Council of European National Top-Level Domain Registries

Report on IETF96

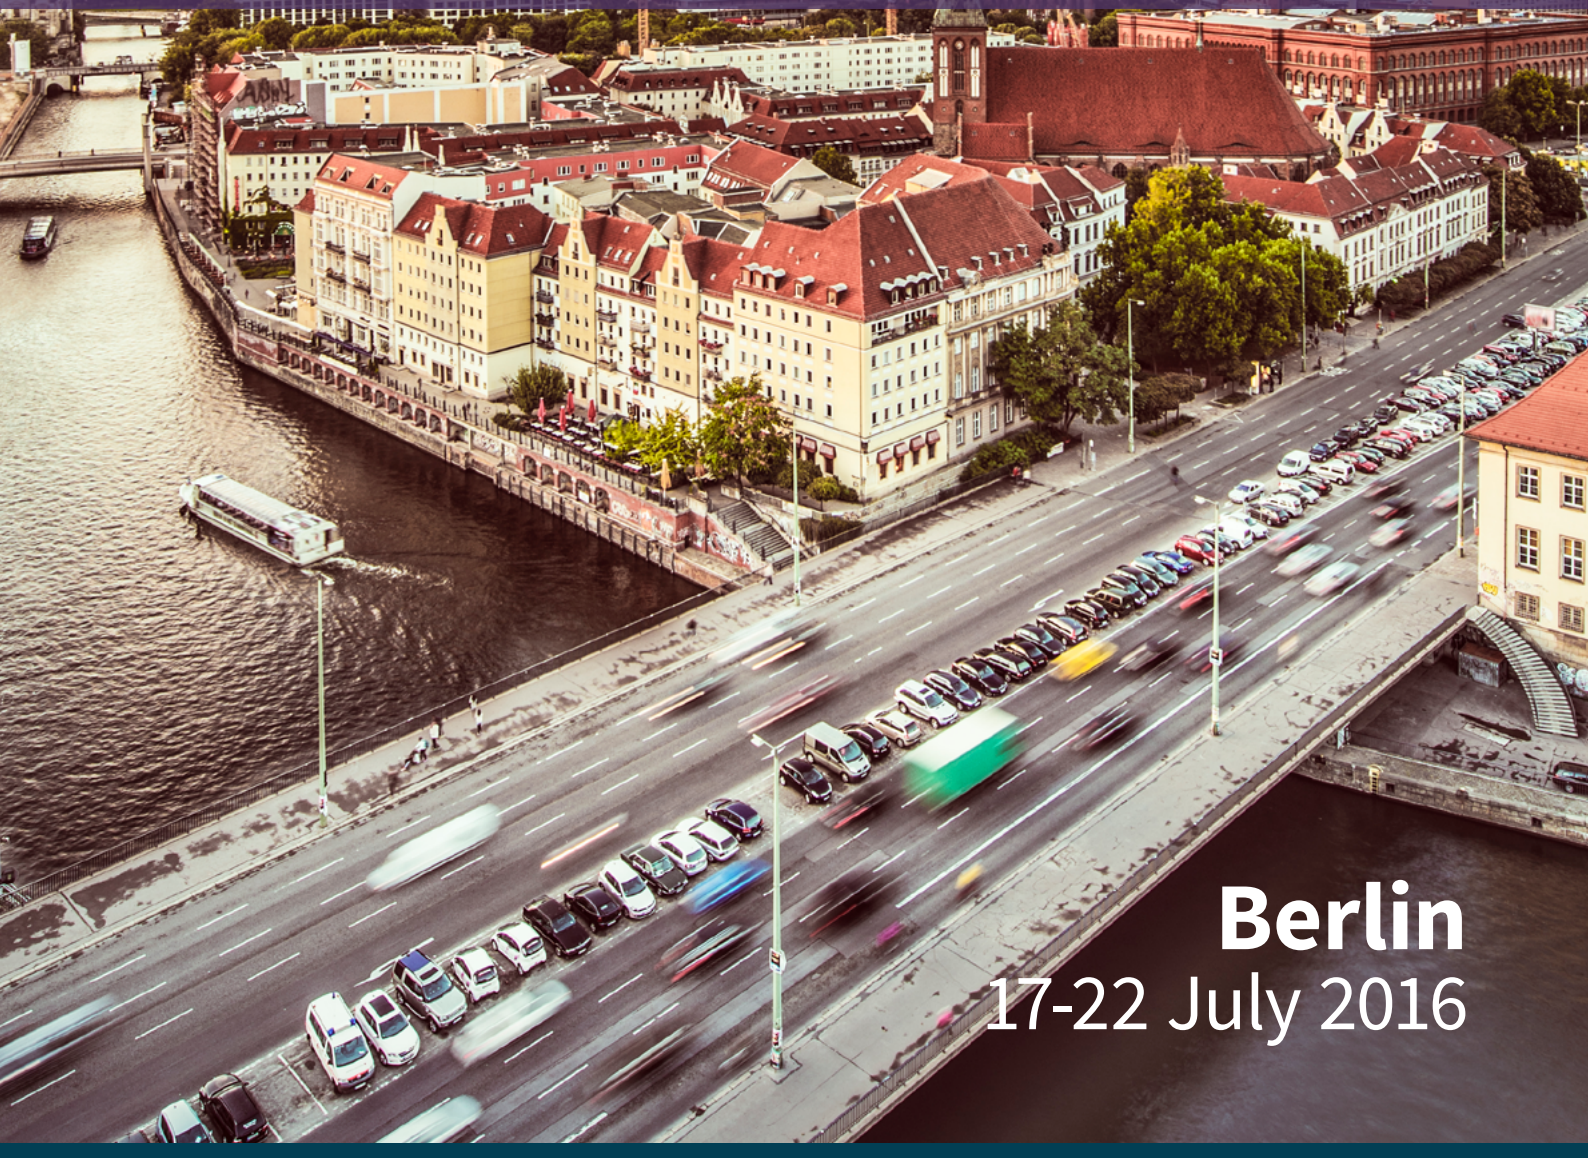

Berlin
17-22 July 2016

CENTR vzw/asbl · Belliardstraat 20 (6th floor) · 1040 Brussels, Belgium
Phone: +32 2 627 5550 · Fax: +32 2 627 5559 · secretariat@centr.org · www.centr.org

Contents

Highlights

RegExt: Using the IETF to create ICANN registration policies

The Registry Extensions WG (RegExt) faces a general problem with receiving a growing stream of ever more complicated documents on registry extensions while not getting much attention from the affected industry. ICANN alone has fed 5 (of 16 documents altogether) into the WG, with up to 60 pages per document. With only about 20 participants during the Berlin session and only very minor discussion on the RegExt mailing lists the WG risked to pass sub-standard quality RFCs, warned Alex Mayrhofer (nic.at).

The ICANN documents especially present the RegExt WG with the challenge that they include policy aspects. Jody Kolker from GoDaddy, the only ICANN registrar active in the RegExt session in Berlin, warned that in the Trade Mark Clearing House Extension the IETF was creating a policy for ICANN, "and it should be the other way round". Scott Hollenbeck (VeriSign) added that the IETF standards process must not be used as a means to rubber stamp policies for which there was no consensus policies yet. While pure informational documents could still be viewed as options, standards track documents in particular could be fed back as must-implement "technical standards". It is interesting to hear this argument from VeriSign, who is one of the main users of the process, as well as from other parties including country code operators.

Policy-making through the back door?

The example for such an IETF policy-making move was given by Kolker with regard to the "ICANN TMCH functional specifications" draft that describes requirements, architecture and interfaces between the ICANN Trademark Clearinghouse (TMCH) and the ICANN registries, as well as between the TMCH and ICANN registrars for sunrise and trademark claims phases. Kolker called the 48-hour window for acceptance datetime for trademark claims notices a problem. For registrars, collecting updated trademark claims notices just before the general registration opens up (after a trademark claims period/sunrise period) could create bottlenecks allowing for registrations for general phase domain registrants – where another registrant had failed during the TMC phase. 48 hours were an artificial timeline not present in any ICANN policy.

Similarly, the registry extensions related to transliterations and translations (of contact info) prepared by ICANN staff received some critical comments. The documents "Transformation of Contact Information Extension Mapping for the Extensible Provisioning Protocol (EPP)" and "RDAP Transformation of Contact Information" add information about language, type and source of translation/transliteration to the registry information.

Some of the information added was redundant, though: for example, for country information the ISO alpha-2 codes were used internationally. The added information seemed to only support for the information to be easy-to-search/google.

Even more data to be fed into the ever more data-rich registry databases included information about resellers. Mayrhofer warned that extensive sets of information about resellers would represent garbage-in-garbage-out-type data of no technical value. With the change from the old Whois to the new RDAP system which became obligatory on 30 July, the development to data-rich registries is prepared. There is some risk that more and more data points will be shifted from optional to obligatory – through ICANN policies or IETF RegExt RFC documents.

Problems of Domain Bundles

Representatives from CNNIC presented the draft Extensible Provisioning Protocol (EPP) Domain Name Mapping Extension for Bundling Registration that intends to allow to package, for example, simplified and traditional domain names together, allowing also for bundle updates.

The bundling as proposed could result in problems with DNSSEC, though, experts warned. Clients unaware of the bundles might open up vulnerabilities as well. There could also be a potential need to allow people to reject automatic bundling, because they could be disinclined to become related to certain types of names (for example the PRC-typical simplified domain name version).

The RegExt Chair Jim Galvin concluded that there was

clear demand to discuss bundling of domain names more generally and also understand how the various current bundling drafts were related (including IDN bundling).

Eliciting more reviews, more face-to-face discussions and on the list is one of the major tasks the WG has to take on. Galvin, who has been practically acting as a single chair for some time now (with Antoine Verschueren not participating in person in the meetings), appealed to the participants to send out documents for review to connections in the sector. While five reviews were recommended in the IETF process, the WG might settle for three, if these were done by people independent from each other. Participants recommended to request industry associations (like the Domain Name Association) to weigh in.

GoDaddy list of drafts

A list of four possible registry extensions was presented in Berlin (briefly) by GoDaddy:

- on validation;
- on unavailable names;
- on fees (there are also several fee-related drafts in the RegExt document list); and
- on an API for third-party service providers.

For validation, GoDaddy proposes a registry extension that will allow a registrar to send in data from a domain name requester to see if they will be able to register. Kolker explained that validation was too different in various registries (length of phone numbers in different geographical regions for example). This would allow the registrar to tell the potential registrant if he will be able to register the requested domain with the data provided beforehand.

Domain Connect will allow service providers to discover a DNS provider more easily and later on modify DNS records. The unavailable names draft addresses intends to allow for transparency on domains unavailable for registration because they are registered, reserved, policy reserved or IDN variant reserved. Better transparency on non-standard fees are the goal of the fee-related registry extension that basically defines the file format for the storage of non-standard domain name fees and related details for a top level domain name registry.

The key-relay-extension already advanced in the IETF WG process to IESG (draft from SIDN) faces a discussion about an IPR claim from VeriSign that includes vague information about the potential licensing regime to be applied. Stephen Farrell (Trinity College Dublin) requested a clarification of the IPR situation before moving to publish the document as RFC: "The IPR declaration says that license terms will be available 'later.' As things stand, I don't understand how the WG can have made an informed decision in that case." VeriSign has made a similar IPR claim for the bundling draft of CNNIC.

A list of all active RegExt drafts is available here: https://datatracker.ietf.org/group/regext/documents/

.home – IETF leadership in damage control mode

In April 2016 the RFC Editor published RFC 7788, the Homenet Control Protocol. A proposed full standard, the document standardizes the use of .home as a TLD for the homenet naming architecture. The authors – Markus Stenberg (former Cisco, now independent), Steven Barth (Linux, openWRT developer) and Pierre Pfister (Cisco) obviously did not take RFC 6761 into consideration as a process for special TLD allocation by the IETF, not to mention the ongoing fight around the RFC.

The original text in RFC 7788 includes:

> "Names and unqualified zones are used in an HNCP network to provide naming and service discovery with local significance. A network-wide zone is appended to all single labels or unqualified zones in order to qualify them. ".home" is the default; however, an administrator MAY configure the announcement of a Domain-Name TLV (Section 10.6) for the network to use a different one. In case multiple are announced, the domain of the node with the greatest node identifier takes precedence."

Terry Manderson, Internet Area Director (and one of the responsible IETF Area Directors, as well as Director DNS Engineering at ICANN (responsible for L-root), said the .home "allocation" slipped through the WG last call, IETF last call, Internet Area director review, IESG review, IANA review and got published. "We had a break in process", Manderson acknowledged. The

authors were not to blame, there was no suspicion that they wanted to 'ninja this TLD through' - or if they did, they did a good job", he said.

Manderson, supported by the IETF Chair Jari Arkko, apologized for the mistake. They asked for a remedy to what was described as "effectively allocating a domain inside an IETF RFC". Arkko said he wanted to see an immediate fix with an RFC Errata and, in parallel, a process dealing with the bigger naming architecture issues.

The following possible options were presented to the homenet WG:

1. Publish a new RFC containing the errata that removes all mention of .home or any other tld, obsoleting RFC 7788.
2. Publish a new RFC removing all mention of DNS or naming in addition to the .home reference, obsoleting RFC 7788.
3. Publish an RFC that explains the procedural failure that occurred and normatively updates RFC 7788 such that it no longer identifies .home as the default.
4. Publish an RFC that explains the procedural failure that occurred, normatively updates RFC 7788 and QUICkly moves for .home or .homenet via RFC 6761

Manderson underlined that the added option 5 proposed by Ted Lemon (do nothing) was not an acceptable option. The IETF leadership, especially Arkko, rejected complaints from Lemon and Stuart Cheshire (Apple) about the mere political nature or "Monty Python reasoning" (Cheshire) of the pushback against 7788 and .home. Some participants warned that updating with keeping 7788 in any form would confuse implementers.

Cheshire, Lemon and the 7788 authors (Barth, Pfister) warned that just deprecating .home would result in implementers using a variety of extensions for the homenet naming, resulting in interoperability, but also the poisoning of additional non-in-use TLDs. Cheshire said that .home was poisoned anyways according to ICANN's studies, so it would be best to continue to use it as the "natural solution".

Procedure-wise, an RFC 6761 process could be performed, but should not take too long. DNSOP Chair Tim Wiczinski said that .onion only took six months to pass. DNSOP Co-Chair Suzanne Woolf said that beside the process break ("there is a WG that is supposed to make this easier"), there were still "technical problems not fully aired".

The hum of the WG on which of the options to take was very inconclusive. Follow-up discussions on the mailing list have still being very slow.

Battle on Special Use Domains continues

Meanwhile it is pretty clear that there is an urgent need to put the discussion over 6761 to a (consensual) end. The DNSOP WG staged another round of discussions in Berlin. But once more, no conclusion could be reached about which of the two competing documents will become the WG documents to develop the future IETF position on special name allocations. The homenet discussion only illustrated how urgent the need was.

The documents have converged, said Woolf in her introduction, but the two author groups (which keep adding authors: Warren Kumari, Google, is now listed on both documents, Geoff Huston joined the adpkja draft and worked on the rewrite) could not agree on how to move forward.

Problems identified by the draft of Alain Durand (ICANN), Peter Koch (Denic), Warren Kumari (Google) and Geoff Huston (APNIC) were split in issues with RFC 6761 itself, and the process of string evaluation through the IETF.

On RFC 6761:

- can be used to reserve any names, not just TLDs, could allow to ban registering specific names in any TLD
- does not mention obligation for requested TLD/ string to be published in form of RFC document
- no clarity on who will carry out the evaluation of applications in the IETF
- no formal criteria evaluation criteria
- leakage, granting application brings no guarantee that special names won't be sent over the net
- no easy to use guidance for those affected by the special domains listed in 6761 registry (only pointer to potentially complicated documents)
- potential waste of space, if intended usage of special name fails

- to complicated process for those interested in experimenting with a special name

On evaluating candidate strings and the relationship to the ICANN process:

- IETF does not have process to evaluate candidate strings for trademark, name collision, other issues (appeal mechanisms: IAB and IESG)
- IETF review process not foolproof (as .home illustrated).
- two parallel processes to assign TLDs: 6761 IETF (adhoc fashion), and ICANN's gTLD program
- significant risk of conflict when both the IETF and ICANN want to register the same (or similar) string, and no cooperation mechanism for such cases
- potential for anti-competitive abuse, using the special use application process to block a gTLD application of a competitor applying for name (or similar name) from ICANN

Problems identified by Lemon's/Droms'/Kumari's document are:

- no formal coordination between the IETF and ICANN name assign functions
- no power (neither ICANN or IETF) to prevent use of strings by somebody
- demand for more than one name resolution protocol (but lack of switch signal between protocols)
- queries for non-DNS names end up to being sent to authoritative servers
- uncertainty of RFC6761 process (took 10 years before 6761 first was used for assignment), slow
- resistance in IETF to assign names outside of DNS (because of lack of switch, sense of IETF/ICANN owning the space, or potential competition and legal dispute)
- mistakes have been made in RFC 6761 assignments, due to lack of clarity
- failure of 6761 to provide assignments for additional TLDs [I-D.chapin-additional-reserved-tlds]
- no process exists to avoid that names are accidentally assigned by the IETF (see RFC 7788)
- use of registry is inconsistent, some specify registry entries and delegations, some don't
- no safe, non-process-violating mechanism for ad-hoc allocation of special-use names in place
- 6761 talks of Domain Name (implying that DNS is used)

Lemon's assessment of the relation and tasks of ICANN/IETF may be rather controversial.

"The assignment of Internet Names is not under the sole control of any one organization. ICANN has authority in many cases, and could be considered in some sense the default. IETF has authority in other cases, but only with respect to protocol development. And neither of these authorities can in any practical sense exclude the practice of ad-hoc allocation of names, which can be done by any entity that has control over one or more name servers or resolvers, in the context of any hosts and services that that entity operates."

Lemon, while underlining that his draft provided a more complete set of problems (and therefore a better problem statement, instead of offering solutions) took a step towards compromise, by offering to use the Durand draft as the base for a combined document. The WG Chairs did not follow-up on the offer during the session. Instead, Suzanne Woolf announced a potential interim meeting in the coming weeks.

One solution that has been offered for some time is the .alt special use top level domain, or in the words of author Warren Kumari, a TLD label in non-DNS contexts or for names that have no meaning in a global context; the possible sandbox for people who want to use a space to experiment could, for some special domains, become even a permanent home, Kumari said. It would not solve the problem of 6761 in general.

QUIC goes IETF

There was a lot of applause for Google, which brought its new transport protocol QUIC into the IETF for standardization. Adoption of QUIC as a work item for the IETF and start of the WG was approved with close to no objection in Berlin. The BoF drew a record crowd of around 370 people, illustrating the interest for Google's protocol.

After developing QUIC for three years, the company obviously thought the code stable enough to head into the IETF standardization waters. The spec as presented now was not sacrosanct, leaving IRTF Chair

Lars Eggert, co-chairing the BoF, stating that it was not "intended to standardize the existing code" as is.

QUIC is based on UDP and is characterized by one round-trip or (for returning servers) zero round-trip times and a promise of better encryption. One major promise is that it protects not only the content but also withdraws parts of the header. Headers would be "fully authenticated and mostly encrypted".

Work items for standardization, TLS-related

An agreed work item put on the agenda during the BoF by the Google developers and co-author Martin Thomson (Mozilla) was the replacement of the QUIC genuine security protocol with TLS 1.3, for which the IETF adoption of TLS 1.3 has paved the way. Thomson pointed out that TLS 1.3 had in fact been inspired (one might even say pushed) by the QUIC security protocol.

Issues yet to be addressed during IETF standardization were poor performance of QUIC in the event a client was forced to downgrade from their preferred version (downgrade attack) and the need to avoid passive linkability of connections (via the configuration identifiers sent for a handshake). TLS 1.3 also did not include QUIC's ability to include a cookie (in HelloRetryRequest).

Other differences between TLS 1.3 and QUIC are discussed in the TLS-related QUIC draft. Standardization now starting can also be expected to check on the balancing of requirements – efficiency vs. security. One example is that a QUIC connection identifier, which allows for session resumption while roaming, allows a passive observer to correlate connections.

One feature underlined during the BoF was that due to the modular approach chosen for QUIC, TLS 1.3 could later be replaced by TLS 1.4 or other crypto solutions in the future.

A mega-working group: one to rule them all?

Work on the QUIC draft has been split in four different documents:

1. a core transport protocol which describes the connection establishment flow control, etc.
2. a document on loss recovery mechanisms (author Jana Iyengar underlined that QUIC learned from TCP in that regard)
3. the TLS document on the crypto handshake
4. a document on mapping Http semantics over QUIC

As the protocol suite spans multiple WGs, transport, security and applications, questions were raised about keeping all the work into one WG, which could result in a mega WG draining other areas and WGs. That would be inconsistent with the IETF's usual work style. Some think that by agreeing on one WG, the IETF condones a layer violation. Splitting the work up in different WGs, on the other hand, is not in the interest of the QUIC authors, as it can result in delaying the work, something Iyegnar cautiously pointed out.

Good (IETF) citizens, supporters

Google's QUIC effort looks like a considerable success at this point. The company kept the protocol to itself for early developing stages, but kept the IETF community in the loop with Bar-BoFs and presentations about progress. With the TLS 1.3 spec now designed to the satisfaction of the QUIC developers and several implementers joining the QUIC bandwagon (Akamai, Microsoft's Christian Huitema, Mozilla), Google decided to take the next step.

According to Ian Swett from Google, QUIC is "used for every major Google site on Chrome, desktop and Android, and many Google Android Apps – Android YouTube QUIC enabling is underway".

Miroslav Ponec (Akamai) said that Akamai has already deployed QUIC to all edge servers for http delivery and that the company was slowly enabling traffic. The long-term plan was to include QUIC as default feature in all Akamai products. For its implementation Akamai has used the Chromium QUIC Code and added Akamai's congestion control algorithms and media acceleration software development kit for app integration on the client-side.

Christian Huitema said he had implemented it for Microsoft, but had nothing to ship for now. His company would only ship fully standardized products. Huitema, who has been an early QUIC supporter, applauded the core QUIC idea of having an "encrypted header" and "very clean stack on top of UDP". He still wanted to work on privacy in QUIC. The attractiveness of QUIC was that compared to TCP over TLS it was smaller and has made an important design

decision to give up some information in the header. That could provide a smoother way to innovation, as middlebox traversal could be ensured.

Next steps for Huitema would be interoperability tests and performance questions. Google figures released in Berlin showed that 93% of connections from Google service users could successfully use QUIC, 7% failed, with UDP blocked for user (4,5%) being the single most reason for failure.

Discussion about compromising with middleboxes was postponed to the Plus BoF. An effort to include a graph in the QUIC WG Charter text on "network path interactions" (proposed by Joe Hildebrand, Cicso) failed during the QUIC BoF.

Another rejection to approve a little "cooperation" with middleboxes

With TLS, and possibly QUIC gaining momentum, the heat is on middlebox vendors as encryption withdraws information crunched by their devices for traffic management, firewalling, etc. But the proposal to establish an IETF WG to create a protocol for explicit cooperation of middleboxes with applications and network providers failed after some spirited discussions.

Encryption and middleboxes, mobile networks

Following-up to the Spud BoF during the IETF92 meeting (March 2015) Joe Hildebrandt and Ted Hardie (Cisco), Mirja Kühlewind and Brian Trammell (both ETH Zurich) came back to present the proposal for a Path Layer UDP Substrate protocol (Plus). Plus is described by the BoF proponents (see proposed charter) as "common shim layer atop the User Datagram Protocol (UDP) to provide a transport-independent method to signal flow semantics under transport and application control, necessary to enable the deployment of new, encrypted transport protocols."

Natasha Rooney (GSMA) who co-chaired the BoF put in a bid for the mobile operator community who shared the same problem: encryption made the task to make best use of network resources (scarce spectrum) more difficult. Rooney underlined the mere technical reasons for the flow-information via a Plus, with the intention to do this "with the lowest amount of information possible, only to allow an intelligent decision" for resource assignment.

A mere trust model based on prioritization – allowing traffic users to flag their needs as loss sensitive or delay sensitive – could not be trusted. DPI was an "interesting", but also "dangerous" option. Rooney had been coming to the IETF since the "Managing Radio Networks in an encrypted World (MarNEW) workshop". Plus was a solution to serve both interests, those of the mobile operators that wanted to manage traffic and users who wanted to receive enough resource for their applications. Eric Rescorla (Mozilla) called the idea of aligned interests of carriers on the one hand and persons sending data on the other side a fallacy.

With the focus on encryption on the rise (and resulting problems of encrypted traffic traveling through the "middlebox internet") and the clear statement that endpoints would control what information they would provide, if any, the Plus proponents got close to receiving agreement to start the work.

Pushing back against shifting powers

In the end several, counterarguments were made to push back against establishing the Plus WG.

A major argument is the traditional IETF position that intervention on the path contradicts the end-to-end argument. Daniel Kahn Gillmore (ACLU) warned when the Plus mechanism could be used "for coercion" to make endpoints expose information in order to get over the network, networks would become even more powerful. The Plus proponents had not offered a full analysis of the potential consequences of this risk.

The ACLU technologist and others also did not fully buy the presented limitation of the Plus protocol to "enable encrypted traffic to flow". The Spud use case document includes a list of nine different use cases.

Brian Trammell perhaps was just a little too honest in acknowledging that there was a possibility to use Plus to require a client to insert a particular kind of metadata in the stream. That was an unsolvable problem, he argued, but Plus would at least bring transparency to the behaviour. The behaviour itself, according to the Plus proponents, was already widespread, without transparency for the user.

One powerful argument against the idea that Plus will only make explicit what has been done implicitly was explained by Christian Huitema (Microsoft). The moment the IETF would approve such an explicit mechanism as an RFC, "you would change the

balance of power between network and endpoint". It would fundamentally change the IETF approach and grant legitimization to a request for the traffic metadata.

The killer argument was finally made by Yoav Nir, who said, speaking for a middleware provider, he thought Plus was after too much information. The times when a flow started and ended were satisfactory from his point of view. The proposal seemed a "huge spec" and asking for "more extensions". After a very long discussion, the BoF ended with a hum favouring working on the scope and charter again. The BoF Chair asked to continue the discussion on the Spud mailing list for now, but one can expect the BoF proponents to be back.

Working Groups and BoFs

DNSOP: More drafts, TCP over TLS implementation work

To ease implementation of TCPoverTLS (Dprive) test servers at Nlnet Labs and Oarc are allowing to test features (TLS fast open, pipelining, providing OOOR, EDNSO keepalive, and for TLS: TLS on port 853, server certificate, EDNSO padding). For recursives servers patches are underway or already implemented. On the stub side the picture is still a little patchier. GetDNS supports all of the features on the stub side, more patching is on the way for several of the software solutions (see a perfect overview for recursive and stub implementation status here). Further Dprive work (Dprive did not meet in Berlin) has been proposed by Stephane Bortzmeyer (Afnic) here.

Ongoing work in the DNSOP include a bis-document of the DNS terminology RFC, which fixes things on which the WG could not find consensus, and also includes some terms not covered in the first version (wildcard, SEP, and so on). The idea of the terminology drafts was to define DNS terms used in various IETF standards, Paul Hofmann explained. Hofmann who is a co-author called on participants to do more reviews. Interestingly, the terms that turned out more difficult to find consensus on were classical DNS terms (like resolver or secure entry point); the newer stuff was easier (sometimes not implemented too much).

Additional work presented in Berlin was a close-to-finalised draft on aggressive use of NSEC/NSEC3 records "to generate negative answers within a range". The draft intended to decrease latency and resource utilization on both authoritative and recursive servers prepared by authors from JPRS and WIDE in Japan, and Warren Kumari (Google) will go to WG last call in the next edition.

New potential work items presented were a mechanism for DNS session signalling (presented by Ray Bellis, ISC), the idea being to have a new session signalling Opcode to carry persistent "per-session"type-length-values (TLVs)". An initial set of TLVs used to manage session timeouts and termination is described in the draft.

Questions posed to Bellis included how the session signalling would interact with the new http2 wire format and also with the proposed Plus-layer (see Highlights). If Plus or Google's Quic protocol would define a session layer then it would be substrate for this, said Stuart Cheshire (one of the co-authors). The document was defining the requirements of the lower layer. Cheshire asked if the document should go for the minimal route (four-byte header followed by payload) or take the EDNS0 Opts approach (would be masquerading as a resource records). There was no discussion, but there is an interest to push this through quickly.

Also presented was the possibility to allow for multiple answers from authoritative DNS servers to "predict" some follow-up answers, prepopulating the caches in recursive and thereby decreasing the latency for end users and also load on the recursives and authoritatives. Questions on the proposal included if this would introduce "any" requests by the back-door and if it could be an amplification factor.

Several documents (not presented) are currently in different stages of the IETF process:

- draft-ietf-dnsop-resolver-priming
- draft-ietf-dnsop-refuse-any
- draft-ietf-dnsop-edns-key-tag
- draft-ietf-dnsop-rfc2317bis
- draft-ietf-dnsop-attrleaf

There is a considerable long list of DNSOP RFCs passed in the first six months of 2016:

RFC 7766 *(was draft-ietf-dnsop-5966bis)*
DNS Transport over TCP - Implementation Requirements
RFC 7793 *(was draft-ietf-dnsop-rfc6598-rfc6303)*
Adding 100.64.0.0/10 Prefixes to the IPv4 Locally-Served DNS Zones Registry
RFC 7816 *(was draft-ietf-dnsop-qname-minimisation)*
DNS Query Name Minimisation to Improve Privacy
Errata
RFC 7828 *(was draft-ietf-dnsop-edns-tcp-keepalive)*
The edns-tcp-keepalive EDNS0 Option
RFC 7871 *(was draft-ietf-dnsop-edns-client-subnet)*
Client Subnet in DNS Queries Errata
RFC 7873 *(was draft-ietf-dnsop-cookies)*

Domain Name System (DNS) Cookies
RFC 7901 *(was draft-ietf-dnsop-edns-chain-query)*
CHAIN Query Requests in DNS

Side Event ZSK/KSK signing

Matt Larson (ICANN) and Duane Wessels (VeriSign) gave another update on the ZSK and KSK rolls, underlining close cooperation in preparing for the rolls. The smooth transition from a 1024 bit to a 2048 bit RSA key for the ZSK – see timeline below – will signal a green light for the second, more complicated roll of the KSK.

ZSK roll – general rehearsal for rolling the key

The ZSK roll using the regular pre- and post-publication system will result in a parallel publishing of the different-sized keys for 20 days and after that VeriSign will keep the key for an additional time span into slot 2 of Quarter 4 2016. Wessels did not give an exact number of days, but in case a roll-back was necessary, the ZSK can be rolled back immediately (signing with smaller key again) during two plus slots. But this would only be done in case of a major and unforeseen technical issue – "not because one ISP experiences a problem", Wessels said in Berlin. For minor issues, VeriSign would recommend to stop validation. On 30 October, the key would be fully removed, after which there was no going back to the old key.

For the ZSK key size roll the technical issues addressed by VeriSign are potential fragmentation or truncation of queries due to the larger key material. Wessels presented figures once more that see issues if at all only for extremely long domain names. Tests can be performed beforehand by everybody over keysize.test.versignlabs.com.

Bigger fish: KSK roll-back

Policy-wise the KSK roll is now a done deal, Matt Larson and his boss, David Conrad, assured operators at the IETF. One thing could still delay the roll: if a major issue is detected during the ZSK roll, it could result in a change to the KSK plans. Operational realities dictated how ICANN did things, they said.

Jim Read, RIPE DNS WG Chair, reported how failure due to DNSSEC validation was difficult to discern as such as it developed slowly (with Caches slowly

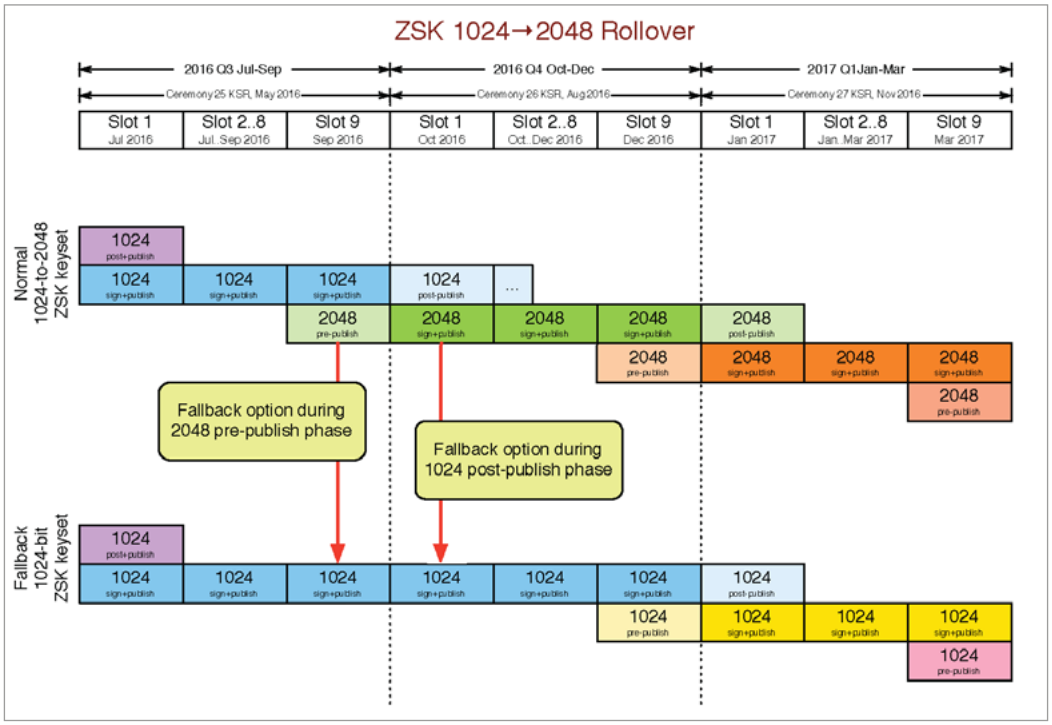

running empty). He told this reporter the biggest issue was the unknown unknown, meaning failure in applications (for example shop systems) where DNSSEC validation had been embedded without the user of the system even knowing.

Larson presented the timeline for the KSK roll, which is different not the least for not having parallel signing periods.

Key rollover from KSK 2010 to KSK 2017 is "tentatively scheduled for 11 October 2017 after KSK 2017 has been published ten regular 10 day-slots in a row. (Key rollover schedule system is organized in phases 9 to 10-day slots per Quarter). Starting 19 September, the DNSKEY RRSet will contain both KSK 2010 and KSK 2017 to allow systems to pick it up either manually or via 5011 automatic key rollover. KSK 2010 will be set to revoked on 11 January 2018 and be removed from the zone on 22 February 2018. In August 2018 KSK 2010 will be deleted from all Hardware Security Modules.

Major steps are:

2016-07-22 The KSK roll project plan made public for review and discussion
2016-10-27 KSK-2017 KSK is generated
2017-02 KSK-2017 KSK is operationally ready
2017-03 KSK-2017 KSK is published on the IANA web site
2017-07-11 KSK-2017 KSK is published in the root zone
2017-09-19 Response size increase due to ZSK rollover
2017-10-11 KSK-2017 KSK is used for signing the root zone keyset
2018-01-11 KSK-2010 KSK is published as revoked
2018-03-22 KSK-2010 KSK is removed from the root zone
2018-08 KSK-2010 deleted from all HSMs
2018-08-31 The KSK rollover process concludes

At the end of the IETF week, on Friday 22 July, ICANN has published five operational documents to prepare for the KSK rollover.

One is on operational steps to take (including the tentative timelines above). The second explains back-out scenarios on trust-anchor changes (1), response size changes and DNS RRSet changes. Criteria for the three would be:

- validators have incorrect trust anchors, problems with tools based on the Internet Draft (or future RFC) describing the root trust anchor XML files" (for 1);
- fragmentation issues, significant retries for DNSKEY records (exempting retries related to fall-back from UDP to TCP) or operators (relying on automated updates) reporting that their validators are not seeing the new trust anchors (for 2);
- higher frequency retries for DNSKEY records by a significant rise in queries from a large distribution of ASNs. Operators (relying on automated updates) reporting that their validators are not seeing the new trust anchors. Reports of distress.

The other three documents explain various tests to be performed in preparation for the roll.

What seems still to be lacking is an extra communication plan for now. What ICANN will provide are two sets of public test environments: a testbed for multiple key rollovers in real-time and one providing continuous key rollovers using accelerated time.

According to the document on external Tests: "The target audience for the real time 5011 environment is DNS resolver operators and is designed for validating deployed software configurations and can be used in production environments. The accelerated 5011 environment is intended for software developers. Because this environment requires modified RFC 5011 timers as well as a special root zone, it should not be used in production environments."

More testing can be done using http://toot-servers.net or http://keyroll.systems (test environments prepared by Rick Lamb and Warren Kumari, Google, respectively).

So many human rights groups at IETF

It becomes rather difficult to follow all working groups in the IETF/IRTF that focus on human rights issues related to the standardization work of the IETF. The most concrete WG at this point is the new meeting venue WG (MTGVENUE) which according to its charter will produce documents to address two issues:

1. "A specification of the geographic IETF meeting policy, currently described as the "1-1-1-*" policy. The policy going forward is up to the working group.
2. A specification that describes the detailed meeting venue selection process and criteria, the contents of which are also up to the working group."

A draft on the meeting venue selection by Fred Baker

was briefly discussed in Berlin. It resulted from the discussion about gay politics in Singapore. The IAOC meanwhile announced that they would not relocated the Singapore meeting, but during the meeting there was a look for back-up locations (for short-notice changes) in case of major disruptions at venues. Such ad-hoc emergency relocations should be mentioned in Baker's draft as well, participants wanted.

Baker's draft for now describes how the IAOC selects venues explaining procedures and also some principles, objectives and criteria for the selection. Political considerations were not part of the criteria, the draft states: "The IETF does not make political statements. We do not decide who is or is not a country, and we do not choose or not choose venues based on political criteria." Instead, participation, geographic rotation and available services (Internet access), hotel and food options are on top. Declared objectives of the IAOC are:

- Advancing standards development
- Facilitating participation by active contributors
- Sharing the travel pain; balancing travel time and expense across the regions from where IETF participants are based.
- Encouraging new contributors
- Generating funds to support IETF operations in support of standards development, including the Secretariat, IASA, and the RFC Editor.

A document on how to balance the different objectives (the IETF meeting selection morale document, so to speak) has been prepared by IAB Chair Andrew Sullivan. Sullivan in the documents puts the criteria in an order of descending importance: inclusiveness, co-location of attendees, network access, safety and security and, last, affordability. The so-called "one-roof"-policy (allowing the meeting and accommodation to be under one roof as much as possible), financially interesting maximal attendance and geographical outreach were non-goals, Sullivan said. Yet the ordering and non-goal approach of the IAB Chair generated controversial discussions.

An additional document in the pipeline of the MTGVENUE WG are one on the "Definition of Participation Metrics for IETF Attendees".

The question about how to solve the "Singapore issue" was discussed despite the "non-political" statement in the MTGVENUE. Alissa Cooper (IAB) said more guidance with regard to selecting venues was still a desideratum, especially as the model to choose six safe and political correct hubs, while discussed, would not happen.

An attempt to understand potential human right criteria for meeting venue selection processes was made with an additional, informational BoF. Former World Bank lawyer Motoko Aizawa from the Institute for Human Rights and Business, which is also working on ICANN's Human Rights obligations and status during the IANA transition presented a list of criteria to consider in the selection processes:

- a country's general human rights track record (via external data bases)
- visible patterns of abuse of rights central to the organizations' values/mission
- listening to reports from local peers and NGOs
- legitimizing effects (question, what will be relation to government during event)
- safety threats to special groups of the participants

Human Rights in Protocols, Internet Access for all

Beside the practical work to hammer out policies how to deal with the IETF meeting selection venue, there are for now two more human rights related WGs who both met in Berlin. The Human Rights Protocol Considerations Research Group had a brief exchange with David Kaye, UN Special Rapporteur for Freedom of Expression, who joined the meeting remotely. Kaye who said there was a "neat overlap of IETF work and his mandate" called for input from technologists for his work on how private ICT sector (including ISP, telecom operators, and equipment providers, but also academic and technical communities preparing standards) implicates freedom of expression.

He underlined the UN had no intention to regulate standards bodies like the IETF. Instead he expressed support for the IETF aligning human rights with technical protocols and said he was concerned that core values of the IETF would likely to be challenged as governments seek to undermine multi-stakeholder governance ideas by governments: "I hope you can maintain this work in the face of challenges to the protocols", Kaye said.

The HPRC RG is preparing to have its first document last called, a long research document on how human rights figure in protocols by Niels ten Oever

(article19) or even how they are violated. Examples of RFCs analysed with regard to their human rights preserving/non-preserving status range from IPv4 over DNS to Middleboxes and DDoS attacks. The document also develops a model for human rights protocol considerations in the IETF work. Ten Oever hopes that over time the RG can become an IETF WG.

There are two drafts (not RG documents, but related) that make an interesting reading, which are Mark Nottingham's follow-up to the IETF remote encounter with Edward Snowden, The Internet is for users and an overview over worldwide censorship technologies by Joseph Hall from the Center for Democracy and Technology (and others) which was presented in the Security Area meeting in Berlin.

Finally, there is a Research Group that looks into the much cherished motto "Global Access to the Internet for All" (Gaia), which according to its charter wants to create visibility for the issue access, shed some light on cost issues in different geographies and on deployment. As the HPRC RG it wants – in the longer term – to provide input to the standards development. In Berlin, examples for low-cost/free WIFIs for refugees and community networks (Freifunk in Germany, Funkfeuer in Austria and wlan Slovenia) in general were discussed. Jose Saldana (University of Zaragoza) presented the Horizon2020 funded Wi-5 Project that proposes an architecture based on an integrated and coordinated set of smart Wi-Fi Aps intending to reduce interference and develop business models for low cost wlan access.

Glass to Glass Internet Environment – Dispatch meeting

Glen Denn, NBCUniversal, and Leslie Daigle, ThinkingCat Enterprises, presented the Glass to Glass Internet Ecosystem (GGIE) work ongoing at the W3C, driven by the GGIE Taskforce. The problem to address, Denn described, was to better organize video distribution through the networks.

With more and more video devices being connected and resolutions going towards 8k, video was the number one bandwidth "eater". Consumption outpaced the growth of bandwidth capacity, despite measures like better codecs, caching and content delivery networks or new transport protocols introduced for mitigation. Without more fundamental steps, quality downgrading would be the result, according to Denn.

Bringing GGIE to the IETF, the proponents intend that the IETF takes on some of the 33 use case scenarios in six categories (use of creator and user identity, , content discovery mechanisms, streaming, identifying content and measuring content, location and accessing content at the network layer, capturing digital video).

For the IETF location and accessing content at the network layer could be of particular interest, as it wants, for example, to define a content URI carrying the video identifier (uri:eidr-s:F1F8-3CDA 0844-0D78-E520-Q), content identifier namespaces, and finally a content lookup service query protocol. Deen presented Mars (Media Address Resolution Services), a service that he said would find caches nearby.

Questions posed to the proponents included privacy and linkability via the "camera identifier", the special resolution services and potential resolution conflicts. More general points that were made were that DRM could stand in the way of locating close-by cached content, and that there was no big incentive for big content creators to join a federated system.

There was a hum on the question of whether or not people thought there was IETF-relevant work in the material presented for GGIE, it was supportive of more work at the IETF. According to Leslie Daigle, "work will continue to pull apart the problem(s) into pieces that can be worked on in one or more groups going forward."

Cryptech presents Krypto Board for Alpha test

After just two years of development work, the Cryptech developer group presented their crypto processor board for alpha testing on the week end before the IETF. While the engine was not yet ready for DNSSEC zone signing – at least for a little larger zone –, according to Alexander Mayrhofer, nic.at, the progress was "impressive". The board, which can be bought for $800 USD at the moment, can be put to work for smaller signing operations.

IETF news

The IETF endowment (now to be found under sustainietf.org) managed to boost its fund to $3.9 million USD through large donations (a million or close to a million) by RIPE, ARIN, Afrinic and ISOC. The goal according to Kathy Brown (ISOC) was to reach $20 million USD in the next two years. The endowment was started to provide another revenue stream for the IETF (in addition to ISOC's contributions from PIR earnings), possibly to be able to keep participation fees stable (or reduce them?).

IETF Chair Jari Arkko updated the plenary of the Berlin IETF on the IANA transition. IETF leadership and the IETF Trust were writing and preparing to execute the contracts to have the IETF Trust act as a home for the trademarks ("IANA") and domain names ("iana.org"). Work was proceeding, Arkko said.

An interesting issue on the IPR management in the IETF came up during the plenary meeting (addressed by Philip Hallam-Baker) which was that the growing trend to use git hub made it necessary to consider storing git hub content in order to being able to answer future IPR legal disputes.

The Nominating Committee currently is looking for a successor to the IETF Chair Jari Arkko. Arkko will step down after two two year terms and never has the word, that the IETF Chair is the organizations CTO (Chief Talking/Traveling Officer, created by Scott Bradner) be more true.

IETF97 will take place in Seoul on 13-18 November 2016.

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 53 full and 9 associate members – together, they are responsible for over 80% of all registered domain names worldwide. The objectives of CENTR are to promote and participate in the development of high standards and best practices among ccTLD registries.

CENTR vzw/asbl
Belliardstraat 20 (6th floor)
1040 Brussels, Belgium
Tel: +32 2 627 5550
Fax: +32 2 627 5559
secretariat@centr.org
www.centr.org

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