Council of European National Top Level Domain Registries

# Report on IETF 93 Prague

20 - 24 July 2015

# Table of Contents

## Highlights 3

## Working Groups and BoFs 10

## IETF News 13

# Highlights

## Two highly different surprise guests: Edward Snowden and Houlin Zhao

A surprise guest, Edward Snowden, joined the participants of the IETF meeting in a one-hour question and answer session following a screening of the documentary Citizen Four. Snowden recommended to the IETF community to standardize neither for governments', nor for corporations' convenience, but instead „make sure that standards, our technology, the systems are working to protect and armour the users' intent.“

### End-to-End, less metadata, next generation DNS

The way forward for standardization, Snowden told, was to think about the classical standard principle of end to end security. This had been compromised and the „idea of a simple core and smart edges had changed to very dumb edges and a deadly core.“ Warning that all middle boxes put in the way of end-to-end communication in the end created potential vulnerabilities, the former NSA analyst got rounds of applause. Snowden warned, the network path „currently is the most dangerous part of the network“. Engineers had „to help the users to get safely across it“.

On meta-data Snowden said, on the way to „pervasive encryption“ which he expects to be very much in place in 15-20 years, meta data should not be forgotten. Intelligence Agencies and companies still would find all kind of reasons to use metadata to track and profile users by association or location. Quoting ex-NSA President Michael Hayden, Snowden said: „We kill based on metadata“, referring to drone strikes by the US military.  As narrow as originally planned, surveillance technology easily got commercialized once former NSA employees left the agencies and started their own business, bringing the very technology to the market.

„If it's creating more metadata, this is in general a bad thing", he said with regard to new protocols. He did mention the recently discussed ideas in SPUD as one negative example that would create another layer of metadata source on user's intent.

As another source of metadata Snowden also mentioned the DNS specifically. If content were encrypted everywhere, but DNS requests were not, users could still be identified via statistical inference. Snowden therefore welcomed the efforts of the DPRIVE and DANE working groups and called the deployment of DNSSEC, while not a golden egg, „better than what we have now“. „When you combine this (DPRIVE, DANE and DNSSEC), you create the next generation of DNS“. In developing these standards, engineers had to consider to „build an Internet that does not only survive for a few years, but for a hundred years or more.“

### Separate „identity“ from „person“

His major recommendation with regard to privacy was: „We need to divorce identity from persona – in a lasting way“.  Snowden recommended to consider alternative naming systems. Contrary to credit card payment, for example, concepts for anonymous payment had to be developed and applied, for example paying for tokens. Users should have a choice to use a „common persona, non-persona or anonymous persona“ for transactions on the net, Snowden demanded.

Globally unique hardware identifiers also constituted a problem for privacy, he said, and also questioned the current use of MAC addresses.

## Stick to the principles

Apart from his critical comments on mass surveillance, Snowden welcomed the efforts of the IETF and called the internet standardization project more open than it had ever been before. If technology becomes a danger it is because we left it to somebody else, he said during the Q&A.

He finally raised concerns about „man-in-the-middling" against the internet standardization process real concerns, reminding on the NSA's interventions on the Dual_EC_DRBG standard and NIST's subsequent withdrawal of it from its draft guidance on random number generators. He had, Snowden said, observed during his active time at the NSA and CIA people being sent into large organisations, not standardization bodies, but large infrastructure companies. Those „mules" had been told that they had to steer the organisation away from this and steer them to that because the latter was more secure, when it in fact had a back-door for the agency. The mules were, in fact „doing bad things for good reasons", and instead of being suspicious against the „guy with the moustache", Snowden proposed the focus on the very principles, like protecting users' intent. If something seems to be suspicious, instead of trying to analyse who was behind it and for what, engineers should just scrutinize the work and find the weaknesses to block it from being adopted.

Snowden received a standing ovation by the attendees and during the IETF week many people called his appearances the highlight of their IETF week, even if it formally was no IETF event, as Mark Nottingham underlined. An invitation to the official technical plenary was obviously seen as too bold a step.

Nottingham had co-organized the event together with Daniel Kahn Gillmor (ACLU). There are now also considerations to invite the whistle-blower to the next W3C meeting.

## Let's not fight, asks the Secretary General of the ITU

The second surprise guest visiting the IETF was the new Secretary General of the International Telecommunication Union, Houlin Zhao. Elected last October, 23th by the ITU member states to succeed Hamadoun Touré, Zhao made his inaugural visit to the IETF, for which IETF together with ISOC had invited him already for IETF 91 in Hawaii. The career diplomat had been working before with the ITU-T (the standardization sector of ITU) since 1986, was Director of the Telecommunication Standardization Bureau (1999-2006) and ITU Deputy Secretary General (2007-20149).

Zhao spoke in Prague during the Internet Architecture Board Plenary. His main message to the IETF community was a call for cooperation between the different standard bodies, regardless of their respective models, intergovernmental or bottom-up. He said: „Some might see the ITU as top-down, but whether top-down or bottom up, we serve the same market." The organizations should avoid to fight for their respective sides, he called on the IETF, never addressing any of the recent disputes like the one one fought over the ITU fork of MPLS.

Reminding the attendees about the formation of the Protocol Supporting Organization, Zhao said, perhaps it was time to open a new page in the relations between the organizations. The PSO, built by IAB, ITU, W3C and ETSI, and meant to rotate directors to the ICANN Board, was disbanded in November 2002. The follow-up body was the Technical Liaison Group, and a director seat downgraded to non-voting, but a proposed Bylaw change finally established an IETF-only liaison to the ICANN Board in 2013. ITU, ETSI and W3C are no longer directly represented in the Board; the IETF-ICANN relation has been strengthened during Fadi Chehade's tenure.

Given these developments, the ITU Secretary General's visit which was greeted with polite applause, albeit rather cool by the IAB attendees, felt like an attempt to regain ground. Yet instead of providing concrete proposals for „cooperation" Zhao took to the symbolic gesture of donning the T-shirt of the IETF 46 meeting over his tie.

# Next round in Special Names controversy

A special name reservation for the Tor project for the name .onion is in IETF last call with comments due until August, 11. But the battle over how to interpret and implement RFC 6761 – Special Use Domain Names - raged on the mailing list before the meeting of the IETF DNSOP Working Group and after. The DNSOP Chairs announced during the WG meeting in Prague that they intend to install a design team that is expected to go over 6761 to clarify questions the WG started to discuss going over the special name applications. For some, it's an attempt to close the door to alternative name registration before the IETF gets in some political trouble with ICANN.

During the DNSOP session in Prague Christian Grothoff, INRIA, presented the applications by Tor (in addition to .onion also .exit, and possibly .tor) and a group of P2P-projects including the Gnu Name System, GNS, (.gnu), the Namecoin Project (.bit), the I2P-Project (.i2p). All are based on protocols distinct from DNS, but using name strings as identifiers. Grothoff explained that the P2P projects had been out there for various spans of time and wanted to use the RFC 6761 process to document the usage of names in an attempt to possibly avoid collisions with DNS names that might be delegated by ICANN in the future. „Would it better if we did not delegate?", was his question to the IETF community. He announced that more special name applications from the P2P community were on their way.

## Technical Community split over non-DNS names reservation through IETF

DNSOP and the IETF at large are clearly split over how to deal with the special names applications based on RFC 6761. The RFC had been introduced by Steward Cheshire, Apple, for the reservation of .local (RFC 6762). The author of the two RFCs defended the process laid out in the RFC. Just as for .local, there were names that were not DNS, but would only resolve locally, a point also underlined by Grothoff who pointed to Microsoft's *.ipv6.literal.net* and *pnrp.net*. Cheshire argued that instead of pretending that „we are the only game in town, we have a role of stewards for the shared namespace".

Stephane Bortzmeyer, AFNIC, recommended to consider the effect a rejection would have on those who developed and used name systems outside the IETF. If those who came to the IETF to document their protocol would be rejected in the first instance, it could result in a loss of trust in the standards body. Bortzmeyer later wrote in the mailing list that he had been „shocked that we have a 'design team' to work on RFC6761bis while we still do not have a consensus on the problems with 6761." Bortzmeyer belongs to those highly critical of a potential back flip on the special names reservation process.

On the other side of the spectrum is, for example, IAB Chair Andrew Sullivan, who got emotional during the DNSOP session calling some of the proposals „attacks on the way that the DNS works. Using the domain name space to subvert the DNS is a bad idea." He did not see a reason why the IETF should grant reservations to those trying to "compete with other people's business models", he added calling for separating the proposals and dealing with them accordingly. To those proposing a „new global identifier system I have to say, sorry, we already have one", he explained to this reporter.

Sullivan, in a lengthy email on the DNSOP mailing list later, qualified his comments with regard to the intentions of the proponents and explained why he was opposed to some names:

*„So, just as local. marks something as to be looked up only with mDNS, onion. and exit. both mark something as to be looked up only under onion routing (or maybe, depending on your view, only using Tor). But others of these proposals, such as bit., mark out a name space and associated protocol that competes with the DNS. It is a fully parallel name resolution universe, applicable to absolutely any network application. My point was that the second class of these basically puts us in the position of approving a special-use registration that is effectively an attack on someone else's business model (ICANN's and that of the various registries and registrars). I believe that draws the IETF into a political battle for which it is unprepared, and that's really why I object to these registrations."*

The „political battle" is a concern shared by a considerable number of IETF participants as the broader discussion on the IETF mailing list illustrates. In fact it looks as the major issue.

Proposals have been made to keep the „technical" reservations in a special zone (like .alt or .external, or even .ring). This idea is opposed by the applicants, as the names applied for have been in use. Would a .alt have been around years ago, Grothoff said referring to an applicant, one project might have considered using it. Changing the rules now in the middle of the game is on the other hand rejected by others like Bortzmeyer.

## Getting rid of 6761 - .local squatting was bad precedent

The Working Group nevertheless seems to be embarking on the road to re-open 6761. With the three-to-five member design team members to be announced soon by the WG chairs, Suzanne Woolf said. Woolf is the SSAC appointed liaison to the ICANN Board and also IAB member. The design team, she said, would not make any decision on a potential bis-version for the RFC, but would present results of the evaluation of problems that WG members had been coming across during the recent debate. David Conrad, ICANN, put it bluntly, calling RFC 6761 and 6762 „formalizing squatting on name space" and a bad precedent.

## Reservation process unclear

In what could be seen as a start of the design work Peter Koch, DENIC, and Alain Durand, Microsoft, dissected some of the issues to be discussed.
Even without a change of the RFC the new applications might fail, argued Peter Koch. The P2P applications, for one, failed to be standards track IETF documents (contrary to .local's RFC 6762). The seven questions to be answered by applicants to receive special names reservations would not, Koch said, serve as a justification to reserve the name.

Several things were unclear in the first place in 6761:

- **what is the process for the reservations? (need for standard track document)**
- **who decides? (the IESG, the DNSOP WG who in fact was not mentioned in the draft, or a special adhoc-WG?)**
- **what results from reservation? (guarantees with regard to non-leakage? Expectation of implementation by audiences mentioned in the draft)**
- **is there a need for a removal processed?**
- **will applicant pick the string or could IETF have a say in which string?**
- **relation of „single root" principle and alternative names (Sullivan's question)**
- **what kind of coordination with ICANN necessary?**
- **can IETF do more than protocol analysis?**

The process would, Koch and Durand literally quoted from section 3 of RFC 6761, create a „higher-level protocol rule, above ICANN's management of allocable names on the public Internet". This in essence seems to be the main fear of many, that the IETF could be drawn into the dog fights around new TLD delegations with a lot of money and lawyers around. Yet the IETF already is in a bad spot, as rejecting .onion after .local looks bad, not rejecting it makes it harder to reject other P2P TLDs. Stay tuned.

# Choosing new TLS signatures and eying post-quantum crypto

Following the Snowden revelations and acknowledgment of NIST on manipulation of certain crypto selection processes by the NSA the Crypto Forum Research Group (CFRG) had been tasked with choosing new crypto algorithms for TLS. After selecting Curve25519 and *Goldilocks* (Curve448) as new cipher suites, the Crypto Research Forum had yet another „beauty contest" to invite proposals for signature schemes. At the same time the WG decided unanimously to start considering post-quantum safe crypto algorithms. The CFRG continues on the path to establish itself as an alternative venue to choose the crypto for the public internet with NIST trying hard to not give in as debates at a recent workshop on Elliptic Curve Cryptography illustrated. The announcement by a NIST employee that the new SHA-3 hash had just been signed by the US Secretary of Commerce did not receive much attention in the IETF so far, one expert said during the CFRG session in Prague.

## Quick decisions on signature schemes

Five proposals for the new signature scheme were presented during the CFGR meeting in Prague, by Dan Brown (presented by Gaëlle Martin-Cocher, Certicom, Blackberry), Ilari Liusvaara, by Dan Bernstein (University of Illinois), Tanja Lange (University of Eindhoven) e.a., by Mike Hamburg (Rambus Cryptography Research) and Watson Ladd (presented by Martin Thomson).

Liusvaara and Bernstein both have come up with comparisons of core features of the proposals, including protocol usage, coding (for example inversion, random numbers or non-standard hash APIs), extras including personalization, firewalling and batch-ability, as Liusvaara described. Bernstein offers a python script to allow for a comparison of the five candidates.

It also tracks the changes to the five proposals. Bernstein on the mailing list did warn against „trying to achieve cryptographic balance through chopping every system down to the minimum size that can reach some predetermined security level". There was a risk to create loopholes for attackers.

CFRG Co-Chair Kenny Paterson promised after the presentation in Prague that a much quicker decision on the signatures scheme can be expected. The selection of the new ECC cipher had taken over a year, despite a sense of urgency. The signatures schemes, Paterson said in Prague, shall be decided upon before the next IETF meeting at least.

### Quantum Cryptography - Do not leave it for later

The CFRG also unanimously agreed to take on a work item looking into quantum computing, humming in favour of William Whyte's proposal on „cheap quantum-safe cryptography without breaking anything".

It was not clear when quantum computers would be available, Whyte, CTO at Security Information said, but it could happen well before the often estimated 30 years from now. With quantum computers RSA, Elliptic Curves and Diffie-Helman algorithms would all become easy to break. Other algorithms -based on different kinds of mathematical problems- are supposed to be resistant to quantum computers
and these could be embedded in cipher suites today. To go right to a quantum-resistant new cipher suite could solve the situation, Whyte said. But neither had the CFRG started to talk about that at all, nor would it be easy to implement.

Whyte proposed instead a „hybrid approach" combining a quantum-safe public key encryption/key-exchange with existing algorithms. The example he promoted was the NTRUEncrypt algorithms (patented by his company), which had been integrated into the ntor key exchange protocol, in an effort to provide quantum-safe secrecy for Tor. Ntor could be instantiated with 25519, for example, too.

Potential alternatives to NTRUEncrypt were the quantum-safe public key crypto systems „Learning with Errors" and McEliece. The latter had very large keys, though. Performance figures from NTRUEncrypt-ntor were satisfactory with regard to the overhead for the client (who bore most of it) and the router. The existing patent met with some reservation at the CFRG session.

Another implementation of the hybrid concept is an „extension" to TLS, featuring a „Quantum-safe hybrid cipher-suite identifier (QSH)" and „extensions for quantum-safe public key and cipher-text". Running code can be found here, according to Whyte. Again NTRUEncrypt has been used.

Questions about the need to push for quantum safe crypto, when existing cipher-suites were still seen as safe, were answered by Whyte and researcher Tanja Lange. Communication secured with the state-of-the art crypto was susceptible to „harvest-then-decrypt" attacks. Lange said if people did want to have sensitive information still to be safe in 10 years, they had to go ahead. She pointed to a EU Horizon 2020 3,9 million Euro-project of eleven partners on post-quantum cryptography that just started in March. NIST held its workshop on cybersecurity in a post-quantum world in April. Whyte gave another practical example on the need to tackle the post-quantum crypto during the IAB plenary on car-to-car communication and the security for it. Cars, he said, had a life time of a decade or more.

## DNS „pretty bad privacy"

DNS privacy is even worse than its reputation, according to Haya Shulman, researcher at the University of Darmstadt. Shulman explored side-channel attacks on DNS queries that persist even after the implementation of DNS over TLS and she also questioned the introduction of TLS based on a lack of support of servers for TCP. Based on measurements over 50K-top Alexa domains and 568 TLDs deployment obstacles for DNS over TLS could be expected for „at least 38% of servers and 12% of TLDs. Shulman due to this fact questions the efficiency of DNS over TLS and similar mechanisms and recommends further studies. She received the IRTF's Applied Networking Research Prize (ANRP) for her paper.

## Side-channel attacks on DNS and what to do against it

Does it even help, if you encrypt queries and answers?, was Shulman's answer to the privacy mechanisms she briefly presented (including DNSCurve, DNSCrypt, DNS over TLS and opportunistic IPsec). Often having the IP address suffices to allow an accurate guess about which domain is accessed by correlation of the destination IP address in the DNS request with the name servers. Co-residence of domains (up to 500 for some servers) can blur the picture.

But other side channels still help to refine the guess, especially request and response sizes, latency between request and response or other, more specific side channels, especially „transitive trust dependencies". The fact that a domain often is served from more than one name server to allow for redundancy and resilience and the name servers in turn are configured to be under different domains  provides an attacker with a dependency graph. Once an attacker has a database about such dependencies he can match queries/answers to it and thereby de-anonymize them.

One possible answer to the side channel attacks and that kind of „fingerprinting" of DNS queries and responses was presented during the DPRIVE WG in Prague: „Padding" - the addition of extra bits to the encrypted queries (before encryption) allows to help blurring the size of requests. „Padding" has been presented as a privacy-enhancing option to TLS in the TLS WG by Daniel Gillmor (EFF), but could, Alex Mayrhofer, nic.at, suggested, be used for the DNS over TLS schemes as an alternative. TLS padding, as Gillmor said, might still take some time to be implemented as a feature of the just to be developed version 1.3 of TLS. It is not clear if it will be implemented in TLS 1.2.

Potential truncation due to the size of the „padded" DNS queries has been mentioned as a security consideration.

## Problems with encrypted DNS

Shulman's second measurement activity checks on the readiness of the installed DNS system for the encryption efforts. She points to several problems: one, the non-interoperability with caches could result in prohibitive traffic overhead for servers; two, encryption could be averse to hiding an authoritative name server behind a recursive- authoritative name server. Finally TCP protocol support in installed software is a limiting factor, according to Shulman.

Shulman reports about a variety of failure modes not only on the client, but also on a considerable number of name servers. Even popular domains were affected, she said. She spoke of 21 distinct fatal failures for the 50k Alexa sites.  Shulman presented an interesting list of „fatal failures with TCP" that affected a lot of popular name servers:

- **After TCP handshake, DNS request is responded with RST+ ICMP(type=3, code=10) server cannot answer (administratively prohibited) for instance: edns-chtn.cht.com.tw 202.39.168.132**
- **After TCP handshake, DNS request is responded with ACK then RST for instance: gerek.accv.es 195.77.23.35**
- **Server keeps resending SYN+ACK for instance: ns7.utoronto.ca 162.243.71.42**
- **After TCP handshake, DNS request is responded with RST for instance: dns1.hessen.de 141.90.2.53**
- **TCP window fluctuations: SYN+ACK with window 0, then SYN+ACK with window > 0 (e.g., 4096) for instance: beloit.edu 144.89.40.1**
- **After TCP handshake, DNS request is responded with multiple small segments e.g., segments of size < 100bytes for response length 557 bytes for instance: ns.CWRU.Edu 129.22.4.1**
- **After TCP handshake, server sends SYN+ACK, then silent for instance: cnsa.vita.virginia.gov 166.67.65.169**

Concluding from this she said, while encryption was „no doubt important" more attention of the

# WGs and BoFs

## DPRIVE

Despite the rather sceptic views expressed by Haya Shulman on the efficiency of DNS encryption efforts (see „DNS 'pretty bad privacy'“), DPRIVE is moving ahead. Publication of the DNS privacy document is imminent (it has been in the RFC Editor queue for several weeks already).

WG Chairs and authors of the core document DNS over TLS (Allison Mankin et al.) are eyeing last call later this year. Features to be added to the draft are TLS 1.2 or better (no installed base), a reference will be made to RFC 7525 (BCP 195 Recommendations for Secure Use of TLS and DTLS). One major, still undecided discussion revolves around the question if a new port should be used instead of STARTTLS to set up the secure connection. While simpler for implementation, several members of the WG called for better data and threat analysis. The WG did not come to a final answer on this question. The evaluation draft, while being said to be important, did not receive as much attention.

There are several issues that DNS over TLS shares with the DNS over DTLS, which hopes to avoid the classical issues of TLS (head-of-line blocking) while still offer a privacy enhancing UDP. Advantages were also, said Dan Wing, who presented the draft, broad use in protocols like WebRTC and speed in session-resumption (with DTLS 1.3, sessions could be started with zero round trips, a figure that has been used to promote QUIC, the alternative transport protocol promoted by Google, see below). Except issues with anycast that were special for DNS over DTLS most problems were shared with DNS over TLS. The issues include blocking of encrypted traffic on port 53, authentication of DPRIVE server and downgrade attacks. Downgrade attacks were also a concern described by Haya Shulman (see above).

DNS over DTLS has already been adopted as a WG draft. Not that much interest was for another solution: IPsec. Instead of potentially complicated changes to DNS IPsec is presented as a way to encrypt all traffic (VPN). The draft to check out was currently developed in the IPsecME WG.

## DNSOP

Apart from the controversial debate of how to tackle the special names reservation policy (see above) DNSOP discussed intensively about aspects of TCP use for DNS, which is driven by the protection it provides against spoofing and amplification attacks, as well as by the overall development to bigger DNS response sizes (DNSSEC and IPv6). Presented by Sara Dickinson, one of the co-authors of the DNS over TLS draft,  the document is intended to make TCP a „required“ part of the full DNS protocol implementation.

Mandating behaviour, especially for a long-standing protocol like DNS, is nevertheless difficult, the WG acknowledged. Risk of idle TCP connections should also be addressed, warned Geoff Huston.
The revived document on EDNS-TCP-keepalive, also presented by Dickinson, is supposed to signal variable idle time-out for TCP connections and shall help to balance UDP and TCP shares of the traffic. Servers could signal what they expect, clients could opt for keeping a connection up longer. Potential beneficiaries could be DNSSEC validating resolvers and Tor.

Finally there was a discussion on approaches for secure trust anchor retrieval for the time after the root zone rolled the KSK. Joe Abley said the retrieval and also a standard automatic bootstrapping of validators were still desiderata with regard to the upcoming roll. The design team for the KSK roll-over will soon, Abley said, publish their proposal to be discussed by the community. He proposed to revive to older documents on trust anchor retrieval and validator bootstrapping. Questions were raised why these would be needed given that there is RFC 5011, which should be promoted instead. Less well received was also a quick proposal by Warren Kumari, Google, who had asked for a potential signaling mechanism of who would break when the key is rolled.

## Dane - baking DANE into the basic infrastructure

The DANE WG is continuing to make DANE more widely usable for validation and authorization of various client applications. In Prague two proposals were presented: one describes a "TLS extension for transport of a DNS record set that is serialized with DNSSEC signatures needed to authenticate that record set". Serializing the RR allows the TLS client to perform DANE authentication of a TLS server certificate without performing any additional DNS record lookups – resulting in less latency, avoidance of middlebox interference and the option to allow for link-layer authentication, too. Users will be web browsers, VoIP and XMPP services or other TLS clients that do not want or are unable to look up DANE records.

The second proposal wants to provide for TLS applications that already use DANE authentication of servers a similar mechanism to authenticate clients. According to the authors Victor Dukovni, Two Sigma, and Shumon Huque, VeriSign, it is an update to RFC 6689. Some of the design patterns could be beneficial for Internet of things design, where large networks of physical objects identified by DNS names can authenticate themselves to a centralized device management.

DANE WG co-chairs proposed to allow for a joint zone/label for OpenPGP (which is in Area Director evaluation) and the DANE S/MIME draft, which is still under consideration.
Two documents of the DANE group are currently in the RFC editor queue: dane-smtp-with-dane and dane-srv.

## QUIC - butt-kicking TCP?

Since a blog post on the Google website in April that spoke about the company being interested in standardizing QUIC, there has been a buzz around their home-grown transport protocol. At a very well attended Bar BOF (250-300 people) during the Prague meeting Ted Hardie, Google, clarified that Google was for the time being not going for a Working Group. Instead, as one of the authors of the protocol said to this author, the company wants to keep Quic under its own control (a WG would result in change control being transferred to the IETF) and do more experiments.

According to Jana Iyengar, who is one of the developers, QUIC looks like "TCP plus TLS plus SPDY over UDP". The main goal of the new transport protocol: speed. The protocol allows for zero round trip times and also features new things like an id number instead of the IP address for allowing a connection to stay up when moving through different networks. The good thing, Iyengar said, was that TCP was there as a fall-back allowing QUIC developers to experiment. That was also the reason why Google did not want to standardize, but instead keep QUIC separate – allowing to kick old TCP's butt. Google, while not asking for standardization asked for implementation by the BoF attendees. Christian Huitema, Microsoft, gave a presentation about implementation.

The light-weight characteristics of QUIC nevertheless have some downsides. While promoting always on encryption, the ID that stays live throughout the networks a user roams through, allows for perfect tracking, warned Daniel Kahn Gillmor. Tanja Lange, crypto expert from the University of Eindhoven, warned against replay attacks resulting from the 0-RTT, that needed to be fixed. Forward error correction is also still experimental, one of the IETF area directors said.

# EPPExt

With many EPP extensions of round one in WG last call or close the EPP Extensions Working Group is gearing up for another round of new extensions. A short Charter discussion in Prague confirmed that the WG will go on to be a stand-by group to discuss new extensions that will be registered in the IANA registry of EPP extensions. Extensions may be registered for informational purposes as long as there is a published specification that has been reviewed by a designated expert. The WG will review those documents that want to become an Internet Standard. Mergers of similar EPP functions may be attempted.

The Chairs presented a list of potential new candidates for EPP extensions. Rik Ribbers (SIDN) presented an extension for secure DNSSEC key relay as one of the last proposals from round one. It was agreed that it was ready for WG last call. A "generic relay" for securely transferring information during a provider change is out of scope.

Ning Kong (CNNIC) presented two documents on reseller handling and reseller mapping as the first of the new document batch.

# IETF News

## Postel Award

Rob Blokzijl, founding member of the RIPE and RIPE Chair for 25 years, received the Jonathan Postel Award 2015. Blokzijl, a physicist by profession, got involved when his boss at the Dutch National Institute for Nuclear and High Energy Physics (NIKHEF) requested that he had to „fix" for him to be able to internet-work with the CERN in the 1980s.  Blokzijl also was instrumental in the creation of RIPE NCC, in establishing the Amsterdam Internet Exchange (AMS-IX), and he served as a member of the first ICANN Board for the Address Supporting Organisation in 1999. In his acceptance speech he acknowledged the change from what he had seen as a pure technical function to an area that also interfaced with politics. Blokzijl is the 17th winner of the Postel award.

## IANA IPR

The IETF cancelled a meeting of the IANAplan WG, the WG in charge of discussing the transition of IANA towards a fully privatized, multi-stakeholder oversight. It was felt that there was nothing new to report, IETF chair Jari Arkko said. The somewhat controversial issue about IANA IPR and IANA domain name was only touched briefly during the administrative plenary on Thursday. The new IETF Trust Chair, Benson Schliesser, Brocade (stepping in for Tobias Gondrom, who was elected IAOC Chair), just re-confirmed the Trust „would be willing to hold intellectual property rights relating to the IANA function, including the IANA trademark and the IANA.ORG domain name".
The CWG transition proposal that stated that **„ICANN will grant PTI an exclusive, royalty-free, fully-paid, worldwide license to use the IANA trademark and all related trademarks in connection with PTI's activities under the ICANN-PTI Contract**" was not discussed at the Prague meeting. Arkko said to this reporter, there was an incompatibility related mainly to the CRISP (RIR) proposal.

Next meeting will take place in Yokohama, 1 - 6 November 2015