Council of European National Top-Level Domain Registries

Report on IETF97

Seoul
13-18 November 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

DDoS attacks: Fears of Regulation or of an "Oligarchonet"

The recent DDoS attacks on DNS provider Dyn were one major topic at the IETF in Seoul. No more technical details were presented on the attack, which saw a hundred thousand cameras connected to the net used as minions for an extended attack. However, the search for solutions on the ever larger attacks is on and fears are growing that regulators might step in. In a dedicated plenary debate, the community tried to look ahead on potential technical, economic and regulatory consequences.

Bruce Schneier, well-respected Internet security guru, did ask for smart regulatory steps during that hearing in the Congress Energy and Commerce Committee (Schneier spoke of market failure to address the lack of security in IoT devices). Henning Schulzrinne, long time IETF participant and once again Federal Communications Commission CTO warned that not much time was left for the community to solve the problem by itself.

The Resolver Club and other defence mechanisms

Nick Sullivan from Cloudflare gave a list of all the potential technical tools at hand to stem the attacks that have risen to hundreds of Gigabit/s magnitude. Depending on the choice of attack (Syn Flood, direct attack on authoritative, attack on resolvers, with classical reflection-amplification being on the decline), defenders had to choose from the tool box – and combine available tools.

For direct to authoritative resolver-attacks, Sullivan recommended fast filtering, smart filtering and getting filtering as close to the source as possible. A flood of requests to authoritative servers from non-resolvers should be treated as an attack, for example. The idea that in times of attacks, authoritative servers should only respond to well-known resolvers and answer other queries only if time and capacity permitted was discussed in various sessions over the meeting.

While Geoff Huston, APNIC, said that legitimate resolvers were a rather well-known group of around 8,000, there were those that strongly disagreed. Filtering based on the membership to the well-known-resolvers club (via whitelisting) would not only be against the idea of neutrality (a violation even of net neutrality?). It would also counter recommendations made more often recently to set up one's own DNS resolvers.

Classical defences include zero routing (blackholing and thereby blocking attack traffic at the expense of blocking regular traffic), spreading the load through anycast or by various routing techniques. Reuse of equal-cost multi-path routing (ECMP) or of options offered by the Berkely packet filter (BPF) mechanism (allowing to set policies for packet selection for either acceptance or rejection) could help, according to Sullivan. Machine learning and automation should help scaling.

Most of the classical defences come with unintended side effects, the much used rate limiting for example could result in triggering repeated requests and thereby amplification.

One alternative new way to help diminish the load on servers was discussed a little more broadly during the Internet Engineering Planning Group meeting on Sunday: it entertains aggressive negative NSEC3 caching (see below).

The bigger problem: How to preserve the Internet design

Andrew Sullivan, currently fellow at Dyn and IAB Chair, in the plenary debate in Seoul underlined that the increasing attacks ironically gave proof of the success of the Internet model. By opening the network to smart endpoints, a certain class of abusers were able to abuse the system. In an interview with this reporter, Sullivan said:

When you put this kind of intelligence at the edge and you have a network of networks, one of the things that happens is you have a certain kind of abusers that can do these things without really being able to be certain who they are. And the trade-off is that we get all the other benefits - the resilience benefit and performance benefit of the Internet, and this does mean that this is a vulnerability that is a part of the design.

The fact that for the Dyn attack, hacked IoT devices (Mirai Software controlled cameras) were used, was just an "accident". CPE devices, unpatched software and operating systems all offered vulnerabilities open for the botnet attacks. Andrew Sullivan said that calling IoT bad and discouraging against connecting the IoT devices of users to the Internet would be the wrong answer and would, in fact, be contrary to the Internet model.

Sullivan warned against two trends developing as a result from recent attacks. One was the stepping in of regulators, potentially unbalanced when considering positive effects and unintended consequences. Just a few days after the IETF meeting, US Congress in fact held a hearing to evaluate how they might react to the attacks through regulation or legislation.

The second trend Sullivan described was the adoption of technical steps against the attacks that would corrupt the original design of the neutral network of networks. Access controls or licenses for providers, white-listing for only large resolvers (discussed in Seoul) would violate the end-to-end and low barrier of entry concepts. There is already a looking problem that only large organisations are able to get the necessary defences in place and stay online under attack. Sullivan said:

"We need to ask ourselves: are we building an oligarchonet? We don't want an oligarchonet. That is not the goal. We had one of this in fact: that was the telephone system. We need to figure out how to improve some of those things."

At the same time, for companies like Dyn or Cloudflare or Verisign, there was no good argument against implementing stricter filtering on a bigger scale. What was good for his employer and other large DNS companies nevertheless was not so good for the greater Internet, but as long as the Internet would not respond to the garbage-problem, the trend "go big or go home" would continue. Down the road, the DNS market could see concentration in a way the email market had developed. For any change in the email, between 12 and 40 people had to be convinced, Sullivan said, otherwise it would never happen. This kind of development was a risk to the Internet.

Possible ways to mitigate

When talking about the role of regulation, the technical community is divided. Sullivan reflected on the calls for regulators to step in – as Schneier did, for example. But calls to make BCP 38 obligatory, for example, would not have helped with the Dyn attack (as there was no spoofing, but compromised machines involved). What he supported much more were calls for "negative" or "backdoor" regulation proposed during the plenary debate. For example, making certain security standards a condition for getting insurance would provide a smart incentive for house owners to select more secure IoT devices (and manufacturers to build accordingly), he said. Liability claims and courts granting damages to victims was also one avenue considered during the plenary. Talking to this reporter, Sullivan also mentioned the high cost of recalls for insecure devices (like the compromised cameras in the Dyn attack), as more market-originating corrective action. In general, he recommended to rely on these "negative" regulation and technical measures in the first place before calling for legislative action.

Potential technological steps, he said, were underway in the IETF homenet WG, which in essence made an attempt to secure the Internet on unmanaged networks – while keeping the option for homenet owners (end users) to participate as a fully operational subnet in the global Internet.

DDoS serving best argument to get on DNSSEC

Another technical measure to step-up protection against DDoS attacks discussed over the IETF week was the NSEC/NSEC3 aggressive caching. Geoff Huston, APNIC Chief Scientist argued at the Internet Engineering Planning meeting (a gathering regularly organised immediately before the IETF main event) that only adding capacity was insufficient in the long run, but promoted the negative caching as one potential step besides more filtering or whitelisting of recursive servers.

Aggressive caching with NSEC and NSEC 3 will allow recursive resolvers to answer attacker generated queries for non-existent domains. With authoritative answers already cached about domain ranges that do not exist, there was no need to send attack queries on to authoritative servers or root servers. Attack traffic therefore would be stopped from traveling further up the tree. Aggressive NSEC caching has been proposed in a draft RFC to the IETF. DNSSEC deployment necessarily has to be in place, though. The proposal was, according to Olaf Kolkman, ISOC, the best commercial argument for DNSSEC so far.

DNS Privacy: Pushing protection up the tree while implementing DNS over TLS from stub to resolver

With standards work completed on TLS-secured querying of recursive resolvers (from stub resolvers), the DPRIVE WG considered how to proceed. There is rough consensus (at least from DPRIVE participants gathered in Seoul) that the WG should go on to secure traffic further up the tree, from recursive resolvers to the authoritative servers/root servers. Stéphane Bortzmeyer (AFNIC) presented the basic issues and questions to answer before work gets underway, namely: should TLS be reused (or should another mechanism be chosen), and how should, if at all, authentication be performed?

It's more complicated: Encrypting traffic between recursive to authoritative

Bortzmeyer noted (and explained in a short draft document) that securing the link between recursive and authoritative resolver was different in several ways from securing the stub-to-resolver part of the tree. For the stub to iterative resolver part, resolvers were known and there were normally few of them, so static key pinning was possible. That was fundamentally different for the resolver-to-authoritative link, where unknown servers were queried.

Concerns were therefore expressed during the session that using TLS for the connections with authoritative servers would open these to attacks because of the overload from connection management. Paul Hofmann, ICANN, warned against the effects of many parallel TLS connections being opened. Bortzmeyer said that he expected TLS 1.3 (with smarter connection management) to be already in place and could in combination with TLS fast open reduce the burden and also potential latency issues. Standardization of TLS 1.3 is currently underway. While the WG will have to make up its mind on the concrete mechanisms, TLS use seems to be getting support.

In the draft document Bortzmeyer also listed options available for authentication. Keys could be put in the name itself (DNScrypt-style) or regular PKIX, as well as DANE (key in DNS). For the latter, DNSSEC deployment was necessary, though. Another issue to be addressed was how to deal with servers offering different authentication mechanisms (for example, one server offering PKIX, the receiving server DANE). With Qname minimization opportunistic mode only might suffice, one participant offered.

The WG will need to re-charter, according to several participants (including IAB Chair Andrew Sullivan and Internet Area Director Terry Manderson). Sullivan said securing traffic to authoritative servers would concern a lot more people and not re-chartering could invite complaints from the wider community later.

Pushing for more Privacy in the DNS, Stubby and possible problems with 853

With the work on encrypting the link from resolver to authoritative soon to be started, results of the first phase (DNS over TLS) still have to be implemented. In what was lauded as an excellent tutorial on the status quo of DNS privacy work, Sara Dickinson (Sinodun), presented on Sunday the background, motivation, remaining technical issues and implementation work of DNS over TLS. Two years after the Snowden revelations, the Dprive WG had reached important milestones. The WG first rebutted the long standing idea that the DNS data was all public (and fair game) with RFC 7626 (DNS privacy considerations) and went on to standardize DNS over TLS standards suite, including RFC 7766 and RFC 7828.

More work is ongoing in Dprive on what profiles users want to choose (strict authentication or fall-back to unencrypted to receive traffic), and on padding (see Dprive WG below).

With DNS over TLS standardized, there is still a need to implement. In an effort to push ahead with implementation, Sinodun together with NLnet Labs developed Stubby. Stubby is based on the getDNS library. Originally developed for to assist key management for DNSSEC, getDNS now provides basic modules for the DNS over TLS key management. According to Dickinson, Stubby allows users to set their profile (do they want to only receive authenticated, encrypted DNS answers or do they want to fall-back to unencrypted to receive DNS traffic).

The application will also list those recursive servers that provide (for test reasons) the TLS encrypted

traffic – so far there are four DNS over TLS test servers (from Surfnet, Daniel Kahn Gillmore at the ACLU), more test server setups (like at OARC and RIPE) might become available. Testing Stubby in connection to these servers should help, Dickinson said, to evaluation potential problems, for example with filtering the new port 853, which has been dedicated to encrypted DNS traffic.

In the tutorial, Dickinson, jointly with Daniel Kahn-Gillmore (ACLU), underlined the need to push ahead with encrypting DNS traffic as being one part to the overall effort to encrypt. Without DNS being secured, other efforts might be in vain as DNS, for example by allowing a peek into the first connection to a mailserver (or DNS server), giving away important metadata. In her opinion, when network traffic was a boat with many places to patch for privacy/security, the DNS was the gaping hole in the boat. In the discussion following the tutorial, there was a proposal to make the DNS problems more visible not only to users, but also to data protection officials – in order to gather support for the efforts to change from unencrypted to encrypted DNS traffic.

Concerns about the Funding of the IETF

The funding of the IETF activity has become an issue of concern – some observers speak of a mild panic, and with the Seoul meeting once more written in red numbers, discussion on funding and the future administrative structure seem to be necessary. Outgoing IETF Chair Jari Arkko, who has tried to push the IETF endowment as a new source of income, announced in Seoul a review of IASA, the IETF Administrative Support Association.

With only 996 participants, the IETF meeting in Seoul fell short in registration fees ($-106,000 USD below budget, actual revenues at $658,000 USD). The sponsorship revenue was also below of what was projected ($-62,000 USD, actual revenues at $491,000 USD). Connectivity cost had to be paid for the Seoul meeting (normally, this is borne by the host, but there were some complications due to local connectivity being filtered).

While losing money in one of the three annual meetings might be overcome by the continued cuts in expenses, 2016 has developed in a financial challenge overall. Both the Buenos Aires and Seoul meetings came in under the expected revenues. The overall shortfall after the first two meetings in 2016 (Buenos Aires and Berlin) was already $501,000 USD. ISOC was already expected to pay an additional sum of $169,000 after Buenos Aires and Berlin. With the bad results from Seoul, another bill of $200,000 USD is expected to end up at ISOC's door step, bringing it up to $369,000 USD.

Obviously, one general problem has been the difficulties in attracting sponsorship money (for all three meetings during the year). In an effort to alleviate the sponsorship issue for the first meeting in 2017 in Chicago (moved from Montreal), Ericsson (employer of the outgoing IETF Chair and one of the multi-annual sponsors) announced it would step in as a main sponsor for IETF98. While the IETF leadership tried to downplay the financial issue, this only illustrates that there might be a bigger issue for the IETF, despite the fact that ISOC is prepared to pay up for the IETF from its revenues from the .org registry (PIR).

A cost factor, some say, are the much enhanced remote facilities, which now include highly sophisticated video remote participation (meetecho).

Arkko's announcement of a review of the 10-year-old IASA support structure may be in part related to it. According to Arkko's announcement, "areas to look at include structure, financing & sponsorship arrangements, organisation, and ways of working". The outgoing IETF Chair also wants to address issues like the selection of venues, selection of personnel and the mechanics and interplay of the IASA and IAOC, Trust, the IAB and also ISOC. If Arkko hands over this project to his successor (potentially the first women in IETF history, Alissa Cooper, Cisco, being the most mentioned candidate in the race) or if Arkko will continue to champion this after stepping down remains to be seen.

All in all, 2017 looks like it will be a very interesting and possibly challenging year for the IETF with a new leadership, both Arkko and IAB Chair Andrew Sullivan stepping down (Allison Mankin, former VeriSign Labs, has already been chosen to succeed Lars Eggers as IRTF Chair), the funding issues and a potential structural reform on the way.

Venues for next year are Chicago, (once again) Prague (for the fourth time) and the much-debated Singapore (IETF100).

IETF dips its toes in the IoT arena with new Directorate

Following in the footsteps of other standards bodies, the IETF is also stepping up its profile on work related to the Internet of Things (IoT) standards. During the Seoul meeting, the Internet Area Directors (for the IESG) announced a new IoT Directorate of the IETF, reasoning that "the interest in IoT technologies in the IETF, and more broadly in industry and other SDOs, is continuing to grow and issues with regard to IoT are being raised." The new Directorate will coordinate ongoing work on IoT-related specifications and become a new focal point. The chosen Chairs of the Directorate are Samita Chakrabarti, Ted Lemon and Ari Keränen.

So far, the Things to Things Research Group (T2TRG) of the IRTF was that focal point with large summary meetings which had a standing link to work in the World Wide Web Consortium (W3C). The Summary meetings have so far provided an overview over ongoing work related to IoT (major IETF specifications include COAP, CORE). During the T2T Research Group in Seoul, a revived draft on security issues was discussed and there were calls to add security threats originating from IoT things to threats to them. In the W3C there is now a proposal to open a new WoT (Web of Things) WG that can do normative standards work (in addition to an already existing WoT Interest Group).

Meanwhile, in 2016, the IAB has held two special workshops on IoT: one on Semantic Interoperability (data model question) in Santa Clara in March and one on IoT Software Updates at the Trinity College in Dublin in June.

ISOC Panel on IoT

During the regular ISOC lunch panel that had chosen IoT as topic, Michael Koster, Samsung/SmartThings warned that there was still a lot of fragmentation and incompatibility in the IoT standardization space with the current hot spot being the data model for IoT. Much development and more communication between organisations had to happen there, participants on the panel agreed.

One of the bigger risks, according to Carsten Borman, was that some very large organisations could push the development towards favouring monopolies. Borman and other participants warned against the jailing or lock-in of things – sometimes called for under the pretext of security.

More communication, interoperability and open standards development – and much more education was necessary. Borman also said there might be a need for some regulation, but there had to be clarity first, which features, or measurement points had to be measured for approval, certified IoT devices. How interoperability testing as done by the University of New Hampshire InterOperability Laboratory, on which Erica Johnson reported, could be used for this was discussed inconclusively.

Patents and Standards: Are IPR claims blocking development of "domain secure transfer patent"?

In a rant, Job Snijders, NTT Netherlands, called on VeriSign to give up resistance to grant licenses for its secure domain transfers technology. VeriSign's contested IPR application has resulted in the RegExt WG to delay a draft on the secure transfers of DNSSEC signed zones. VeriSign was effectively blocking DNSSEC deployment, Snijders complained.

The proposal on the key-transfers has been on the agenda of the RegExt WG for some time. While there are various alternatives to solve the issue of how to transfer a DNS secured domain from one provider to another, there was no standard yet said Snijders. Currently, many providers just take off DNSSEC before a domain can be transferred, only after the new provider has taken over the domain, the signature is used again. This created a gap in security that could easily be targeted by attackers. Snijders called on VeriSign to get the issue resolved one way or another.

Several domain experts consider the application highly likely to be a failure, as there are several earlier IETF documents around on the issue, for example a document from as early as March 2011 from several DENIC authors. VeriSign's application at the European Patent Office has been withdrawn already. The one before the USPTO is still pending and VeriSign had currently no intention to withdraw, Scott Hollenbeck from VeriSign confirmed to this author. As long as the patent application is pending, licensing is difficult, experts point out.

The secure transfers patent application was only one

of several IPR issues that came up during the DNS related WG meetings in Seoul. Other IPR issues might influence the work on a standard format for zone captions (see DNSOP WG) and several RegEXT WG related documents (see here, here and here). Delaying work that has IPR claims is not obligatory, said one of the Internet Area Directors in Seoul. While RegEXT has delayed a document on DNSSEC secured domain transfers, the DNSOP WG now decided to continue work on a similar document on transfers.

Working Groups and BoFs

DNSOP WG: Discussion on special names postponed again

The DNSOP WG postponed discussion of the special names application procedures. A problem statement has finally been adopted, listing the issues with the parallel tracks to delegate names on the top level, but discussion was deferred to yet another interim meeting (no date set yet) before the IETF98 in Chicago. At the same time, the homenet WG will not wait any longer for the DNSOP WG to launch their application for a special name for .homenet, Internet Area Director Terry Manderson announced. And more pressure on the special names applications process might come from another proposal from Stuart Cheshire, Apple.

Cheshire made a case for ipv4only.arpa as being a special domain (non-DNS) use and therefore eligible for an IETF special use domain registration. Cheshire said that his draft would only clarify the "special-ness" of ipv4only.arpa. IPv4only.arpa was specified for how a client can discover its network's NAT64 prefix (RFC7050). No classical DNS query was necessary, Cheshire said.

The time delay with regard to the special name topic in DNSOP might certainly result in several special names being approved before the WG finally makes up its mind on the controversy.

Long list of new proposals

From the newly presented proposals in Seoul, only one received overly positive feedback and put on the go ahead list: it is a proposal to standardize an exchange format for large DNS packet captures (currently PCAP and PCAP-NG are used in practice). During the Seoul session, most commentators said this would be extremely useful.

All other new work proposed for WG adoption was put in the "discuss more before adoption" basket with some debate sparked over a document for "delegation requirements" (useful, but difficult to find consensus on, said Peter Koch, DENIC).

With regard to the DNSSEC automated zone transfers proposal from Matthew Pounsett form Rightside, which also shows up in the DNSOP proposed document list, a debate is ongoing on how to deal with patent applications (see in RegExt below). During the DNSOP WG two IPR claims were mentioned, one with regard to the zone transfers of DNSSEC secured zones and another one with regard to Dickinson's draft.

A DNS related Bar Bof explored possibilities to specify http transport for DNS.

DPRIVE WG: Profiles and Padding

The DPRIVE WG can be expected to re-charter soon to add new milestones to its charter, namely encrypting DNS queries-answers between resolver and authoritative name servers (see highlights). Two topics still on the agenda are two documents to complete the DNS over TLS work completed so far.

One is the DNS-over-TLS profile document, which contains profiles from very strict (encrypted and authorized only) to lax security (fall-back to cleartext in order to receive DNS packets). The profile document describes a strict and an opportunistic profile for DNS over TLS, with strict asking for authenticated encryption, while opportunistic would allow for unauthenticated encryption. There was rough consensus in Seoul that a decision on how failures for authenticated encryption (in the strict profile) should be up to the local site. A "jedi-like profile"/ obligatory "hard fail" should not be part of the profile document. The document is in WG last call

and will be sent off to the IESG soon.

The second topic concerned a follow-up document to RFC 7830, the EDNS Padding option that allows clients and servers to add a variable number of bytes to encrypted DNS messages to prevent fingerprinting from message sizes. With the new follow-up document 7830 author Alexander Mayrhofer (nic.at) delivers an overview of five different padding strategies. For all five:

- zero padding,
- fixed length padding,
- block length padding (where padded length is a multiple of a chosen block length),
- random length padding (message padded with a random amount of padding)
- random block length padding (random choice between block length's plus block length padding)

Mayrhofer discussed advantages and disadvantages. The document can be expected to become a WG document.

DomainBundle BoF: Cloning for IDN variants or multi-dimensional identities

In a BoF chaired by Jim Galvin (Afilias) and Ning Kong (CNNIC) another attempt was made to gather support for specifying the bundling of variants of a domain name in one registration. The need for bundling variants in special IDN character sets was presented for:

- simplified and traditional Chinese (互联网中心.cn = 互聯網中心.cn)
- Greek Sigma (Νίκος.gr (xn--kxawhkp.gr) = ΝΙΚΟΣ.gr (xn--uxachku.gr) = Νικος.gr (xn--uxachkp.gr)
- Czech diacritics (á č ď é ě í ň ó ř š ť ú ů ý ž)

Beside the bundling of two (or few) variants from one-character set, there are also ideas to make bigger bundles, for example for name variants with or without dashes. Richard Merdinger from GoDaddy introduced the idea that for the future, there was a need to bind multi-dimensional identities spanning multiple TLDs (instead of ricksrestaurant.com, in the future also include ricks.restraurant).

So far, solutions used include parallel DNS (same name servers for all names in bundle), DNAME and CNAME, according to an overview given by John Levine. Levine listed the various problems too, and new possible solutions:

* BNAME, basically a new resource record type, combining CNAME plus DNAME

* Clone, authoritative server synthesizes name1, name2...records, clone-aware synthesizes, too

A different issue is addressed by Ted Hardies ArcPointing. It tries to open ways to allow different resolution systems (special names, for example) to sit side by side, with a method to indicate the context of resolution for a name. Hardie's document proposes "a registry for such alternative resolution contexts as well as a set of pointer resource record types useful for allowing conformant resolvers which query for the name in the DNS to be redirected to the appropriate alternative resolution context."

Hardie himself was one of the vocal critics of the DomainBundle BoF, warning that, for example, in Chinese, simplified variants could be the result of conversion from different traditional characters. The question therefore would be who would be able to receive the bundle or would all these be bundled together? 发 = 髮 fà, but also 发 = 發 fā

Participants of the BoF pointed to the recent closure of the Dbound WG that had addressed a broader set of the same problem area, but had been unable to find consensus. It was questionable therefore that DomainBundle could succeed.

What got some nods in the end was the proposal rather to experiment with the clone proposal at individual registries level. The initiators intend to continue the discussion and clarify the use cases. Alex Mayrhofer, nic.at, proposed instead to experiment with cloning, for example, before trying to standardize.

REGEXT

A proposal to create a new object especially for resellers was heavily discussed in the REGEXT WG. The proposal from Linlin Zhou (CNNIC) for a reseller object in WHOIS or RDAP shall enable "enhanced reseller features at the registry level" and allow for tracking of reseller financial information, reseller security and reseller reporting.

Most participants warned against the complexity. Alex Mayrhofer (nic.at) questioned why a reseller should "suddenly be policing things on the second level". Mentioning resellers could be an option, said Rick Salz from GoDaddy, but he would not support creating a full object as it would escalate the role of resellers. By introducing a full reseller object, the group would make policy through technology, he warned. The debate illustrates that, as has been observed several times, that the REGEXT WG works on a very fine line between policy and technology.

Once more, Jim Galvin, who has been chairing the meetings of the group virtually single-handedly (since Co-Chair Antoine Verschueren has not participated in person for quite some meetings now), called on the group to do more reviews on the many extension documents channelled through the group. 13 active WG documents are currently in the production pipeline and a long list of new ones are in the waiting line, including several new ones on RDAP (see here).

For the first time, the Co-Chair of the Human Rights Considerations in Protocols Research Group participated, questioning some documents for lack of privacy considerations.

The WG itself also briefly considered how to experiment with interim work and potential design team work (during the IETF meetings).

IAB, IAOC and Trust News

Appointments to various ICANN-related functions

The IANA transition created a Root Zone Evolution Review Committee (RZERC), and the IAB appointed Jim Reid as their representative. Tim Wicinski was re-appointed to the ICANN NomCom. Paul Wouters was re-appointed as a liaison to the technical liaison group. The IAB appoints one IAOC member for a two-year term. The IAB is appointing members for the IAOC: as a consequent, appointees are also Trustees of the IETF Trust.

IANA IPR

Trust Chair Tobias Gondrom (Huawei) gave an update on the transfers of IANA IPR to the IETF Trust. IANA IPR has been assigned to the Trust and there are now several agreements in place to administer the new arrangement: community agreements between Trust and the operational communities (IETF, RIRs and ICANN) and three license agreements from Trust to ICANN for sub-licensing to the PTI.

ICANN has submitted the documentation to the USPTO, transferring the trademarks to the Trust and the Trust is in discussions with a Registrar that is willing to enforce the domain registrar requirements in the license agreements "following concurrence of the agreement with ICANN the IANA domains will be transferred to the Trust".

The next IETF will be taking place in Chicago on 26-31 March 2017.

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