted that as the DNS continued to evolve, more extensions would be defined using UDP packet formats. If these could just be wrapped and carried over UDP it would be much easier than to do ever new additions for the respective extension.

There is considerable interest, not the least from some large providers. Stephane Bortzmeyer, Afnic, provided an overview over implementations (listed on this GitHub site). They include an operative DoH server from Google, Akamai and Clean Browsing. There are also several “toy servers” which implemented DoH. With five different server software used, and four being publicly distributed, a high degree of diversity was already achieved. Implementations were “no big deal”, Bortzmeyer said, but some issues in the draft still had to be clarified for the sake of non-DNS experts (e.g. the HTTP crowd).

Opportunistic DNS

Daniel Kahn Gillmore, American Civil Liberty Union (ACLU), presented what he described as a trigger for discussion for next steps. Instead of only using DoH for the marrying of HTTP and DNS, he recommended a kind of a push-mechanism for DoH. The server would place IP addresses for names not requested into responses, for clients to use down the road.

Once cached the additional names locally, no additional DNS requests would be necessary, the mechanism would benefit privacy, latency. Also in order to benefit from authentication, a push toward more DNSSEC might be incentivized, a welcome side effect, according to Kahn Gillmore.

The “push DNS” will not necessarily become a WG document; instead, it has to be settled at the respective WGs. The WG Chair for HTTP, Mark Nottingham, invited to lead the discussion at http during the next IETF.

Message layer security

Interesting new work was started at the Message Layer Security BoF. The soon-to-be-established WG wants to standardize an asynchronous group key management for groups from two to thousands. While TLS allows to secure end-to-end connections, MLS is expected to specify a key establishment protocol for various messaging groups independent from the transport and application used (including chat, sip and possibly even mail). Proposed and presented by authors from Cisco and Google, with Co-Authors being from Facebook (and Whatsapp), Wire, Inria and Twitter, the level of interest in the work is huge.

Basic elements of the MLS concept are an authentication (initial key) and a delivery services (delivering messages, adding and removing group members) which can be independent from each other. Interoperability of different applications is not intended, but expressions on potential federation (for authentication) seems to vary in the original documents. No new message protocol shall be established, but rather existing ones (like Cose) should be re-used.

At the heart of the new protocol lie the security features. According to Richard Barnes (Cisco), these are not only forward secrecy (earlier communication content protected after compromise), but also post compromise secrecy (PCS, communication content will be protected after a certain point after a compromise happened). Standardized, secure key management for group communications has been a desideratum for some time, an engineer from Matrix.com underlined.

The BoF already hummed on a Charter point for the work that underlines that a “visibility” extension (meaning protocol included decryption capability for third parties outside of the group) are excluded. Wording on this point changed over the discussion and there were some calls to not address this at all.

The architecture document is here, the base spec is here. The crypto behind the concept can be checked out in this academic paper on the “Asynchronous Ratcheting Tree”.

IASA 2.0

For some time, the IETF community has tried to make up its mind about the future relation to the Internet Society (ISOC). At its London meeting, the IASA 2.0 BoF decided to go ahead with setting up a limited liability company (LLC) for the IETF administrative operations. As a subsidiary the ISOC.org will take control of its funding and contracting. The IAOC, IETF administrative oversight committee, can be replaced by a Board of Directors (see the draft by Brian Habermann e.a.).

The way to populate of the Board, the interface to the community (Advisory Council?) and other details (including the potential ending the IETF Trust and keep the IPR of the RFC series with the LLC) are details to be discussed in a future Working Group. See the graph for possible details on the new bodies and the graph below giving the overall picture:

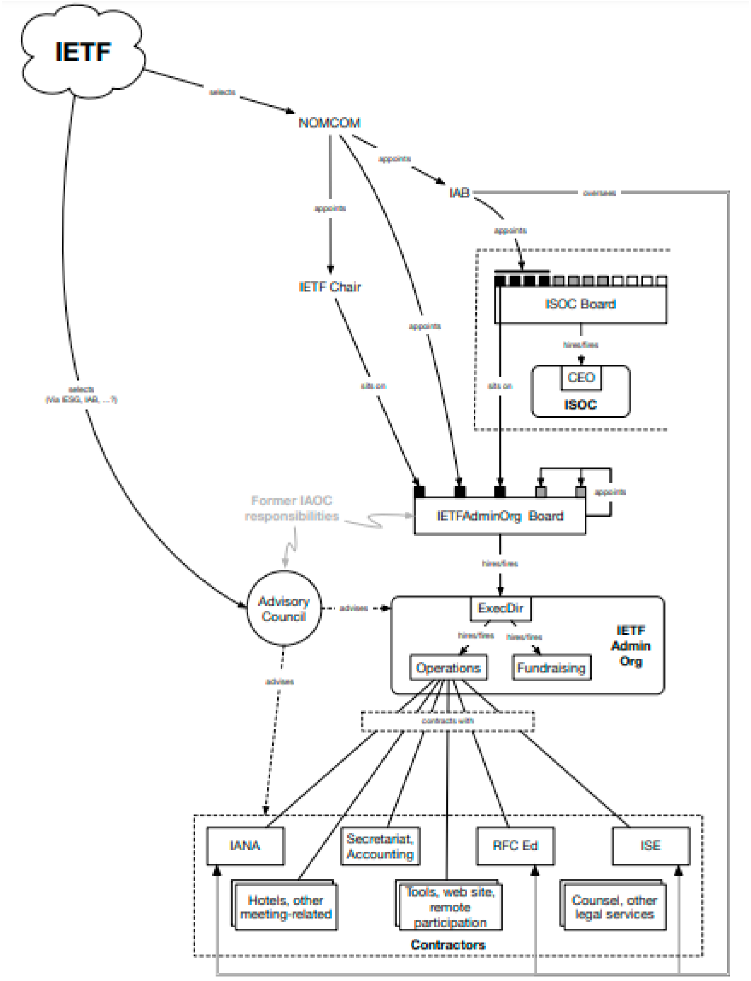

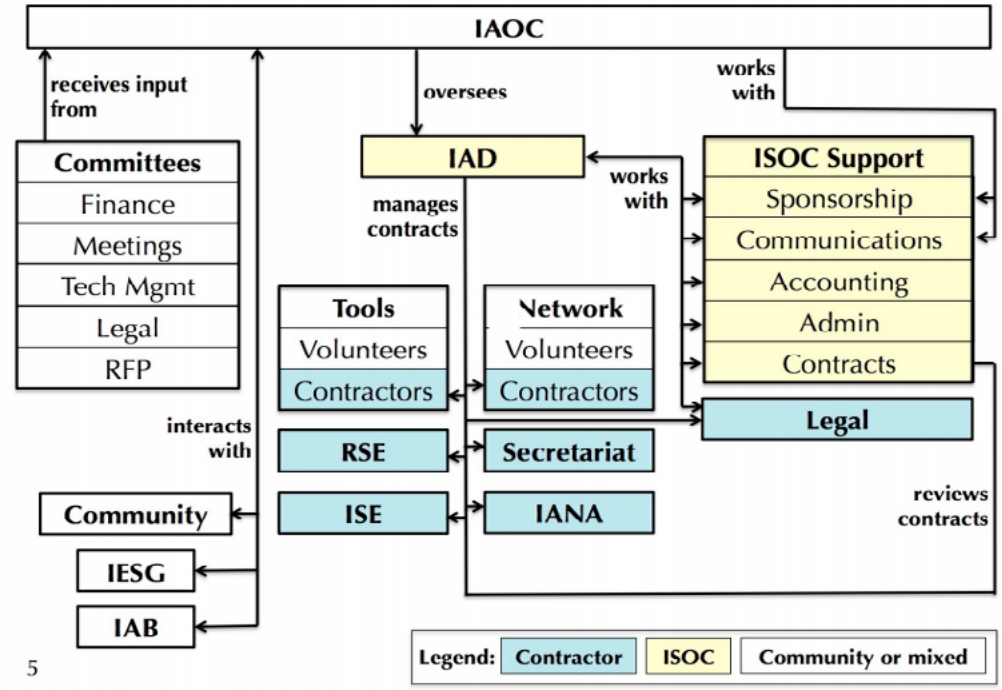

The decision for the subsidiary model (against keeping the status quo or cut the ties to ISOC completely) has to be confirmed on the mailing list, IETF Chair Alissa Cooper announced during the Plenary meeting. A WG will work out the details and make an update to the old BCP 101 and IASA.

The legal texts for the LLC will be written outside of the WG, by lawyers (the IETF during the plenary meeting presented their two new lawyers: Brad Biddle, Biddle Law PC, David Wilson, Thomson Hine LLP).

While consensus was declared for the LLC model, not everybody voted in favour of the new IASA structure. Avri Doria, who is Chair of the Human Rights Protocol Considerations (HPRC) Research Group told this reporter she was concerned about liability for the standards body.

Budget issues unresolved

In one respect, the organisational reform will not change much. While the IETF will be able to contract and hire and will not rely on ISOC to act as a legal roof, a big chunk of the money to spend will continue to come from ISOC.

Andrew Sullivan, new IAOC Chair, presented the budget and explained the financial gap the standardization body faces once more for 2018. While expenses are static and the budget for 2018 is roughly the same as in 2017 ($7M), meeting attendance has been going down. For 2018 the IAOC therefore calculated that there will be a gap of $300,000 from paid attendance. The amount will be borne by ISOC again, but the IAOC, in an effort to make the IETF less dependent from ISOC financially, will raise the meeting fees (more than 10 percent in 2019 from currently $700, starting in 2020 three percent annually).

There was some discussion on the meeting fee rise, with two participants from African countries (both ISOC fellows) warning against not meeting the target to make the IETF more global and more inclusive.

The effects of IASA 2.0 on the financial situation is not clear, according to Sullivan.

To keep sponsorship money coming in, a new sponsorship fundraiser has been hired (Ken Boyden).

IANA.com was transferred from ICANN to the IETF Trust on March 8. The transfer of IANA.org and IANA.net was completed the week before Easter.

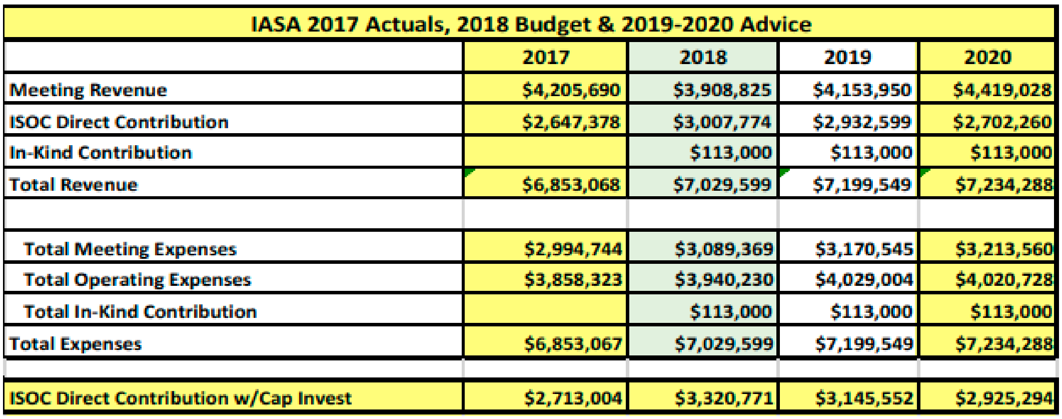

IASA 2017 Actuals, 2018 Budget & 2019-2020 Advice				
	2017	2018	2019	2020
Meeting Revenue	$4,205,690	$3,908,825	$4,153,950	$4,419,028
ISOC Direct Contribution	$2,647,378	$3,007,774	$2,932,599	$2,702,260
In-Kind Contribution		$113,000	$113,000	$113,000
Total Revenue	$6,853,068	$7,029,599	$7,199,549	$7,234,288
Total Meeting Expenses	$2,994,744	$3,089,369	$3,170,545	$3,213,560
Total Operating Expenses	$3,858,323	$3,940,230	$4,029,004	$4,020,728
Total In-Kind Contribution		$113,000	$113,000	$113,000
Total Expenses	$6,853,067	$7,029,599	$7,199,549	$7,234,288
ISOC Direct Contribution w/Cap Invest	$2,713,004	$3,320,771	$3,145,552	$2,925,294

IEPG on DNS: quite some recommendations

The Internet Engineering and Planning Group, established to create an interface between protocol engineers and operators and regularly meeting before the IETF meeting had a considerable number of DNS presentations on its agenda.

Five recommendations for DNS operators emanating from the results of academic work were presented by Giovane Moura (SIDN Labs). The recommendations are:

R1: all authoritatives should have similar latency

R2: Routing Can Matter More Than Locations

R3: Detailed Anycast Maps of Catchments Requires Active Measurements

R4: When under stress, two strategies

R5: Shared Infrastructure Risks Collateral Damage During Attacks

Truncated Responses mitigation

A proposal by Geoff Huston and Joao Damas wants to mitigate problems with truncated responses. DNS over UDP doesn't work on IPv6, and fragmentation is widely unsupported. The ATR concept is to send one packet with TR (truncated) flag behind a truncated packet. If the client receives the fragmented answer it will ignore the ATR packet. If the fragmented answer doesn't reach the client, the ATR probably will and the client will switch to TCP.

Removing EDNS workarounds

ISC (Bind), .CZ (knot), NLnetlabs (unbound), PowerDNS will remove workarounds for broken EDNS0 implementations and only allow standard responses after 2019-02-01. For tests go to:

https://ednscomp.isc.org/ednscomp/

The open-sourced test suite: https://gitlab.isc.org/iscprojects/DNS-Compliance-Testing

Crippling DS records

If there is a sha-256 ds records, the ds records with sha-1 won't get used by resolvers. After several TLS went bogus because a ds record for a non-existent, dnskey was introduced, ICANN wants to mitigate the fact that there is currently no prescription of the prevention of the failure mode where DS with SHA1 is ignored in the presence of SHA2.

Roy Arends presented ICANN's recommendations:

- Be consistent is using digest types in DS records
- Use the same digest type(s) for every KSK.
- Don't rely on your parent to figure it out for you.
- Its 2018. You don't have to use SHA1, you can safely use SHA256.
- Do not roll the KSK and the DS digest type at the same type, either roll the KSK OR roll the DS digest type
- If there is a DNSSEC Best Current Practices 3, this

should be added.

- There are 8 top level domains which are SHA1 only. All others are either SHA2 or dual SHA1 and SHA2

Other interesting content from IEPG

Two measurements are interesting: one is the [RIPE IXP Country Jedi project] which allows to show how User-to-User connectivity inside a country is happening:

These maps show the IPv4 paths (left) and IPv6 paths (right) seen in traceroutes. Indirect links in traceroutes (ie. with hops inbetween without answer, or no geoloc) are shown with dotted lines, direct links with lines with long short alternating pattern.

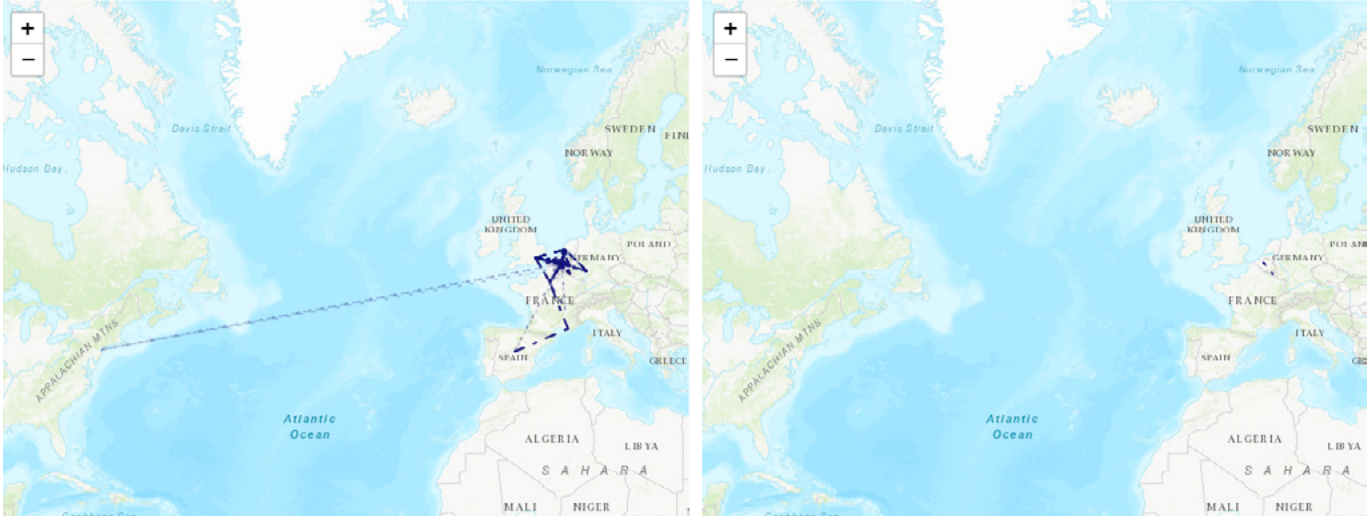

The other is public DNS resolvers; what [roots do resolvers] use? The measurement of ICANN illustrates that there are 20 percent of "strange answers". Of the 25,881 addresses looked at:

- 16,835 returned a response (65%)
- 13,826 returned the expected SOA record
- Of these expected SOA records:
- 13,800 returned expected SOA serial (at most 2 days off)
- 5 had a different formatted SOA serial number (1520976703)
- 21 had a serial number than was out of date (eldest is 2012041813)
- 3,009 returned an unexpected (completely different mname, etc) SOA record.
- from those that responded, 22% (3,009 out of 13,826) have other roots configured.

IETF102 will be held on 14-20 July 2018 in Montréal, Canada.

CENTR is the association of European country code top-level domain (ccTLD) registries, such as .de for Germany or .si for Slovenia. CENTR currently counts 54 full and 9 associate members – together, they are responsible for over 80% of all registered domain names worldwide. The objectives of CENTR are to promote and participate in the development of high standards and best practices among ccTLD registries.

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