



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

# Report on IETF99

**Prague**  
17-21 July 2017



# Contents

## **Highlights** **3**

---

IETF and the politics of standards	3
DNS – XPF record, really?	4
Quic progress & information exposed on the wire	4
The fight about an escrow key in data centres for TLS	5
A bouquet of new DNS Transport options – how to choose?	6
Venues: IETF changes meeting venue to avoid the US	8

## **Working Groups and BoFs** **10**

---

NetSlicing BoF and 3GPP	10
Security Area Advisory Group – Chances of quantum computing	11
DNS	12
Homenet: simple naming document “close to done”?	13
DNSSD – Departure from multicast and better privacy	14

# Highlights

## IETF and the politics of standards

The debate about the political nature of standards and potential ethics in standardization has become a visible thread in IETF discussions since Snowden. While attendance to the dedicated IRTF working group, the Human Rights in Protocol Considerations ([HPRC](#)), is very much limited to a number of highly interested developers, concerns over privacy in designs is coming up regularly in several other working groups (WGs).

Following Dave Clark's presentation on the "tussle" – the dilemma between conflicting interests and rights in technology development (see [CENTR Report on IETF98](#)) – the Chair of the HPRC WG, Niels Ten Oever, and the former IAB Chair and well-known DNS expert (formerly Dyne, now Oracle) Andrew Sullivan got together to clarify the relation between politics and standards in a new [draft RFC document](#). Presented at the Prague meeting, the document explores the different conceptions of the politics/standards or standardization relationship on a position spectrum from "technology is neutral" to "standards are politics with other means".

Regarding the document's purpose, authors and supporters considered it a possible "politics and standards" crash-course read for the IETF community, which might be driven away from the original "[Research into Human Rights Protocol Considerations](#)" that is on its way to drafting the first RFC of the HPRC WG.

### "Code is not law"

Well-known tech policy professor Milton Mueller (Georgia Tech) challenged the idea that standard developers and technologists could exert a big influence on rules for communication on the internet. At best, technology had a mediatory function for human rights (HR). Mueller rejected Lessig's Code is Law argument. "Code is not law", he said, adding that law was often "overwriting" technology options. Mueller pointed to the development of the CALEA wiretapping legislation against the denial of the IETF (in RFC 2804) as one prime example. In the end, HR protection was an institutional and political effort, not a technological effort.

So while it was good for the developers "to be aware" of the issues, their influence was limited. Together with Farzaneh Badii, Mueller is working on a paper that questions the idea of "Advancing Rights via Internet Architecture". Mueller and Farzaneh speak of a "Requiem for a dream".

Additional arguments for a need to "wake up" from that dream made by Mueller were that assessment of rights was only possible *ex-post* (instead of being figured out *ex-ante*), internet design was too stable already to do big things in design, and HR were complex and included balancing of conflicting interests. At the same time, Mueller warned against the temptation to make decisions without input from other stakeholders. A politicization of standards could therefore result in a questioning of the legitimacy of technical developers to design and could also draw other groups, especially regulators and governments, into the standardization process.

The latter was clearly challenged by Ten Oever, who pointed out that governments have long been participants in standardization, including at the IETF. NIST and NSA, for example, are regular participants and over the years, through their own personnel or sponsored representatives, have taken core roles in the IETF/IRTF. Georg Mayer (CT Chair of 3GPP) from Huawei and Bob Hinden both underlined that considerable design changes have been and are under way for mobile networks and encryption. The potential to disable or enable monopolies via standards could also have an effect on HR, said Philip Hallam-Baker (Commodo). Finally, Allison Mankin, IRTF Chair, made a pointer to discussions on how to potentially regulate algorithms (before they regulate us).

## Standards & Politics in practice

There is much politics in decision-making for competing standard designs, as was aptly exhibited during IETF99.

In three different WGs, there were passionate debates on requests presented by operators that according to privacy and security experts, failed to reach IETF standards for secure and privacy-friendly communications. The requests were: (1)

for a new XPF resource record in the DNS WG that would put personally identifiable information (PII) in DNS packets, to help load balancing; (2) to expose information on round trip times in the new Quic transport protocol for traffic management; and (3) to include the possibility to use a static Diffie-Hellman Key in the new TLS 1.3 standard to allow to see traffic to and from data centres on the wire, for troubleshooting.

Never before had there been so much tension at an IETF meeting on a single issue – the conflict of interest between privacy/security and companies’ operational interests, observed Sara Dickinson, DNS and DNS privacy expert at Sinodun. Brief recaps of these discussions can be found in the DNS, the TLS and the Quic WG summaries below.

## DNS – XPF record, really?

Adding personally identifiable information about your DNS customer – without him being aware – is already done by several DNS providers and baked in the DNS machinery by some vendors as an option. During the second of two DNS WG sessions (see below), two proposals were presented that are said to address issues of DNS network management: “[Client ID in forwarded DNS Queries](#)” as well as a new resource record [DNS X-Proxied-For](#) (XPF) were discussed.

For Client ID, the authors (including David Lawrence, Akamai) underlined the need to allow for “customized DNS responses”, like for example “parental control”. For XPF, which has been written by authors from ISC and PowerDNS, the reasoning is the use of proxy devices and the negative of hidden source addresses for load balancing.

While XPF is intended to sit between the load balancer and the actual server and should, in theory, stay in the server’s premises, Connection ID would sit between the end-user machine and a provider. The latter, as Stéphane Bortzmeyer, DNS Expert from Afnic explains, could be considered as much more dangerous. Nevertheless, both proposals go into one direction: both add meta-data to DNS queries and enable pervasive monitoring.

The discussion over these drafts certainly clarifies that adding personally identifiable information from IP addresses to MAC addresses or, as the Client ID draft proposes “other defined Identifier-type values”, is regular practice for some operators, including

vendors like Cisco, Nominum or PowerDNS. Therefore, they would like to get a standard document approved – or even an informational document with an IETF stamp. However, many in the DNS ccTLD community and the privacy expert camp are worried that this will just counter the efforts to make the DNS more privacy-friendly and privacy-legislation-compliant.

The fight over how to balance ease-of-use for operators versus a better protection of privacy is by no means exclusive to the DNS community. It seems to be a common trend at the IETF these days, with the TLS WG and the Quic WG both being locked in duels on these positions for extended parts of their sessions in Prague.

## Quic progress & how much information should the new transport expose on the wire

For many, Quic is one of the big things at IETF, as it is the first attempt to come up with a successor to TCP in quite some time. According to [figures presented](#) by Jana Iyengar (Google) during the Measurement and Analysis for Protocols Research Group (MAPRG), 35 percent of all web traffic and 7 percent of all internet traffic travels via Quic today.

Using UDP as a substrate, Quic integrates transport protocol with immediate encryption (goal: TLS 1.3, and 0RTT on resumed connections) and also promises to do away with TCP head of line blocking through the use of multistreams and strictly ascending numbers for packets.

## Running Code: First Interop with Google, Mozilla and others

The Quic WG met for a first Interop meeting just before IETF99 to test five implementations of the new transport protocol that Google has brought to the standards body after several years of testing it on their networks.

At the Interop meeting, besides Google, four other implementations showed up: Mozilla, a Microsoft implementation by Christian Huitema, an implementation from WireShark and another little one by Quic WG Co-Chair Lars Eggert. In essence, they achieved a handshake to establish a basic Quic connection and a close. In the first of two WG sessions, there was a discussion about how ambitious

the second Interop should be, with Mozilla and Google pushing to include at least a small application, if not even parallel or multiplexed streams.

Iyengar urged to agree on the wire format of the Quic packets as soon as possible to prevent middleboxes ossifying on the Google Quic, causing implementation problems for the IETF Quic, which indeed looks different on the wire. The header format was one aspect of Quic that was changed during the Quic WG's first year of work. The flags Google had in its Quic headers for example were removed, yet these very flags, according to Iyengar, are something the middleboxes use to detect Quic.

Issues currently under discussion in Quic are the mapping of http on Quic, with some people also warning not to focus on the http mapping alone, but to make Quic a real generic protocol. Another discussion underway is the one over uni-directional or bi-directional streams. Nevertheless, the most controversial issue at the moment is, as mentioned, privacy considerations.

### How privacy-invasive are passive RTT measurements?

The advances in transport encryption with TLS practically baked into Quic is welcomed for security/privacy reasons (and its efficient provision). But the step to also encrypt parts of the headers and only leave few elements in the header's visible part, namely a *Type* (5), a *Version* (32) and a *Packet Number* (8/16/32) field, with a *Connection ID* being only optional, is cause for concern for network operators. They will lose information filtered from TCP headers and used for network management, namely queue and congestion management, as well as trouble-shooting.

In a debate similar to the one in the TLS WG and the DNS WG, operators lined-up in the Quic WG to request a mechanism that would give some of their ability back to measure Round Trip Times (RTT).

To come to a decision, Ian Swett (Google) presented [four options](#) on how to proceed: (1) do nothing; (2) packet number echos; (3) one spin bit set per RTT; or (4) identical bit value for an RTT of packets. Weighing the pros and cons of keeping the status quo, which only makes the handshake RTT visible but nothing more, Swett acknowledged that network operators and innovative middleboxes might "attempt to infer

RTT" or use other ways around, or even block Quic packets. Several operators, as well as Brian Trammell (ETH Zurich), confirmed this.

2. Packet number echo: "the sent packet exposes a packet number and the peer echoes that packet number back on ack-only packets"

3. Spin bit idea: "one packet per round trip sets a spin bit in the header to up (1) others are sent

with the bit down (0), which is echoed by the peer"

3a. Identical bit: "the connection initiator sends packets with a spin value of up, the peer reflects the spin in response packets, and the initiator flips the spin"

Options 2, 3 and 3a were all rejected as re-enabling passive monitoring and surveillance, and should not be allowed, especially due to potential abuses in the future, regardless of current justifications. Iyengar offered that Quic would at least already solve the issues currently addressed by queue management through Quic's advanced behaviour with regard to multiplexing and easing traffic flows.

As after 90 minutes, both sides could not come any closer in position, a design group chaired by Ted Hardie (Google) was agreed upon. It will try to evaluate the effects of not allowing for some information being made available in the Quic headers and, on the other hand, also the privacy effects of giving in to network managers' request.

Quic will have another intersessional meeting before the IETF100 meeting in Seattle in October. The intersessional will have both an Interop and several days of regular WG meeting. Around 60 people show up at these intersessionals, according to Lars Eggert.

### The fight about an escrow key in data centres for TLS

The Transport Layer Security Working Group (TLS WG) is close to finalizing TLS 1.3, successor standard to TLS 1.2. Major features include: mandatory forward secrecy, one-round encrypted connection establishment, and having a feature to prevent downgrade attacks (from 1.3 to 1.2). After a second WG last call RFC, author Eric Rescorla (Mozilla) nevertheless asked for a few more weeks for testing the new standard after earlier test rounds (at Mozilla, Google, and, without transparent data given,

Facebook) revealed a rise in failure rates.

The measured failure rates were between 1 and 10 percent, according to Rescorla. At Google, according to one source, connection establishment with TLS 1.3 failed in 5 percent of cases.

Martin Thomson (Mozilla) said middleboxes of two (large) vendors were thought to be the problem. Because of this, developers want to continue testing. One potential result could be another tweak of the draft standard text. The basic fix considered is to change the content type of the Server Hello. The WG would have to make a final consensus call on that tweak. The vendors in question were not revealed.

### **Sniffing on all traffic at the data centre**

The hot potato in TLS, however, is not on the further tweaking of the specification to squeeze TLS 1.3 encrypted packets through middleboxes. Instead, it is about how much the new TLS should be tweaked to allow network operators to control traffic in their data centres. A [proposal for a static Diffie-Hellman key](#) to allow for key escrow in the data centre has ignited an epic battle in the TLS WG. If integrated into the proposed standard RFC, it would break TLS forward security to allow operators to sniff all data to and from their servers.

Essentially, the proposal wants to have a static key (to be rolled frequently) made an option instead of the ephemeral keys that are mandatory according to the TLS 1.3 proposed standard draft specification.

The proposal was laid out by Matthew Green, well-known security expert, who had been hired by a number of companies from the US banking (and network security) industry. Other authors include former IETF Chair Russ Housley (VigilSecurity) and former IETF Internet Area Director Ralph Droms. So far, the authors only include US companies that have received support by NIST, the US Network and Information Security Technology Agency. NIST had gathered the industry group at a workshop to formulate their proposal and also announced at the Prague IETF meeting that it would present a proposal of its own to solve the monitoring problem.

“It’s wiretapping”, ranted Stephen Farrell, former Security Area Director who collected a long list of arguments objecting to the IETF even continuing to discuss the proposal.

The concluding hum revealed a close 50-50 split of the participants.

The TLS back-door discussion can be expected to continue, despite some opponents declaring the issue “dead” after the Prague discussion. On the other side, in a conversation after the session, members from the proponents group also declared “victory” as they had expected the WG to decide more clearly against the back-door key option.

### **A bouquet of new DNS Transport options – how to choose?**

DNS over TLS, DNS over Http and DNS over the new Quic – all were presented at the Prague meeting. Some observers like Alex Mayrhofer warned against pushing the newest transport competitor “Quic” at the cost of DNS over TLS. The concern is that the competition could result in making implementers hesitate even more on putting the DNS privacy-enhancing DNS over TLS into practice.

Sara Dickinson (Sinodun) qualified the concern. Being at the same time one of the developers of a DNS over TLS stub resolver software package and an author of the still rough “DNS over Quic” draft document (together with Christian Huitema, Microsoft), she said to this reporter that Quic could be a very interesting solution for the resolver to authoritative server part of the path, for sheer efficiency reasons. In the medium term, DNS over TLS was still necessary to implement. She also thought that making the effort to implement DNS over TLS first was beneficial for those later turning to Quic, as the effort to move to Quic with integrated encryption would be much more manageable. It was also certainly a first step toward creating awareness for DNS as a privacy-sensible service.

In a short note to the author, Erik Kline from Google made a similar statement, writing “we’re in the process of getting DNS-over-TLS working and integrated first. All these alternative transports will require more work (and necessarily reuse much of the TLS integration, so it makes sense for us to try TLS first). [...] It might be that operational experience with DNS-over-TLS informs how alternative transports progress. Some measurements will probably need to be done to compare reachability on port 853 with port 443, for one thing.”



## DNS over TLS implementation steps

There are 12 recursive resolvers accepting TLS encrypted queries at this point in time, with a new server at the Korean Internet Exchange, [KINX](#), being the latest addition. The IETF also ran an experiment for DNS over TLS during the IETF meeting. Adoption is still slow (so slow that Stéphane Bortzmeyer, Afnic, said he had been hesitant to push ahead further work on the recursive to resolver privacy documents). However, a representative of Dyn said during the Prague meeting that at least he had the intention of also working on implementation. Hackathon implementation work can be found [here](#).

## DNS over TLS Software

Funded by the Nlnet, the DNS privacy project continues to track implementation in the software packages for recursive resolvers – see overview on [dnsprivacy.org](#). Both Unbound and Knot check most of the boxes. Ondřej Surý confirmed that serving TLS requests upwards from the resolver was on the to-do list for Knot to check another box. For the time being, for BIND, a stunnel proxy is necessary to implement DNS over TLS.

On the Stub side, work on stubby is advanced by Sinodun with packages worked on including Mac, Microsoft and Linux. An easy to use GUI development underway (that will allow to turn DNS over TLS on the laptop, desktop) will become available around the next IETF. A GUI for Linux is not planned for the moment, due to cost constraints and the idea that Linux users would be geekier and able to use the command line Stubby version.

Also presented during the Hackathon at Prague was an Android DNS over TLS implementation. Ben Schwartz from Google's New York office did a demo of working DNS-over-TLS in a custom Android build during Bits-n-Bytes. The DNS-over-TLS work is being done in AOSP (Android Keyboard).

Despite these steps, implementation remains slow and experts describe the situation as a hen-and-egg situation with large implementers like Google (which already is using DNS over HTTP) waiting for more demand from users, while users waiting for the large DNS providers to step up. Answering a request from this report, Lennard Poettering, Red Hat and Linux sd lead, explained that systemd-resolved was not intended to take the lead of DNS development,

but wanted to be a good client implementation of successful DNS technologies. DNS/TLS had not arrived at that point yet, as there were not enough implementations. "If DNS/TLS are implemented broadly, we can support it directly in systemd-resolved." While security was important and DNSSEC, for example, had been implemented, DNS over TLS was not there yet. Questioned about what constituted success, Poettering gave the following examples: if Google used it for its public DNS servers, if Deutsche Telekom used it on the DNS-Server for T-DSL, if the DNS Proxy of a FritzBox used it, or the Red-Hat Server for a VPN.

## DNS over HTTP2

DNS over HTTPS is an idea nurtured by the "browser people". It tries to make DNS more fully available to applications. Paul Hoffman (ICANN) said that the motivation behind DNS over HTTPS is that "web browsers can easily only deal with IP addresses, apps can only deal with IP addresses". Web-based applications wanting to use DNS features like DANE, DNSSD service discovery currently have to use browser extensions. At the same time, DNS over HTTP2 was the most practical mechanism for reliable end-to-end communication. TLS provided integrity and confidentiality, and HTTP eased transit through proxies, firewalls and authentication systems. Hoffman and Patrick McManus (Mozilla) propose to use "GET" or "POST" to wrap the DNS queries (in either message body or body). Using the GET method is friendlier to many HTTP cache and was smaller.

According to the draft:

*A query for the IN A records for "www.example.com" with recursion turned on using the GET method and a wireformat request would be:*

```
:method = GET
:scheme = https
:authority = dnsserver.example.net
:path = /.well-known/dns-query? (no CR)
```

```
content-type=
application/dns-udpwireformat& (no CR)
body=
q80BAAABAAAAAAA3d3dwdleGFtcGxIA2NvbQAAQAB
accept = application/dns-udpwireformat,
application/simplesdjson
```

The same DNS query, using the POST method would be:

```
:method = POST
:scheme = https
:authority = dnsserver.example.net
:path = /.well-known/dns-query
accept = application/dns-udpwireformat, application/
simpledns+json
content-type = application/dns-udpwireformat
content-length = 33
```

```
<33 bytes represented by the following hex encoding>
abcd 0100 0001 0000 0000 0000 0377 7777
0765 7861 6d70 6c65 0363 6f6d 0000 0100
01
```

The draft is not the first on DNS over HTTP, [see earlier proposals](#) from Hoffman, with earlier versions using http instead of https. Now https (http2) is supposed to be the primary choice. While Hoffman was highly critical of the DNS over Quic development, warning against potential confusion over implementation of DNS over TLS, he said DNS over HTTPS was a way for browser and app people to use DNS more fully and securely. The assumption also was that DNS over HTTPS could be implemented quickly by somebody who was running a large web service without a large additional effort to run a recursive resolver over their http2.

The interest at Google is documented through Google's implementation of [DNS over HTTPS](#).

On the other hand, privacy was not the primary concern, even if it was the case that if DNS over https got implemented "all web traffic would be private." The proposal by the WG Chairs of the DISPATCH WG to send the draft over to DNS Privacy nevertheless was rejected by Hoffman.

Regarding the concern that more information would end up at the browser vendor's data centre, Hoffman said that currently, people were not choosing their recursive resolver anyway and in turn, did use services like Google.

## DNS over Quic

DNS over Quic is the newest alternative for DNS transport. Introduced during the DPRIVE WG by Christian Huitema, the argument for it is that it will combine DNS over TLS like encryption features with advantages for the transport, especially by allowing ORTT connection resume and elimination of head of

line blocking.

With the Quic WG is currently still working on the base specifications – and http transport as adopted milestone for the WG – the addition of other protocols to the Quic agenda was rejected during the Quic WG session. Huitema argued that Quic should not be specified without keeping an eye on other "transport-customers".

During the DPRIVE WG, one major concern regarding the proposed draft was that it was geared toward the stub resolver to recursive resolver path only. For Quic transport, this split should not be reiterated and was perhaps not necessary, said Andrew Sullivan. The parallel offers of different DNS transport variants could also confuse implementers.

## Venues: IETF changes meeting venue to avoid the US

Just a day before the IETF meeting in Prague, the IAOC published its decision to change the meeting venue in July 2018 to avoid potential issues with immigration to the US. Having chosen San Francisco as the venue for IETF102, the IAOC had to cancel the contract with the conference hotel and negotiate a new contract, but will be able to recover the cancellation fee when going back to the SF hotel for a meeting during the next years. The meeting will now take place in Montréal, Canada, a week later than originally planned for San Francisco.

The reason for making the decision given by the IAOC was the unclear situation for immigration to the US after changes to US border policy and subsequent Court decisions. This had resulted in an "atmosphere of uncertainty", so the IAOC decided to play it safe and change the venue. Results from a survey also showed that 15 percent of 211 respondents had decided not to travel to the Chicago meeting in March 2017.

Interestingly, Prague was chosen once more as a venue for a European meeting in 2019. For North American meetings, Canada could become the prime venue, as it has been for several years in recent times.

## Meeting venue selection policy

As a conclusion from the debate over meeting venues, a [draft document](#) has been prepared to formalize the future selection of venues. The debate was triggered



by a complaint from IAB Chair Ted Hardie over the choice of Singapore as the venue for IETF100 in spite of anti-LGBT legislation in place there, but was also influenced by the debates over the US immigration policy hiccups.

The Meeting Venue draft includes the declared aim “to minimize situations in which onerous entry regulations inhibit, discourage, or prevent participants from attending meetings, or failing that to distribute meeting locations such that onerous entry regulations are not always experienced by the same attendees.” The document also calls to avoid meetings in countries with “laws that effectively exclude people on the basis of race, religion, gender, sexual orientation, national origin, or gender identity.”

The draft RFC includes criteria for venue and hotel selection and a step-by-step process for the process. Early publication of potential meeting venues and an invitation for community comment have been included to allow for greater transparency and participation. The document is still under discussion.

Roles in the selection process have been proposed for:

- 1. IETF Administrative Oversight Committee, IAOC (oversee and select IETF meeting venues, instructing IAD to work with ISOC to write contracts, making sure participant concerns about particular venues are weighed);*
- 2. IETF Administrative Support Activity, IASA (performing the meeting selection process under the oversight of the IAOC);*
- 3. IETF Secretariat (part of the IASA under the management of the IAD);*
- 4. IETF administrative Director (coordinates and supports the activities of the IETF Secretariat, the IAOC Meetings Committee and the IAOC, managing meeting budget); and*
- 5. IAOC Meeting Committee (participating in venue selection process, tracking meeting’s sponsorship program).*

## **IASA 2.0**

While the roles are fixed for the venue selection process, the very structure of IASA is also under discussion with a design team talking about how to potentially restructure the IETF’s administration. With these restructuring debates under way, there will be no immediate replacement for Ray Pelletier, first and long-time IAD (and person in charge of preparing meeting venue selection). Pelletier is stepping down from his position before the Singapore meeting.

The three options for the future administration presented by the design team in Prague were: (1) IASA PlusPlus (minor changes, keep structure in place); (2) ISOC subsidiary; or (3) Independent organisation.

# Working Groups and BoFs

## NetSlicing BoF and 3GPP

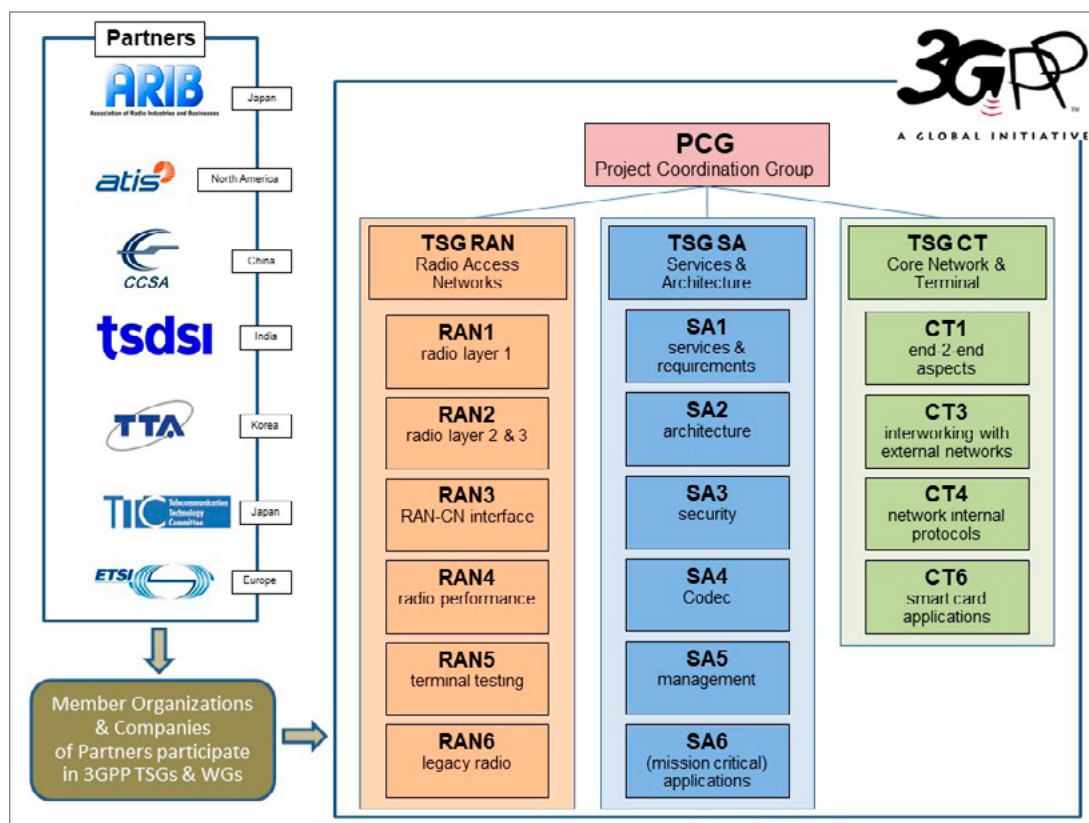
With a lunch meeting and a BoF on Netslicing, the IETF and mobile standardization body 3GPP tried to better align their work, especially with a view to the ongoing 5G standardization process. As 3GPP is re-using many IETF protocols, Georg Mayer, Chair of 3GPP Core Network and Terminals (for the structure of the 3GPP work, see graph below), called on the IETF to quickly come up with any issues the IETF developers felt had to go into the 5G core base specs. Yet 3GPP still has to decide which IETF protocols it will use, Diameter, HTTP1, HTTP2 or even Quic (although Mayer said it might be too late for Quic). In any case, the base spec (Version 15 of the 5G document) had to be ready by June 2018. The timeline is also fixed due to preparations for the next [ITU World Radiocommunication Conference](#) in 2019, when use cases for the next round of frequency allocations have to be defended.

## Network slicing

At 3GPP, network slicing is a core concept for 5G expressing the idea that instead of a single network in the future, slices are expected to be cut-off for different users and use cases, also allowing for different traffic quality. The network slices concept was a reaction to the differentiation of customers: instead of telecom operators and their customers, there now were sensor network operators, the car industry or users of drones or smart city architects, each with a distinct set of requirements for their networking. Instead of one-network-fits-all, slices could be offered to these customers.

The BoF about network slicing mainly revealed that 3GPP and IETF's concepts of network slicing were vastly different. Ted Hardie, IAB Chair, pointed out that there was a need for both organisations to agree on common definitions on slicing, as well as other terms. Next step in the attempt to better cooperate might now be a draft on terminology.

For more information on the 5G standardization process, see [requirements](#). Some of the base specs also include reference [23.501](#) and reference [23.502](#).



## Security Area Advisory Group – Chances of quantum computing

Kenny Paterson, Co-Chair of the IRTF Cryptoforum and well-known Crypto expert from Royal Holloway University of London, challenged the quantum computing hype and even offered the provocative idea that the focus given to (and money spent on) quantum technologies could be a distraction from the fact that cryptography on the net was not in good shape nowadays. Paterson pointed to large scale measurements that exposed “lots of Diffie-Hellman parameters which have unknown provenance” and about “major flaws” discovered in TLS due to bad implementations.

### Crypto-Apocalypse: now or never?

On the other hand, a potential “cryptocalypse” through quantum computers was difficult to predict. Predictions from the last decade that quantum computing was imminent had in fact not materialized.

The IBM announcement of a 17 Qubit machine in 2017 rather illustrated persistent limitations. A 17 Qubit machine was a remarkable engineering success, given the instability of Qubits to noise, heat, and other inconvenient factors. Still, 17 Qubits would not allow to threaten state-of-the-art crypto. According to Paterson, the only people who were able to say how advanced the technology really was were those working in the field, yet they had also some self-interest in keeping the field alive.

Paterson told this reporter he would estimate that there will be a quantum computer in his lifetime that was capable of factoring 1024 bit RSA numbers at 10 percent. On the other hand, he is member of one of the teams filing a proposal to NIST's Post-Quantum Crypto Project ([deadline is 30 November 2017](#)). The project aims to identify a quantum resistant new crypto algorithm. NIST has asked for public key algorithms, digital signature schemes and key exchange mechanisms, so all PKI techniques.

Paterson said he expected proposals from the various camps: isogeny-based elliptic curves, lattice based curves, code based curves and multi-variant equation based curves. “I think we will see all four boxes being populated and then it gets very interesting on how to choose between these things, because they are radically different in terms of performances, maturity and trust levels”, Paterson said to this reporter.

Quantum resistant technology, according to him, were hash based signatures, as they had very nice properties and were well understood, while much less cutting edge (or bleeding edge, as Paterson said).

To questions regarding trust in NIST after the manipulated Dual ECRNG, Paterson pointed to the completed Post-Snowden NIST review and the subsequent organisational changes. NIST had for example decided to hire more crypto experts of its own to become more independent from other agencies, namely the NSA, he said. Moreover, he expected the crypto community to watch the quantum resistant crypto selection “like a hawk”.

### Post-Quantum Crypto focus – a distraction?

Paterson agreed that there was a small probability that quantum computers would make all public key cryptography worthless, but still he talked about all the money, time and efforts thrown at it as worthy of reconsideration.

During the Prague IETF meeting, Stephen Checkoway, as member of a larger group of US crypto experts including Matthew Green and Eric Rescorla, presented their findings about the vulnerability embedded deep in to Juniper NetOS devices. As a follow-up to Juniper's announcement that the devices were accessed and a parameter used for the computing of Random Numbers for Ipsec had been changed, the researchers started a large reverse engineering project to go back firmware by firmware until they found a whole set of changes to the NetOS software, put in together in 2009. The researchers found that, contrary to Juniper's assertion, X9.31 (a second PRNG) was never used due to the reuse of the output buffer and the global index variable. Otherwise, the manipulation of parameter Q would not have mattered. The combination of vulnerabilities, all put it, deliberately or not, in the same firmware version in 2009, made all VPN traffic emanating from the NetOS devices vulnerable.

Paterson therefore concluded that “it is helpful to direct a whole amount of scientific academic research resources into something (quantum computing) that may be important one day or may not be, depending on what happens with large scale quantum computing. But it keeps us all busy here, while the real action in terms of securing the internet is over there.”



## DNS

Different variants for DNS transport were discussed during the Prague meeting in various WGs (see “Highlights” above), causing OPS Area Director Warren Kumari to joke about the need for a new working group called “DNS over new Transport” (DONT).

While these various transport proposals and the DNS over TLS implementation seem to go in the direction of privacy, there are also several drafts currently on the discussion list of the DNSOP WG that seem to push in the opposite direction (see also highlights above). The DNSOP WG once more looked into a rather long list of proposals, and is considering to hold mini-interim meetings on individual items.

A draft related to the transport topic is the one on [session signalling](#). It is supposed to enable session instead of per-packet signalling to reduce overhead resulting from the per-packet signalling mechanisms (EDNS0). While there is still some discussion over the opcode format, the authors from ISC, Apple, Sinodun and Salesforce do propose a new format with TLV (type-length-value, instead of RR Codes). The new format will require updates to all sorts of tools (logging, storage formats, any tool that wants to parse a DNS message). The RFC will give a first list of TLV messages.

Sara Dickinson from Sinodun said during the presentation that in a way, the session signalling will change the standard DNS message format (RFC 1035). There were even questions if this kind of development would trigger questions about a more clear-cut design for a DNS2.

Ondřej Surý, CZNIC, said there was an issue with implementing a new format, but also underlined that a lot of DNS improvements like padding would be left aside. Christian Huitema asked how the request to process messages “in order” will fit with Quic (or even UDP) that process the messages as they are received. The draft is still to be discussed further and could also become a topic for a mini-interim.

Issues on the DNSOP agenda also included status updates on several drafts.

“[RFC 5011 Security Considerations](#)” clarifies waiting times for the use of new DNSSEC keys when these are rolled. It describes the “math behind the minimum time-length that a DNS zone publisher must wait before signing with only recently added DNSKEYs

and also the minimum time-length that a DNS zone publisher must wait after publishing a revoked DNSKEY before he can assume that all active RFC5011 resolvers should have seen the revocation-marked key and removed it from their list of trust anchors”. The draft in its current version clarifies that there is an add-wait time (about how long one has to publish a new key before one can only use that key) and a remove-wait-time. The current root zone KSK roll is already under way. According to the security considerations, it has 30 days add-wait time, old rrsig validity 21 days, old dnskey TTL 2 days. With the math of 5011 it would mean 56 days. That is much shorter than planned by ICANN for KSK roll. The same is true for the revocation time: according to the new 5011 times, the hold-down-time would be 26 days (the 30 days add-wait-time are subtracted). The ICANN announced time for this is 70 days.

Questions brought up included if an interval or wall clock time should be used for the add wait time. Wes Hardacker said he thought it was ready for last call.

[DNS packet capture format \(C-DNS\)](#) wants to ease storage and transmissions of large packet captures by pairing questions and answers and minimizing the size of the packet capture files. At the same time, full DNS message contents along with the most useful transport meta data shall be preserved. The capture format is intended to help traffic monitoring applications. Open questions included what to do with malformed packets. The authors also asked the WG to make the case for additional use cases (and data they want to be captured). The IPR situation (IPR lies with ICANN) remains unclear.

The “[DNS terminology](#)” draft was briefly discussed and Paul Hoffman asked for additional review to finalize it. This document will be the successor to RFC 7719, and thus will obsolete RFC 7719. There might also be a need for a third one. The DNSOP Chairs considered the terminology draft as one of the potential items for a mini-interim.

An intended update to RFC 2845 was discussed as a reaction to the recent TSIG vulnerability in BIND and Knot. DNS vendors getting together at Prague decided that section 4.5 was the source for the implementation error leading to the TSIG vulnerability.

Algorithm negotiations in DNSSEC shall allow DNS clients to specify, in order of preference, which

algorithms they want to use. Responding servers shall use the client's preferred algorithm that they support. It shall allow supporting choice of algorithm and algorithm flexibility at the same time.

A well-updated list on the many documents under consideration in the DNS WG can be found [here](#).

## DPRIVE

In DPRIVE, apart from DNS over Quic (see Highlights above), there were notable presentations on padding and Demux.

For the padding draft, presented by Alex Mayrhofer, nic.at, there was a brief discussion about the added [“recommended strategy”](#). After feedback, preferability (or not) of totally randomized padding that would prevent analysis on block-counts was briefly discussed, but rejected for now. Daniel Kahn Gillmor (ACLU) who [analysed impact of padding](#), looked at real packet traces from Surfnets resolvers, applied a simulation of padding to the packet and came out with recommendation that you should pad queries to a block size of 128 bits, and answers to 468 bits. Padding to 128/468 means that 93 percent of packets have the exact same size. “Cost” was also calculated by DKG based on an attacker who is interested in one query-answer – how much of the other packets will be in the same size bucket. Random padding, which some recommend, would be much harder and could easily become pseudo-random and leak data. In the future, padding policies might change, but the draft would be a good and necessary starting point to avoid implementers going into different (and perhaps failure-prone) directions.

The DPRIVE WG also considered a controversial proposal by DKG (ACLU) for DNS over TLS to “squat on port 443” to make it indiscernible from https traffic. A simple “demultiplexing server” should distinguish between DNS and HTTP packets arriving, based on the first few bytes sent by the client on a given stream; once a choice has been established, the rest of the stream is committed to one or the other interpretation. DKG has an implementation on a Debian server, but acknowledged it was “the horrible idea of the day” (and that DNS over https would be the better solution).

The TLS and DTLS profile drafts are in IESG review. The next step – privacy on resolver to authoritative path – still needs to be addressed.

## Homenet: simple naming document “close to done”?

After being stuck for some time due to the delegation of a “special TLD” the homenet WG now hopes to advance, but still has quite some issues to resolve. Relations to the DNSSD work are now described in a roadmap document that tries to give an overview of the status quo of services discovery in local and homenet zones. The homenet WG has a new Co-Chair, Barbara Stark from AT&T, who is stepping in for Mark Townsly (Cisco), with Ray Bellis (Nominet) remaining. Interestingly, Stark made a comment challenging the statement “the DNS is uniform”. She said the DNS is not uniform, or at least it was a bad concept (for the parallel name spaces of homenet and global DNS).

The new choice for the homenet domain, homenet.arpa instead of .homenet, is on its way to getting finalized after additional rounds made with some efforts going into the behaviour of homenet.arpa and DNSSEC. To address the dilemma that DNSSEC validating resolvers could drop queries for what they find as insecure delegations, [version 11](#) of “Special Use Domain ‘home.arpa’” bans recursively forwarding example.homenet.arpa queries “to servers outside the logical boundaries of the homenet with the exception of DS lookups for ‘home.arpa.’”

With the special TLD dispute finally settled, Ted Lemon tried to advance naming and discovery services during the Prague meeting. The draft he proposed to the WG to accept as a WG document, [“Simple homenet naming and service discovery architecture”](#), combines domain look-up on the internet, publish services reachable anywhere in the homenet and discover services in the homenet.

Lemon declared as non-goals for now:

- publication of a DNS zone for the homenet in the DNS
- making service discovery available off the homenet
- allowing off-homenet services to publish services in the homenet
- securing homenet using DNSSEC

He explained that he saw multi-homing as the remaining challenge, but all other issues were already mainly addressed in a set of documents, which were not discussed in detail during the Prague session:

+ [draft-sctl-service-registration-00](#) According to the draft, the DNS-SD Service Registration Protocol shall provide “a way to perform DNS-Based Service Discovery using only unicast packets”. It was “largely built on DNS Update [[RFC2136](#)] [[RFC3007](#)], with some additions.”

+ [draft-sctl-discovery-broker-00](#) “The Discovery Broker is an intermediary between the client devices and the Discovery Proxies. It is a kind of multiplexing crossbar switch. It shields the clients from having to connect to multiple Discovery Proxies, and it shields the Discovery Proxies from having to accept connections from thousands of clients.”

+ [draft-sctl-dnssd-mdns-relay-00](#) According to the authors, this extends the current discovery proxy for MDNS service discovery by describing a discovery relay “which allows discovery proxies to provide service on links to which the hosts on which they are running are not directly attached.” the two parts of the protocol are: “connections between Discovery Proxies and Discovery Relays, and communications between Discovery Relays and mDNS agents.”

There was some support in the WG, but also a considerable amount of discussion. With regard on multi-homing, Andrew Sullivan asked why a host should have a theory at all regarding which ISP to address in the multi-homing scenario. His concern, he explained, was that the way the naming document works right now, the host (at least, and maybe the application) needs to have a theory about which ISP is going to be used for a connection. This was not going to happen, and ideas about work that came out of MIF could help, was “mostly hope rather than a plan”. Significant work still had to be done on the naming document.

New Co-Chair Barbara Stark and David Schinazi offered the idea that happy eyeballs could solve the problem.

Juliusz Chroboczek and others talked of “many moving pieces” in the documents. The roadmap prepared by Stuart Cheshire, presented in DNSSD, called for a major change in the field of the local/homenetwork by moving away from multicast toward completely unicast.

Another presentation by Lemon addressed questions of encryption in the homenet. While there is no written draft as of now, Lemon did propose that each of the nodes in the homenet should generate a public-

private key pair and distribute its public key parts to the homenet nodes. This would allow using DTLS, instead of only having the shared secret approach now available in the Homenet Control Protocol (HNCP).

While he underlined that the PKI approach alone did not bring security, it would allow identification of the nodes one talked to. There was some support for addressing the problem. Some participants recommended to consider baking encryption/security into the base specifications. Lemon argued that the keying draft could be finished much faster, though, so a split might make sense.

An attempt to implement homenet for a dual-homed homenet scenario presented by Lemon revealed considerable problems. From the three available HNCP (homenet control protocol) implementations (hnetd, pysyma, shncpd), Lemon chose hnetd on OpenWRT and shncpd on Ubuntu. After the setup, the OpenWRT router lost its upstream IPv4 address because DHCP was no longer providing it, and “the IPv4 RFC1918 prefix I’d allocated was de-configured on all interfaces for no obvious reason”. While several participants pointed out that OpenWRT was working well for them, some recommended to look deeper in the issues.

## Homenet and DNSSD – still different?

DNSSD and Homenet seem to get closer in their work items, at least sharing common assumptions: DNSSD is trying to make service discovery work well across more than just the link-local network, while Homenet takes as a basic assumption that the “homenet” could be (not must be, but could be) multi-homed, in which case some things are not *\_always\_* on the local link.

## DNSSD – Departure from multicast and better privacy

At the Prague meeting, Stuart Cheshire recommended to move away from multicast for DNSSD service discovery and brokerage. Multicast should still be supported in order not to leave devices that have been using it for 15 years behind. But the concept was overburdening larger enterprise networks with hundreds of Wifi clients and also home networks (homenets) with more than one link. The group therefore should break away from multicast and embrace unicast only.



In a roadmap, Cheshire describes the scenery for the evolving DNSSD services, with service discovery and service registration, and especially a new concept for a central “discovery broker”.

The broker presented by Ted Lemon (Nominum) and Cheshire shall allow the bundling of multiple domains into one and directing queries from different clients to it, instead of different discovery services. The broker would function as a meta-discovery server, taking on the queries from outside servers and querying the various discovery servers for them. To the discovery servers, the meta-discovery/bundle/broker server looks like a client. This introduces more of a hierarchy into the local/homenet. Cheshire went as far as saying that two drafts catering to extend multicast advertisements across links should not be pursued.

Nonetheless, backwards compatibility should be ensured in the discovery functions. Cheshire and Lemon are considering to split up functions for discovery into a discovery broker and a discovery relay, with the latter being modelled along the concept of enterprise network core DHCP servers.

For registration in the name space, legacy mDNS devices shall be able to use a DNSSD hybrid discovery service, ([draft-ietf-dnssd-hybrid-06](#)). According to Cheshire, the document is dependent on DNSPush (for asynchronous change notification instead of polling) which in turn is dependent on DNS Session signalling, currently under heavy debate in DNSOP (see below).

For active registration of devices, a new services registration protocol is put into another new document ([sctl-services-discovery](#)) by Cheshire and Lemon. It is based on DNS update (RFC 2136). An EDNSO option is used to specify what additional information should be carried (sleeping server and wake-on-LAN magic packet bit pattern that can be used to wake it up) to allow power efficiency, or power saving respectively.

Finally, participants of the CoRE WG presented work on CoRE Resource discovery and promoted a [document on a mapping of DNSSD discovery and CoRE discovery mechanisms](#) already in use. CoRE resource discovery and the coarser grained DNSSD service discovery were complementary “in the case of large networks, where the latter can facilitate

scaling”. According to the authors, this document shall define “a mapping between CoRE Link Format attributes and DNS-Based Service Discovery [RFC6763] fields that permits discovery of CoAP services by either method.”

## Better Privacy

Before the presentation of the next round of drafts for the “new” DNSSD path, Christian Huitema presented the two privacy-related drafts for DNSSD. To enhance privacy, nodes publish instance names (hashes) for every pairing they have. Using the shared secret, they start TLS sessions. This concept was preferable to a PKI system, Huitema said, working down the issue list, because the public key was a unique identifier and would be disclosed during the TLS handshake. PKI provided implicit client authentication. The shared secret concept, on the other hand, allowed for an anonymous exchange.

Other issues discussed were the time slots synchronization, with the nodes publishing instance names every five minutes. Synchronization to about 4-minute intervals was necessary to avoid the nonces/hashes getting stale. There was a question about how much the time interval could be stretched for ease of use. Huitema defended the choice of the short interval, the time based nonce controlled the computing load and mitigated DOS attacks, he argued. Stretching the validity would allow tracing devices on their way through networks, he said. To ease potential problems at the edges of the interval, for the first minute of a new interval the old hash had to be checked for the last minute of an interval there has been a check for the new hash.

More possibilities for fingerprinting results from the publication of as many instances as one has pairings. By counting the number (and matching it to published instances), fingerprinting might be possible. A potential countermeasure could be padding with fake instances. An additional draft proposes to use QR codes as an alternative for discovery and verification. The question for the WG was if QR codes should be made part of the main draft or should remain in a dedicated draft. The WG still has to make the decision how to split, with potentially three separate documents possible: problem analysis, specification of pairing and specification of QR codes.

**IETF100 will be held on 11-17 November 2017 in Singapore.**



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](mailto:secretariat@centr.org)  
[www.centr.org](http://www.centr.org)



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