Council of European National
Top Level Domain Registries

Report on

IETF94 Yokohama

2-6 November 2015

Table of Contents

IETF94 from 10,000 feet (Executive Summary)

IETF94 in Yokohama was marked by a lot of talk about measurements and the potential benefits from measurements for the development of protocols. The ease of access to networking data called for by some participants would also create privacy problems, some warned during an opinionated discussion in the plenary. A measurement workshop of the IRTF had preceded the IETF week, adjunct to the Siggcom meeting.

Linking IETF/IRTF work to other bodies seems to be becoming a trend. During the meeting representatives of the mobile sector (GSMA) and IAB/IETF met to decide how to follow-up on a joint workshop in September, that has looked into the rise of encrypted traffic in the networks. GSMA also is considering giving data access to those interested in measurement (see especially HOPS Research Group).

A reflection of IETF work and potential "transitions" from legacy to new protocols (technical and organisational) is reflected in a draft written by Dave Thaler. Changes with regard to the format of meetings are also continuing, after the introduction of the Code Sprint and Hackathon (in addition to "Bits and Bytes") the IETF cuts back on the plenary time, combining administrative and technical plenaries.

In 2016 the IETF will be, for the first time, meeting in Latin America, willing to spend more than usual for reasons of internationalization and diversity. Another expression of change for the once techie-only organisation is the imminent establishment of a Human Rights Research Group under the preview of the IRTF.

With regard to the Domain Name System, the review of 6761 is now underway, driven by a design team. The IESG wanted the community to "fix" 6761, IETF Chair Jari Arkko said during the plenary. In essence there is a concern by parts of the IETF that the organisation might be (ab-)used as an alternative provider of name space (see special names). DNS over TLS is getting close to Working Group Last Call (see DNSPrive WG report), meanwhile there is yet another push to add trust by logging the DNSSEC signed root.

Highlights

IETF and special name reservation: As big a deal as it gets

The Special Names RFC 6761 will be reviewed by a newly established Design Team. Members were picked by the WG Chairs and finally confirmed during the Yokohama meeting: Ralf Droms (only DNS-outsider), Alan Durand (ICANN), Peter Koch (Denic), Joe Abley (Dyn) and Ralf Droms. Apart from Droms who was named according to WG Co-Chair Suzanne Woolf, these are the authors of a presentation on a possible review of the Special Names Registry of the IETF.

Design team selection

Lack of clear communication about the design team selection process was expressed during the WG meeting. The WG Chairs had failed to follow-up with an undisclosed number of people which had asked to be on the Design Team. Co-Chair Suzanne Woolf apologized during the Yokohama meeting. The selection of people chosen (two ICANN, or ICANN close, one a former WG Chair who had expressed reservations against the special names process) lends itself to a more conservative outcome.

Design Work processes

During the WG in Yokohama it was somewhat unclear how the WG would be involved in the next steps. According to Droms' explanation, the Design group would do the review work, report back to the group and push discussion on the mailing list, rather than use f2f time during IETF meetings. IAB Chair Andrew Sullivan on the other hand called the RFC 6761 review "as big a deal as it gets" for the Working Group and therefore "deserves time and attention of the folks who wish to attend". It is not fully clear how the WG will proceed now, but given the considerable interest and the comments in Yokohama – also by Jari Arkko who said the IESG had handed the issue back to the community for a solution – time will be allocated for discussion. Other drafts (including .bit (bitcoin), gns (Gnu Name System), etc.) have been put on hold for now.

The substance of the matter: architectural, technical and organisational aspects

Koch and Durand said in Yokohama during their presentation of a problem statement draft that they intended to look at 6761 from an architectural, technical and organisational aspect.

Architecture-wise there is an acknowledgement that the name space is not just the domain name system, but that there are other, different types of names, using more or less different protocols. Localhost allows for internal name resolution (without reference to DNS), Local is used by Multicast DNS, Onion is used "to construct names that designate anonymous hidden services reachable via the Tor network using onion routing." The TLDs were not used as TLDs in these cases, but as "protocol switches" (DNS to something else) and considerations could be given to catalogue them or use an .alt as an indicator for the switch for alternative resolution choices.

Technically the question was how users would get to know how to treat the special names to avoid potential problems, for example the leaking of .onion or other non-DNS names. In the case of .onion, privacy/security issues for the onion users could be a result.

Conflict between RFC 6761 and the IETF-ICANN MoU (RFC 2860)?

From the organisational point of view, RFC 6761 had somewhat shifted the relation of the IETF and ICANN. According to RFC 2860 the IETF seized its role as a policy body for the DNS and ICANN and the latter now had a process to delegate names, albeit a heavy and costly one. A second open round for new TLD applications was just under consideration, Durand said.

While the MoU had made an – albeit vague – exemption from the ICANN's prerogative as the name space body in charge, RFC 6761 had slightly changed that. For the review now the intra-organisational problems (IETF lacking a process to reserve names due to 6761 vagueness) and the inter-organisational problems (was 6761 as such a violation of the MoU?) had to be considered.

Questions around a potential invalidity of 6761 were quickly rejected as a potential "rat hole". An opening up of 2860 – in order to renegotiate the roles of ICANN and IETF – were quickly dismissed by Arkko ("IESG asked WG to look at 6761. RFC 2860 is not on the table.")

Options for resolutions

Options presented in Yokohama included the revocation of the 6761 process, an option not included in the written draft, as Stéphane Bortzmeyer (AFNIC) noted. Bortzmeyer called on the design group to rather go back and fill the gaps of the 6761 process in order to make the reservation of special names on technical grounds operational. Another option could be the creation of a special TLD-Zone (like alt., see proposal by Warren Kumari, Google). But that would still make a process necessary on how to decide about which reservations to grant.

Control instead of Trust: DNSSEC logging

The Public Notary Transparency (TRANS) WG is preparing to send its core specification on Certificate Transparency (RFC 6962 bis) to Working Group last call (WGLC) by December. For monitoring the logs, the WG has adopted a document proposing three mechanisms to detect misbehaviour (see Gossip draft), check of Signed Certificate Timestamps (SCT), the pooling of Signed Hash Trees (SHTs) at https servers, and finally direct sharing of SCTs and STHs with trusted auditors by http clients.

Certificate Transparency (CA) shall allow third parties to check public logs of certificates edited by a CA in order to detect false certificates. With the core spec close to be done the WG now is looking into attack scenarios, but also into the extension of the transparency by logging mechanism beyond the CA system. In particular a so called Bar-Bof, an informal gathering, considered beginning work on logging DNSSEC (DS) signatures.

Extend Certificate Transparency to DNSSEC

Logging as a way to allow monitoring signatures has been described as a general mechanism from the start. During the IETF Yokohama a group of a dozen people gathered to discuss the potential for using public logs of DNSSEC signatures to allow checking on the potential misuse of keys by a zone owner. According to a first draft the attack scenario looks like this:

"a zone owner that has been compromised or compelled by a third party can hijack a child zone to return different DNS data that is indistinguishable from DNSSEC validated data from the child zone by using its own DNSKEY to sign DNS data on behalf of the child zone. It could deliver this modified DNS data to only selected regions or individuals, making this attack very difficult to detect by the legitimate child zone."

The first draft by Daniel Kahn Gillmor and others was very briefly advertised in the TRANS WG in Yokohama. Discussion was moved to a Bar Bof meeting where participants agreed that a first step to test a mechanism should be the central root zone to avoid complications of a larger zones like .com where possibly cooperation with a private provider like VeriSign could be necessary.

Kahn-Gillmor, one of the document co-authors, said that the mechanism would allow to detect potential split versions of the signed root zone. The transparency mechanisms in general allow only for detection, not prevention. According to the draft proposal zone owners must submit DS resource record to one or more preferred logs before publishing them (the log entries must themselves to be signed for authentication). They would be accompanied by additional resource records (DNSKEY RRs, DS RRs, and RRSIG Rrs). An accepted DS RR would result in inclusion in a monitoring log.

Wes Hardaker (Sparta) agreed that it was possible to log everything the root signs, while he was not sure it would scale for large zones like com. For the root zone – due to strong monitoring in place for the root zone – he did not expect that anything would be detected. But logging the root could be a good proof of concept for DNSSEC logging.

Co-Signing

Public logging (for PKI certificates or DNSSEC signatures) is not enough, proposed Bryan Ford from the Swiss Institute of Technology in Lausanne, as a powerful attacker could get access to the private key of CA and log server which were nothing less than single points of attacks.

Ford presented the idea of witnesses for the logging. Changes not co-signed by all witnesses, who could be trusted third party organisation like the EFF or the CCC, would be detected and exposed by the latter. Again, manipulation would still be possible, but would be detected. Co-Signing can in principle be compared to co-signed public PGP-keys. An implementation of the co-signing concept had been tested with Google and revealed that computation time was reasonable, Ford said.

Certainly problems remaining were, for example the sloppy handling of certificate revocation. In 74 full IPv4 HTTPS scans the researchers found that 8% of the certificates served had been revoked, and that obtaining certificate revocation information can often be expensive in terms of latency and bandwidth for clients. Checks on 30 different combinations of web browsers and operating systems revealed that browsers often do not bother to check whether certificates are revoked (including mobile browsers, which uniformly never check). In Google Chrome the CRLSet only covers 0.35% of all revocations.

Motivations for certificate providers to adopt the more arduous co-signatory concepts according to Ford could be mere competition, the possibility to deter authorities interested in keys (as without co-signing the split would be noticed) and the market power of browser companies (not served in Chrome when not co-signed).

Workshops and BoFs

EPPEXT WG morphing into REGEXT WG

The Extensible Provisioning Protocol Extensions (EPPEXT) WG is in the middle of re-chartering and setting new milestones, including potentially four or five batches of new extensions (see list). The WG will continue to discuss EPP extensions (brought by registry operators) intended to become proposed standard documents. Proprietary extensions or individual submissions intending to receive an informational or experimental status are free to apply for registration in the EPPEXT registry after expert review according to RFC7451.

Group 1 WGLC completed by February 2016
-allocation token (draft-gould-allocation-token)
-change poll (draft-gould-change-poll)
-epp-rdap mapping (draft-gould-epp-rdap-status-mapping)
-reseller (draft-zhou-eppext-reseller and draft-zhou-eppext-reseller-mapping)

Group 2 WGLC completed by May 2016
-verification code (draft-gould-eppext-verificationcode)
-service messages (draft-mayrhofer-eppext-servicemessage)
-nv mapping (draft-xie-eppext-nv-mapping)

Group 3 WGLC completed by September 2016
-fees (draft-brown-epp-fees-05)
-bundling (draft-kong-eppext-bundling-registration)

Group 4 WGLC completed by January 2017
-IDN Table Mapping (draft-ietf-eppext-idnmap and draft-gould-idn-table and draft-wilcox-cira-idn-eppext)
-Relay (no draft yet, split from keyrelay)

With the new charter, the WG will also take up the work on the Registration Data Access Protocol RDAP (originally developed in the Weirds WG). The more general nature of the WG will be expressed in its new name Registry Extension WG (REGEXT).

Several drafts, while not official WG drafts, were floated on the EPPEXT mailing list, namely the first version of the ICANN gTLD profile and a possible document to map statuses of EPP and those registered for use in RDAP (including potential gaps). Another potential individual submission (potentially through Barry Leiba AD as shepherd) is a draft on data escrow, which has been pending for some time.

All-things-ICANN WG?

The short Charter and milestone discussion revealed that there are some concerns on the scope of the new REGEXT WG. Peter Koch (Denic) questioned the idea that EPPEXT/REGEXT might become the new “all-things-ICANN” WG. The IETF should instead focus on protocol development.

The concern obviously is that ICANN could be granted a special one-stop-shop at the IETF and in turn bring IETF-stamped “standards” back to the Registrar-Registry community from there.

With regard to ICANN’s gTLD profile, Kaveh Ranjbar (RIPE) prior to the meeting had begged for caution (on the mailing list). The IETF should not be taking up a document which in essence was outlining ICANN’s “policy” towards its Registries and Registrars. Scott Hollenbeck, VeriSign, argued that implementation requirements (based on community consensus) could well be informational RFCs. A similar document is a draft tabled by ICANN on the Trademark Clearinghouse specifications.

Hollenbeck and Ning Kong (who stepped in for the Jim Galvin), Afilias, and Antoine Verschueren (former SIDN) to chair the Yokohama meeting held several hums to ask if the EPPEXT/REGEXT WG would want to continue to also discuss the informational/proprietary registrations (or leave them to the expert review process) and if the working group agreed to divide upcoming work into groups. There is some disagreement over the latter, not least because some thought flexibility to allow to take up RDAP (or other?) work should be maintained.

No agreement could be reached either over taking up an extension for a special Chinese registration policy with questions being raised over how the Chinese regulator would react if the WG would change the procedures. A WG is in charge of change control to a draft once it "owns" a document which is intended as standards track RFC.

Koch questioned the motives for creating ever longer lists of IETF standard track documents for the EPP extensions due to what he called potential "external incentives to have standard track labels on these documents that are not in line with IETF documents". Incentives could be financial (bonus for RFCs by companies), or a possible interest to add legitimization to technical/organisational solutions.

The next extension to be adopted is expected to be the Keyrelay document which foresees a mechanism to exchange DNSSEC keys in case of registrar change.

DPRIVE: Countdown for DNS over TLS – And what about implementation?

The major protocol project of the DPRIVE WG, DNS over TLS, is in Working Group last call and should become an RFC early next year. The WG basically agreed to have profiles cut out of the document (they will be dealt with in an extra document) and StartTLS removed. For the time being there will be no option to upgrade a DNS over TCP session to a DNS over TLS session. It was agreed that this change would make the protocol more straightforward and "light-weight".

DNS over TCP enabled servers will listen to a new port has been reserved with IANA (Port 853) for negotiation of a TLS connection. Ideally TLS and DTLS (which has been given up as a separate document, see below) will be able to refer to the same document.

Things to be explored would be the use of additional authentication mechanisms beside TLS, for example DANE. Client authentication (in addition to server authentication) could be developed further.

Questions remain for some potential implementers. Wolfgang Beck from Deutsche Telekom asked about the number of parallel sessions that could be run, as for DTAG they did run 1 million in parallel per server. The document was giving some advice on enhancing performance, according to Duane Wessels (VeriSign Labs, and one of the authors).

Christian Huitema from Microsoft pointed out that deployment for large scale providers which would be arduous and there was a potential security issues in merging the security for a large number of credentials.

With regard to implementations, Sarah Dickinson reported that Unbound was done with DNS-over-TLS (since 1.4.12). Tools available were digit, getdns. The DNS Privacy team won a prize during the IETF Hackathon preceding the IETF94. Work from the Hackathon to be consulted includes:

DNS Privacy topics
– getdnsapi extension (call debugging) implemented with changes so user learns transport/privacy results
– edns0-client-subnet privacy election
– edns0-padding option (client side is done)
– Check TLS at Recursive - node.js application

DNSSEC topics
– DNSSEC roadblock avoidance – proposed new extension for getdnsapi
– CDS/CDNSKEY

For more see a list of Hackathon projects related to DNS here.

DNS over DTLS given up for now, fragmentation to be explored

The further development of DNS over DTLS was basically given up in Yokohama after some discussion on whether it was worth to pursue it while allowing for fragmentation and reassembly of packets too big. Dan Wing presented the DNS over DTLS version, which was considered a possible variant to privacy friendly DNS over UDP. Yet Wing acknowledged that in order prevent frequent fall-back to DNS over TLS ("so why bother to have DNS over DTLS at all?"), fragmentation and reassembly had to be performed, something seen as to complicated by many (including Andrew Sullivan, IAB Chair, and Paul Hofmann, one of the latest additions to ICANN's techie department). Wing agreed to give up on further developing the draft which also got a new IANA port (port 853) to avoid middle box issues.

DPRIVE new work?

There was not a lot of discussion on the document by Allison Mankin on Evaluation of Privacy for DNS Private Exchange (which is a WG document) and on a completely new document on "stateless DNS encryption", mainly as it is a different approach from the one pursued currently with DNS over TLS.

With one major issue, one privacy mechanism for DNS querying soon ready, the WG also briefly discussed what to tackle next. Ideas were to look into securing queries also between recursive and authoritative server, as more and more people might use their own resolvers (Bortzmeyer, Gillmor), EDNS padding (Gillmor), where there exists a draft already, and also measurements to follow-up the eventual update of DNS over TLS (Tim Wicinski).

Paul Hofmann (ICANN) recommended a more conservative wait-and-see development approach. Allison Mankin (VeriSign Labs) instead underlined the group should use the current window of interest for DNS privacy, as waiting might result in losing the current momentum.

DNSOP – a list of new proposals

Beside the big discussion over the special names reservation process (RFC 6761, see Highlights), the DNSOP group briefly discussed a list of proposals and accepted work on several issues.

1. An older proposal by Joe Abley, providing a mechanism for an authority server to signal that conventional ANY queries will not be supported for a particular QNAME. ANY requests can be used for legitimate purposes like debugging or checking the state of a DNS server. It can also be used to retrieve MX, A and AAAA RRSets for a mail domain in one single query. The problem is that ANY-queries more and more are abused to mine full data sets, or worse, to use them as an amplification factor in DDOS attacks. Returning to smaller answers allows to mitigate such attacks.

2. A short document by Paul Wouters and Olafur Gudmundson describes a mechanism for in-band-signalling of DNSSEC status changes, allowing a child to signal to the parent to turn DNSSEC on or off for its domain using CDS/CDNSKEY. The status changes may be necessary when a domain owner changes his registrar. So far, not being able to enable trust via an easily automated mechanism was hindering DNSSEC at scale by anyone that does not have automated access to its parent's "registry". While there were some concerns to allow for DNSSEC to be turned off via this, overall there was support for this to be developed, while bootstrapping issues has to be checked.

3. Duane Wessels proposed an EDNS option (OPT meta-RR [RFC6891]) that will allow end-system resolvers to tell a server in a DNS query which DNSSEC keys they use to validate the expected response. It allows to measure acceptance and use of new trust anchors and key signing keys (after a key or algorithm roll over). Questions raised were how cached information was handled, with some concern that this might be a source for an attack. Possibly this could be limited to the root zone only. A potential alternative as fare as compliance with 5011 (roll-over procedure) was concerned, could be a draft by Warren Kumari which was not discussed.

Also presented, but not adopted yet:

4. A document presented by Shane Kerr (for a group) that proposes to allow for fragmentation of DNS messages instead of on the IP layer. The objective, according to the draft, "is to allow authoritative servers to successfully reply to DNS queries via UDP using multiple smaller datagrams, where larger datagrams may not pass through the network successfully."

5. A proposal to revive one part of an older (2010) proposal by Paul Vixie on "Stopping Downward Cache Search on NXDOMAIN" to make it clearly obvious that if bar.domain does not exist foo-bar.domain does not exist either. There was no consensus on going ahead on this.

TCPInc: IETF unable to decide

The TCPInc WG is expected to come up with additional opportunistic security for TCP connections, but has been unable so far to make a choice between two different proposals on the table.

The two candidates are TCPCrypt, developed by a group of researchers at Stanford University (Andrea Bittkau, Mike Hamburg and others, see here), and TCP TLS Option (Erik Rescorla, Mozilla). The main difference according to Sean Turner (TLS Chair) between the two originally has been that TCPCrypt was in the kernel, while TCP TLS Option was claiming that TLS was well deployed and known to users. According to Turner Rescorla in the meantime has also worked to have it in the kernel, too. The basic goal, opportunistic security for TCP, and the basic concept to establish the connection are quite similar[1].

[1] Connection establishment for TCPCRYPT:
"The initial key exchange works as follows. Each machine C has an ephemeral public key, K C . When C connects to a server S for the first

TCPCrypt has been further developed to include a TCP Encrypted Negotiation Option (ENO). TCP-ENO is a TCP option used during connection establishment to negotiate how to encrypt traffic, according to the draft proposal, it can be implemented incrementally and allows for fall-back to unencrypted TCP and middlebox traversal. TCP TLS Option is also using ENO the TLS 1.3. TLS 1.3 is expected according to TLS WG Chair Sean Turner to be ready by the end of the year and getting being vetted by academics before final standardization during a conference planned for February (TRON workshop organized by ISOC).

The TCPCRYPT authors are pointing to some implementation underway and called on the WG to make a decision as they were unable to sustain further work without a commitment because financial resources for the project were running out.

While many experts, outgoing Transport Area Director Martin Stiemerling and IRTF Chair Lars Eggert urged the WG to come up with a choice for either one of the solutions, there are some supporting to go on to have parallel implementations. The WG Chairs Mirja Kühlewind and David Black after offering to the WG to either converge both proposals into one, select one or go ahead with both, favour to allow both going forward.

While the IETF cannot make up its mind about a path for a secure TCP option, Google is advancing with an individual submission of Quic as a UDP-Based transport protocol that also supports opportunistic security. The Quic draft has been assigned to the Transport Area Ads.

CFRG: Is post-quantum crypto next?

The Crypto Research WG is close to finish its selection of new algorithms for TLS. Answering the request of the TLS WG the group has selected two new elliptic curve algorithms, namely Ed25519 and Ed448. Discussion on hashes to use is still ongoing for Ed448, because there is a concern that SHA3-512 could not be performant enough. But Co-Chair Alexey Melnikov said that the hash selection will be done by the end of the year, completing this round of algorithm choice.

The request for the new curves which was put to the CFRG by the IETF TLS WG has to a large extent been an IETF reaction to revelations about Bullrun and related NSA programs that seek to weaken encryption. The National Institute of Science and Technology (NIST) which was used by the IETF as the resource for new crypto algorithm material has lost its position as a (practically sole) crypto provider for the IETF for now.

If the CFRG which after years had focused more on research will now stay in business to provide crypto for the IETF protocols, remains open. Yet there was a call to use no crypto that had not undergone public review.

There were only some answers to the Co-Chair's question what work the group would address next. Ideas discussed very briefly were the reduction of unencrypted meta-data (Bryan Ford), a potential standardization of padding (Stephen Farrell, Trinity College) and post-quantum crypto (Ford, Swiss Institute of Technology and Daniel Kahn-Gillmor, ACLU).

NSA recommendation: don't bother to go to elliptic curve crypto, yet

The recommendation of the NSA to wait moving to elliptic curve crypto for now, if one had not done so, and instead wait for a post-quantum crypto suite (which would soon be available) was greeted with suspicion at least by parts of the community according to Marcos Sanz (Denic). The statement was rather fuzzy, said former IETF/IAB Chair Russ Houseley (who was supported by the NSA during his tenure). It well acknowledged a potential need to mitigate attacks, but asked to go to large keys for that instead to switch to elliptic curve crypto. Yet given the efficiencies of EC there were many applications that needed to move.

time, C chooses a random nonce, N C ; S chooses a random secret, N S ; the two exchange the following messages, also shown in Figure 1:
HELLO
PKCONF , pub-cipher-list, [cookie]
INIT 1 , sym-cipher-list, N C , K C , [cookie]
INIT 2 ,
sym-cipher, E NCRYPT (K C , N S)

Connection establishment for TCP TLS Option:
SYN + TCP-TLS
SYN/ACK + TCP-TLS
ACK TLS Handshake Application Data (over TLS)

HOPS and the measurement drive at the IETF

Originally the BoFs on “How broken is the Protocol Stack” (HOPS) intended to work on a new transport protocol that would allow for better traversal of middle box-fenced areas on the net. Instead of inviting deep packet inspection the protocol should give away some information about packets. Yet the idea for a “Spud”-protocol was not well received for its meta-data adding additional layer. The HOPS organizers now advance the adoption of a HOPS Research Group under the roof of the IRTF.

The new HOPS Research Group (name is expected to be changed) intends to bring academics and operators together to allow to use measurement studies and take results back into facts-based-protocol development.

The Group gathered in Yokohama and chaired by Brian Trammell and Mirja Kühlewind (both ETH Zürich) talked at length about confidentiality assurances it could give to network operators so that they would open their data gates to the researchers. Chatham House Rules or even closed RG meetings were addressed, with concerns being raised by the IAB Chair Andrew Sullivan (how would undisclosed data help to persuade people to work on problems) and the IRTF Chair Lars Eggers (a researcher might need to agree to an NDA on data by network operators, but could anyway present anonymised findings to the group).

A representative from the GSMA, Natasha Rooney, reported that GSMA was working to formalize a process for collecting data and sharing it with the group. That was in the interest of the operators who would benefit from the expertise and protocol development in the IETF consecutively. GSMA and IAB met during the IETF week (a follow-up to the earlier Marnew Workshop, presented also in the SAAG meeting) to establish a closer collaboration.

During the HOPS meeting there were presentations on measurement projects like the ARK Measurement Infrastructure of CAIDA, a planned path impairment observatory and the use of crowdsourcing as resource for measuring and informing protocol design. A group at the University of Madrid had used the crowdsourcing platform Microworkers to lure users (for small amounts of money) to perform the connection to the server running the experiments on the status of encryption of the respective users.

Measurements and measurement-driven protocol development (or measurement-informed protocol development) was also the topic of the Technical Plenary section of the now combined IETF/IAB/IESG Plenary on Wednesday. Brian Trammell used the question “can the Internet be run over UDP?” as an example to consider potential existing gaps despite the many measurement platforms and tools already available. The second presenter, Alberto Dainotti (CAIDA), called for better and more data on the BGP environment and announced a BGP live measurement hackathon in February (contact bgp-hackathon-info@caida.org).

There was during the plenary considerable backlash there against the call for bigger data collections to feed the measurements. Concerns were raised in particular about potential privacy problems resulting from the attempt to measure more and more detailed data on the networks.

Standardization also for users, not only big vendors: ISS BoF

The only BoF in Yokohama was initiated by a group of Chinese academics from former IETF host Tsinghua University who clashed with a noisy group of IETF standardization “superpowers”, Microsoft/Yahoo, Cisco, Google.

While focusing somewhat on measured synchronization inefficiencies (of varying degree between different proprietary systems) the problem to address is the lack of interoperability for users and the need for developers to deal with various APIs. “With a standard sync protocol provided, a third-party client that supports multiple Internet storage services is easy to implement since APIs provided by different providers would be unnecessary or at least simplified”, the problem statement reads.

During the discussion about the value of IETF standardization in this area, reactions were pretty blunt: it was of no interest to the large monopolies in the field, namely Dropbox (400 million users, 4 percent of traffic). He would know no reason why he should implement it, said one Yahoo representative. Without implementation the work would be a waste of time, said Richard Barnes (Mozilla).

While the BoF ended on a low note, there were several hints that a standard might be interesting for individual user cloud solutions, enterprise solutions and/or open source cloud storage providers (https://syncthing.net/).

IETF Chair Jari Arkko after the BoF reacted with a call (on the mailing list) that while some existing large storage vendors might not be interested (“this space is big”) including private clouds, enterprise markets, Internet of Things, secure storage services, and more. Arkko added: “I would also observe that vendors are not the only ones specifying what solutions should be. We often have a situation where users/customers want standards that they want to place in their

RFPs so that they can create a more competitive environment for the services that they need."

Asked by this author if Baidu, which provides free storage service in China would be interested in implementing a standard, they said that with the current ban of Yahoo and Google services in China, Baidu was in a too comfortable situation. Would competition be allowed, Baidu might feel more compelled.

The ISS BoF notably was the only BoF in Yokohama which resulted in a question during open mic time during the plenary session. IETF leadership pointed to the load of new work ongoing, three new WGs are just about to be accepted by the IESG.

Human Rights Research Group

After two consecutive BoF sessions and another IRTF session in Yokohama the forming of an IRTF Research Group on Human Rights is now imminent (charter is here). The two co-chairs Nils ten Oever and Avri Doria presented a video clip that compiled extracts from the interviews the group conducted inside the IETF community which illustrates that the topic by now has become (close to) mainstream. The clip features IETF leadership including IETF Chair Jari Arkko, ISOC Board member Scot Bradner, and also many developers from the DNS community, including IAB Chair Andrew Sullivan and AFNIC Researcher Stéphane Bortzmeyer reflecting on the relation of human rights and standardization. Technical protocols, as Arkko put it there, always can take one path or another.

While there is still some concern in the IRTF/IETF group that the scope of the HR group lacks clear focus, one of the first drafts presented in Yokohama was an informational draft calling for protocol work to be performed to serve the end user on the net in the first place. The draft by HTTP2 WG Chair Mark Nottingham clearly could be seen as a reaction to the call by Edward Snowden to the developers to help users with better security and control of their communication.

The users' interests first-call nevertheless would be a change from the current perception the IETF has of its own work, as it serves to address problems of operators, vendors and providers in the first place. According to IAB Chair Andrew Sullivan, it was unclear who the users were. This was also a difficult question in the case of machine-to-machine communication.

Next IETF meeting: 3-8 April 2016 in Buenos Aires

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 52 full and 9 associate members – together, they are responsible for over 80% of all registered country code 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