Report on

IETF95

Buenos Aires

3-8 April 2016

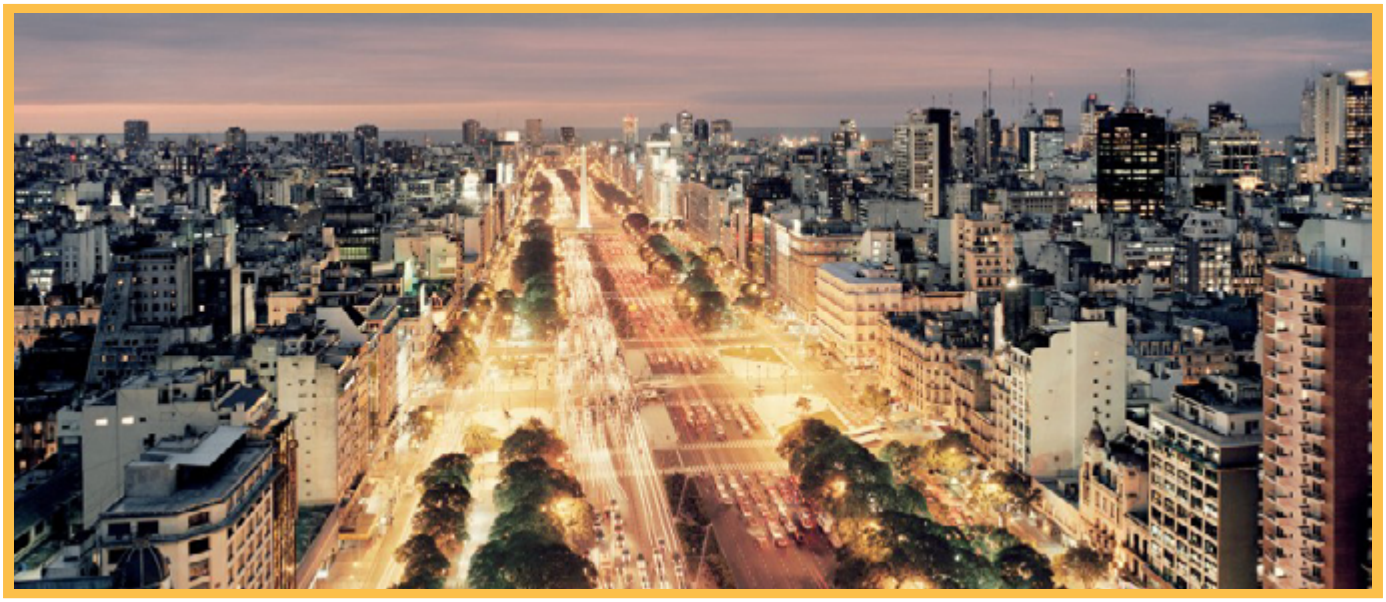

Table of Contents

IETF from 1,000 feet – Executive Summary

The IETF in Buenos Aires saw a record number of birds-of-a-feather (BoF) meetings and, to some extent, already implemented what IETF Chair Jari Arkko has put out as a proposal for future IETF work. Apart from traditional Working Group meetings, Buenos Aires seems to have given a preview of a potential new function of the IETF becoming a convener of communities around vertical topics. Example sessions in Buenos Aires were the ones on "Intelligent Transport Systems" (ITS), on Mobile Networks and Encryption (Accord), and in the field of IoT (Thing-to-Thing Research Group).

IETF as Convener?

On Internet of Things (IoT), the Internet Architecture Board had managed to bring together a number of companies, organizations and standardization bodies that had not met in that way before, Dave Thaler from Microsoft reported during the plenary in Buenos Aires. Together with Arkko's call for the IETF to cooperate more intensively with the open source community in the "Trends" document, these developments could change the face of the IETF into a more outward-looking organization that also explores new working mechanisms.

The push possibly results from a concern that the IETF could become an old guardian of an old protocol world that takes a backseat to ongoing work on applications.

"What's in a name" dispute ongoing

For the domain name community, the handling of special names, developments around RDAP and conversations in the meeting halls on the ZSK development were most interesting.

For one, an ARCING BoF saw broad consensus on the need for a new document that will define the concept of "domain name". Domain names were only one subset of internet names and intier systems, and there was no clear definition of domain names in the respective RFC standard documents, explained Ed Lewis of ICANN. Lewis took the lead in locating "domain names" in the various IETF RFCs and will now develop the "definition" draft document.

Ted Hardie explained the need for context as an indicator for which system an internet name had to be queried/resolved in. Announcing resolution context for a given name through a tag or prefix (as in IDN resolution) could help to distinguish between DNS names and other internet names. The ARCING work is a follow-up and in reaction to the requests for another 11 top-level domains under the IETF "special names" procedure (RFC 6761).

Interestingly, contrary to the attempt to somehow avoid a potential conflict between ICANN gTLD procedures and the IETF special names delegations, there are at the same time requests in Homenet for an additional IETF allocation for homenet special-use TLDs like .homenet and/or .hnsd. Meanwhile the requests for special domains originally made jointly with the one for the allocated .onion (.gnu, .bit, .exit, etc.) is not addressed at the moment. The respective document was taken off the DNSOP WG list of current and expired documents.

RDAP using OpenID for federated authentication and layered access

The RegEXT (formerly EPPEXT) WG, besides short presentations on additional extensions sought after by VeriSign, CNNIC and SIDR, was presented with VeriSign work on the use of OpenID as a method to authenticate users for a layered access to the Whois follow-up system. Until June, users can experience the difference between authenticated and non-authenticated access to VeriSign Whois information. Authentication providers for the experiment are Gmail, Hotmail, CZ.NIC and Verisign Labs.

Rolling of the Root and Zone Signing Keys

According to information from participants to the OARC meeting that preceded the IETF meeting, the signing of the Zone Signing Key was announced by VeriSign. The announcement of the time table for the rollover of the Key Signing Key also was reported to be imminent by somebody close to the process.

Transport Layer Security (TLS) Protocol Version 1.3 ready soon

TLS 1.3 made a big step forward by arriving to consensus over key handling, one participant said after the session. The follow-up version of MPTCP is also expected to be finalized shortly.

IETF in Latin America for the first time and debates on “venue selection”

The IETF meeting in Buenos Aires was the first in Latin America and although well-attended was below the budgeted number of participants. It also saw a heated debate in plenary over criteria for venue selection, with participants focusing very much on the issue of criminalization of gay relations in Singapore following a complaint by Ted Hardie (Google). Singapore is the chosen venue for the IETF100 meeting.

Highlights

6761: the fight is on

The fall-out from the allocation of .onion as a special use top-level domain (SUTLD) has resulted in close to "religious" debates over the future of the special name registry, RFC 6761. After a design team including Alain Durand (ICANN), Peter Koch (DENIC), Joe Ably (Dyn), with Warren Kumari (Google) as an addition, prepared a first problem statement, an alternative was presented by Ted Lemon (Nominum).

Both drafts in essence address the problems with the 6761 allocation mechanism experienced during the .onion decision: lack of clarity on which body is to make the decision, conditions for the allocation, etc. Both also mentioned deprecation/discontinuing 6761 as one of the potential mitigation options.

Yet the older proposal from the design team is more sceptic of 6761 and the way it came to pass. Lemon's counter proposal now questions the "purity" arguments and seems more eager to find a way forward on 6761, not the least perhaps to allow for a .homenet special TLD in the Homenet WG in the near future.

Meanwhile, another short draft by George Michaelson (APNIC) is simply requesting to shut down 6761.

Two competing drafts on the special name problem

Lemon is heavily criticizing Durand, Koch, Ably and Kumari for "mixing different problems" in their problem statement. In his opinion the following types of SUTLDs are perfectly in line with IETF consensus for now:

- *names that should be resolved using the DNS protocol with no special handling*
- *names that should be resolved using the DNS protocol, but require special handling in the resolver*
- *TLDs like .local which act as a signal to indicate that the local stub resolver should use a non-DNS protocol for name resolution*

Against the background of the IAB technical comments on the unique DNS root, RFC 6761 itself and several other documents (.onion document, local names, IANA name collision study and some more), Lemon addresses three kind of "architectural purity" concepts, several "squatting/legal problems" and potential security questions.

On purity, the equation of domains with DNS names is implicitly rejected (if DNS was the only protocol to be used for resolving domain names, .local, .onion and etc/hosts would need to be depreciated, Lemon states). If universal uniqueness (every domain means the same everywhere) would be made a condition, mDNS would have to be deprecated. Lemon considers the protocol switch as a potential future mechanism required from resolvers (to avoid leakage and privacy problems). Potential competition or legal issues from landrush/squatting and the of risk being sued for routing around ICANN could be addressed by clarifications of use cases for the IETF SUTLD and application process with a potential cooperation with ICANN in an effort to have help in vetting individual names consistently. Allowing experimental use SUTLDs under a .alt could allow to make allocations for a limited time. This could prevent committing strings indefinitely to projects that might later fail.

Included in Lemon's version of the problem statement is ongoing work he is doing on names architecture in homenet, which looks into a potential additional carve -out of special domains (.homenet, .hndb?).

"It is not the case that all special-use TLDs can be expected to be non-DNS names; for example, the Homenet Working Group is likely to propose the use of a special-use TLD for use on the homenet in cases where the homenet does not have its own globally unique name allocated. This would nevertheless be resolved in the same way as an RFC1918 reverse query, by sending a query to the local resolver using the DNS protocol."

Lemon's "Homenet Naming Database" (HNDB) aims at allowing for advertising homenet services inside or beyond the local network – while avoiding name collisions and providing the delegation and management of the homenet names in an automatized fashion to the dummie user (see homenet WG below).

On 24 April, .home was declared as the default TLD for homenet queries in the newly adopted RFC 7788. But interestingly, a delegation by IANA was not mentioned in the final RFC. This makes the status of .home unclear – it is not allocated and not applied for through 6761. At the same time there could be potential complaints from those who have applied for .home in the ICANN new gTLD process (applicants include Google, TLD Holdings, Uniregistry and more).

The earlier problem statement by Durand, Koch, Abley and Kumari starts from the same notion of "lack of clarity" of 6761 and problems arising from it, and looks at some problems in more detail.

It addresses architectural, technical and organizational issues. Architecturally the basic question to answer was "are we talking about one name space with different resolution protocols or independent name spaces?". It could indicate switches to other name spaces like in .onion, a switch to another protocol like mdns or Tor or call for local use of a name as in .local or .home. Neither of the switch practices was sufficiently documented, the authors state.

One important point in the technical issues are how to manage expecations of the new special TLD managers. Reservation of a special TLD domain by the IETF for example does not guarantee that DNS queries will not be sent over the Internet, warns the draft. Leakage of private data was therefore one problem to address.

With regard to organizational questions, sloppiness of process is identified as the culprit: the first successful 6761-candidate, RFC 6762 (.local), had only be passed as an individual submission with no documented discussion. .onion (RFC7686) on the other hand was discussed in DNSOP and the IESG before approval – and resulted in the extensive debate over SUTLDs now. The authors also point to the adoption of RFC 6761 as an individual submission RFC in the first place. Critics either think the IETF/IESG should have better vetted the SUTLD base spec or blame ICANN for not calling for more discussion.

The authors of the original problem statement finally point to the ICANN application and vetting process, one the IETF would have difficulties to emulate, they think.

The inclusion of .home in RFC 7788 without even considering a 6761 application procedure could add to the ADPK-JA-argument that the IETF is not behaving consistently.

Debate and next steps

Two main lines of argumentation can be distilled from the controversial debate now underway in DNSOP. One considers the opening of the SUTLD application process as a chance to keep the IETF open to the growing number of application developers who develop new alternative naming systems and allowing them to interface with the DNS which could take more of a "back seat", according to IAB Chair Andrew Sullivan. The core homenet/local camp (Nominum, Apple, e.a.) even warned that "there was no law that we use DNS" (Stewart Cheshire) implicating that big application providers certainly were in the driver's seat. "Consistent IETF behavior" and avoidance of a clash with ICANN (MoU RFC2860) is the main argument from the "purist" camp. David Conrad stated that ICANN was not taking an official position on the dispute (despite the fact that Durand is author of the more critical problem statement). With ICANN not trying to give their interpretation of RFC2860, the IETF community has to slug it out.

While several participants applauded Lemon's draft as being easier to read and more clear, DNSOP Co-Chairs Suzanne Woolf and Tim Wiczinski called the Design Group closed after the meeting and asked for comments on both problem statement drafts with the intention to have one draft picked before the Berlin meeting.

One additional draft to help with "switching" was posted by Stephane Bortzmeyer (AFNIC). The draft recommends a DNAME records to allow for NXDomain answers from the SUTLDs.

Coexistance of Name Resolutions – Magical context switch too hard to find

The dispute about the allocation of special names in the Domain Name System by the IETF (according to RFC 6761) started a broader discussion about "Alternative Routing Contexts for Internet Naming" with a first BoF held in Buenos Aires. It is still unclear if there will be another WG creating BoF at the next IETF or if the drafts coming out of ARCING will be handed to other WGs (for example the DNSOP WG).

Being inclusive in the DNS or not – and who takes the lead?

At the first ARCING BoF in Buenos Aires, Suzanne Woolf (DNSOP and ARCING Co-Chair) drew a line between the ARCING discussion and the follow-up steps for the 10 additional applications for special names that were out of scope. At this time the drafts for .bit, .gnu, .i2p, etc. have fallen off the document list of the DNSOP WG. Woolf said to the CENTR reporter that she was unable at this moment to talk about the next steps for the .onion co-applicants.

Instead of discussing how to proceed with the 6761 applications, which Woolf said was "out of scope", ARCING intends to address the larger issue of alternative identifier systems on the net.

IAB Chair Andrew Sullivan noted that on the overall goal, the IETF had to be concerned with interoperability given that people were designing new naming systems faster than the IETF could react. The IETF itself has developed or at least

accepted several naming systems already with multicast DNS, the handle system or onion routing. The desire to create new identifier systems would not go away, but instead continue to grow, so "we're going to need to have a way to make these identifiers useful", Sullivan said.

Steward Cheshire (Apple), author of RFC 6761 and of a proposal to delegate .home as another special TLD in the Homenet WG underlined the IETF had no option not to follow-up on industry initiatives. It is certainly interesting to compare the .home and .bit, .gnu, .exit and the .home application status quo which are good examples of how larger companies have better chances of influencing architectural discussions.

Definition for domain names and a resolution switch mechanism

The BoF discussed three documents, including the one on a definition of domain names, which received broad support. Author Ed Lewis (ICANN) observed the lack of a "formal and written" definition and explains the historical development of the concept derived form the host name system and heavily influenced by SMTP. Lewis presents definitions from Diestel and Lyman Chapin in his draft which can be expected to be discussed further in either a new WG or the DNSOP.

Ted Hardie's draft focuses on the issue of context sensitivity and the broad idea that it would be helpful to know where a name presented had to be resolved. The ideas on how to allow to present a name together with a hint to where it should be resolved were briefly touched upon a DNS subtree solution or a string delimiter as in the IDN resolution context (RFC5891). The .alt-subtree while requested by some during the session as a quick solution for the outstanding applications for .gnu etc. is seen by some as exiling them into a "ghetto", Hardie said. The string delimiter would "be used to construct faceted URI schemes, one aspect of which contained the usual protocol indicator and the other the resolution context."

The two other options, just continuing as is with 6761 special names applications or fixing either the number of gTLDs or the number of future resolution contexts, were not regarded as viable by Hardie. In the draft Hardie concludes:

"There are clearly trade-offs among the available alternatives, as each has its own drawbacks as an indicator of resolution context. Given, however, that the existence of multiple signals could generate even further interoperability issues and operational concerns, the creation of multiple signals is undesirable. Any system which allows Internet names from alternate resolution contexts to be used in common protocol systems can likely be made to work, provided its drawbacks are accounted for and mitigated appropriately."

Two more systems for context switching were referenced in Buenos Aires: the naming system switch developed for unix systems (nsswitch.conf) and DNS classes. For the latter however, IAB Chair Andrew Sullivan proposes to deprecate the respective RFC.

The third draft explores properties of an ideal naming service. Author Brian Trammell (ETH Zurich, IAB member) listed the following properties:

- *federation, unity, transparency and revocability of names and their uniqueness,*
- *the authenticity of delegation and response (including the one for negative answers),*
- *dynamic consistency and support for explicit inconsistency and*
- *explicit support for trade-offs among latency, efficiency, traceability, consistency.*

Trammell noted that an ideal system would in fact look rather similar to the DNS. Yet enhancements were possible for example with regard to authenticity by requirements to have signatures as mandatory. This would make the only slowly deployed DNSSEC unnecessary, according to Trammell.

A DNAME record for all special use TLDs has been prepared to solve the "explicit inconsistency" issue by Stephane Bortzmeyer.

Rolling the Keys

While not a topic on the official IETF agenda, there was some news on rolling both the zone signing key and the root signing key.

According to one source close to the process the time table for rolling the root signing key will be announced shortly. The final proposals from the design group had been sent to ICANN recently. After a public comment period ICANN has decided to go ahead and prepare for a roll later this year, the source said. The announcement was said to come out any day now. It is still to be seen how ICANN will prepare for the necessary communication.

Meanwhile during the OARC meeting that preceded the IETF Geoff Huston gave another presentation about issues to be faced by the roll, namely the problem that validating resolvers will be cut off from resolving signed domains. One major problem is that there is no way to measure the level of failure in the DNS network globally.

In another presentation VeriSign announced its plans to roll the ZSK at the end of the month. Acting as the root zone maintainer under contract with the US government for the time being, VeriSign had cleared the step with the NTIA and ICANN, according to one source. The ZSK being “hidden” under the KSK signature will not shed any light on the KSK roll, experts say. An unanswered question also is if VeriSign does consider to roll the .com and .net keys soon. Both have not been rolled before.

Many experts think avoiding the roll for so long (or not putting roll over time tables in the IETF specifications before the root was signed) was bad. It is expected that the day of the rollover will see highly visible failure rates.

Working Groups, BoFs and Plenary Bits

REGEXT

REGEXT WG remains notable in two regards – it is not working like your typical IETF WG but is approving extensions to registration protocols (EPP and RDAP are the next step). Not only are these highly registry-registrar specific issues, but sometimes it is unclear what comes first between registrar and registry policy at ICANN (then to be "approved") by the REGEXT WG or technical speficiations that could be sent back to ICANN as "IETF standards". In Buenos Aires the WG decided for now to wait for the policy development at ICANN with regard to an RDAP proposal presented by ICANN staff.

The re-chartered and re-named EPPEXT WG, now REGEXT WG, has a long list of documents as milestones in its new charter. The listed registry extensions do not need to be put to the test of "approval as a WG document or not". They automatically have become WG documents. The timetable stretches from February 2016 to June 2017 (see below), with the last topic of third party providers' access to registrar/registry databases. On the latter Olafur Gudmundsson (Cloudflare) did propose a concrete document (Third Party DNS operator to Registrars/Registries Protocol), that is not yet in the milestones.

Heavy users of the process for the time being remain VeriSign and CNNIC. ICANN is also a frequent presenter. Other registries with one or two proposals in the current milestone list are ICANN, CIRA and CentralNIC.

CentralNic's "fee extensions" draft briefly discussed three options with regard to objects to include – either keep the existing syntax and have no relationship between objects listed in the main body (A), include a single element, the fee information calculated for each object in the main body (B) or allow for multiple elements and the fee information calculated for each combination of object and extension (C). Pointing to possible DDOS on fee extension information, Alexander Mayrhofer (nic.at) asked for a limited C option that would allow registrars to check for create and renew pricing in one single command, for one "period" element per "command".

Other drafts briefly discussed in Buenos Aires were three drafts of CNNIC, one from the milestone list on reseller extension (draft-zhou-eppext-reseller-mapping) and two new ones on the verification for EPP contact and on domain name mapping. Two new drafts were also presented from ICANN, both related to RDAP (draft-lozano-rdap-nameservers-sharing-name, draft-lozano-ietf-eppext-registrar-expiration-date). But participants and Chairs both decided that it was premature to adopt the expiration date draft because ICANN still had to decide on the respective policies.

Federated autentication for layered access to RDAP

RDAP is already the topic of several drafts. In a highly interesting presentation, Scott Hollenbeck from VeriSign described an ongoing pilot project to use OpenID for federated authentication to allow for layered access to the future RDAP database, which is planned by ICANN (Verisign and others?) to be the follow-up protocol for Whois. For now, using credentials from Yahoo, Google, CZNic and VeriSign allows getting a larger set of RDAP data. At the same time, Hollenbeck underlined that VeriSign was not interested in becoming an ID/credentials provider.

June 2017	Submit for publication an informational RFC with requirements for a registration protocol for third-party DNS providers
February 2017	Submit for publication “CIRA IDN EPP Extension” draft-wilcox-cira-idn-eppext
February 2017	Submit for publication “EPP IDN Table Mapping” draft-gould-idn-table
February 2017	Submit for publication draft-ietf-eppext-idnmap
October 2016	Submit for publication “Allocation Token Extension for EPP” draft-gould-allocation-token
October 2016	Submit for publication “Change Poll Extension for EPP” draft-gould-change-poll
June 2016	Submit for publication “EPP Domain Name Mapping Extension for Bundling Registration” draft-kong-eppext-bundling-registration
June 2016	Submit for publication “EPP China Name Verification Mapping” draft-xie-eppext-nv-mapping
June 2016	Submit for publication “Verification Code Extension for EPP” draft-gould-eppext-verificationcode
April 2016	Submit for publication “EPP Reseller Mapping” draft-zhou-eppext-reseller-mapping
April 2016	Submit for publication “Reseller Extension for EPP” draft-zhou-eppext-reseller
April 2016	Submit for publication “Registry Fee Extension for EPP” draft-brown-epp-fees
April 2016	Submit for publication “EPP and RDAP Status Mapping” draft-gould-epp-rdap-status-mapping
March 2016	Submit for publication “TMCH functional specifications” draft-ietf-eppext-tmch-func-spec
March 2016	Submit for publication “Launch Phase Mapping for EPP” draft-ietf-eppext-launchphase

DNSOP WG: DNS classes under discussion – Domain delegation checks and attempts at hardening the system

Beside the big topics, the conflict around special use TLDs and the future of the IETF’s approach to naming (see highlights above) a long list of active drafts and some new proposals were reviewed in two working group sessions in Buenos Aires. A controversial proposal that was discussed was the idea of preventing “classes” from being used in the DNS in the future and as a consequence shutting down the respective IANA registry. Several drafts discussed in BA try to help countering DDoS and amplification attacks in one way or another, for example the NSEC/NSEC3 aggressive negative caching, clarification of NXDomain answers and the limitation or even refusal to answer any-requests.

Three draft proposals by Cloudflare (and partners) are now in or close to WG last call: problems for DNSSEC deployments resulting from non-compliant infrastructure (Roadblock Avoidance), the limitation of the “any”-answer from resolvers (or refuse any-answers completely as proposed in a second document) in an effort to avoid large answers abuse for amplification, and the introduction of a “delete DS record” option and an “enable DNSSEC validation” option for CDS and CDSKey. The latter shall allow the signalling of a DNSSEC validation stop or start from the user to his DS provider.

Next up for the WG last call according to DNSOP Co-Chair Tim Wicinski is a clarification of NXDOMAIN answers that will set as a rule that a caching DNS resolver should stop searching in its cache as soon as he encounters a cached NXDOMAIN answer.

Another attempt to stop DDoS attacks on DNS servers is the so-called NSEC/NSEC3 aggressive negative caching. It relaxes the condition of exact matches for negative caching (a.example.com, b.example.com). According to the draft proposal from K. Fujiwara (JPRS) and A. Kato (Keio/WIDE), “when a query name has a matching NSEC or NSEC3 resource record in the cache and there is no wildcard in the zone which the query name belongs to, a full resolver is allowed to respond with NXDOMAIN error immediately.” And that matching procedure may be applied to all ancestor domain names of the query name. An issue to consider is that newly delegated domain names in a zone might have to wait for the expiry negative information to be effective. Another alternative approach (cheese-shop) will not be further reviewed by the WG.

Based on the work for a new tool to check the delegation of zones, the Zonemaster prepared by experts at IIS and AFNIC, Patrik Wallström (IIS) and Jakob Schlyter (Kirei) proposed a document listing DNS delegation requirements. The document contains all requirements for a “well-behaved DNS delegation of a domain name”, and at the same time can be a base for a set of delegation tests. Some discussion in the WG about what the scope of the document should

be were highly welcomed there – only as an informational document, a best practice or normative document. If it was normative, some participants requested a thorough check for consistency with normative behaviour for DNS delegation scattered in various RFCs. There was also a discussion with regard to scope in terms of which set of requirements would be broadly acceptable – only a minimum set or one that tries to be exhaustive.

Another issue discussed in Buenos Aires was the possible closure of the DNS Class registry at IANA. Andrew Sullivan (Dyn/IAB Chair) proposes to discontinue the registry because it has not worked well. Opinions diverged, with many experts agreeing, but an equally large group of members arguing against closure, because the problem was not the DNS classes: "the problem is badly written software." (Wes Hardacker, Parsons). Hardacker and others like Mark Andrews recommended to consider writing "operational guidelines that different classes need to be delegated in parallel." Some problems could be solved by using classes.

Also under discussion in Buenos Aires were the "EDNS Key Tag Option" and a DNSSEC algorithm update. The EDNS key tag option is intended to allow "validating end-system resolvers to signal to a server which keys are referenced in their chain-of-trust" (see draft). The extensions shall allow zone administrators to monitor the progress of rollovers in a DNSSEC-signed zone.

The DNSOP WG again had two sessions in Buenos Aires and has a considerable list of new RFCs just published or close to be published. Between IETF94 and 95 a record number of 5 RFCs have been published:

- draft-ietf-dnsop-root-loopback RFC7706
- draft-ietf-dnsop-dns-terminology RFC7719
- draft-ietf-dnsop-5966bis RFC7766
- draft-ietf-dnsop-qname-minimisation RFC7816
- draft-ietf-dnsop-edns-tcp-keepalive RFC7828

Three drafts are in the RFC editor queue:

- draft-ietf-dnsop-rfc6598-rfc6303
- draft-ietf-dnsop-edns-chain-query
- draft-ietf-dnsop-edns-client-subnet

DPRIVE – Looking beyond opportunistic privacy

With the DPRIVE base specification "DNS over TLS" done and the additional DNS over DTLS approaching, WG last call is setting its eyes on next steps, especially on the authentication of recursive servers and on usage profiles (strict, opportunistic and no privacy options).

Authentication had been decided by the WG to be dealt with in an additional specific document which is now reviewed by the group. It is intended to add to the opportunistic privacy profile and the stricter out-of-band key pinned privacy profile described in the original DNS-over-TLS base spec. Authentication of the server direct configuration of the recursive server and DHCP are considered with verification via either X.509 or DANE being obligatory.

Discussion in BA mainly focused on how users would deal with various profiles and fall-back mechanisms and how much opportunistic privacy was broken. Christian Huitema (Microsoft) on the other hand argued that there was a need to allow users easy switching for usage scenarios (for example when querying a DNS server from an office network or an Internet café).

The question was also briefly posed whether the WG should wait for more implementation before proceeding. A draft on measurement of DNS over TLS-deployments is also prepared.

Potential security pitfalls of zero roundtrip handshakes (which can expected to be in the new TLS 1.3 specification) are:

- that the data was not forward secret
- no guarantees for non-replay between connections was given
- in case of a compromised server an attacker could tamper with the 0-RTT data without detection

Homenet

The Homenet WG is moving ahead in several aspects, having decided on Babel as a routing protocol and moving ahead on the naming architecture. A considerable amount of time during the Buenos Aires session was dedicated to naming.

The homenet naming architecture shall address all issues of naming:

- Provisioning of a name space in which names can be published and services advertised
- Associating a name within that name space to the set of IP addresses on which a host is reachable
- Advertising services available on the local network and associating those services with names published in the name space
- Distribution of names published in that name space to servers that can be queried in order to resolve names
- Correct advertisement of name servers that can be queried in order to resolve name
- Timely removal of published names when they are no longer in use

Considerations are key for the the homenet naming architecture, namely provision of a kind of a hybrid system: some homenet services should be only advertised to users on the homenet, others could be advertised more broadly to users outside of the home network. Services are expected to be advertised off-network for some links and not others.

Automatization should also help to avoid mistakes, according to Lemon's proposal. *"All of the operations mentioned here must reliably function automatically, without any user intervention or debugging. Even to the extent that users may provide input on policy, such as whether a service should or should not be advertised outside of the home, the user must be able to safely provide such input without having a correct mental model of how naming and service discovery work, and without being able to reason about security in a nuanced way."*

Naming conflicts should also be resolved automatically. Multihoming shall be supported "and therefore support for multiple provisioning domains [6] is required to deal with situations where the DNS may give a different answer depending on whether caching resolvers at one ISP or another are queried."

Interestingly, just days after the IETF95 meeting, RFC 7788 was announced, which settled on .home as the string for names in .homenet, without considering potential collisions or the whole process of special names allocations by the IETF (not to speak of an ICANN application for .home).

Sunsetting IPv4 – madness or good signal?

A fierce debate developed in the IPv4 Sunsetting WG over actually calling for an end of life announcment for IPv4 by the IETF. Lee Howard (Time Warner) proposed the draft on declaring "IPv4 historic". While some warned against what they said was sheer madness, after some debate there were quite a few experts who said that a strong signal by the IETF to get on with IPv6 would be welcome.

According to RFC 2026 a specification for which there is a new version becomes "historic". With IPv6 around and growing, Howard said to this reporter that in the US, it was not unrealistic to expect that 80% of IP addresses would be IPv6 in the next 5 years. He stated that the IETF should stop wasting work on transitioning technologies for which the Sunsetting IPv4 WG was originally created after too many proposals were brought to v6Ops. The consequences of declaring IPv4 historic would not only mean no further updates for IPv4 – and in fact not even bug fixes.

Producing bullshit was no task for a serious standardization body, warned Geoff Huston (APNIC): other standards bodies might just like to take over from the IETF. A stop of v4/v6 dual stack developments would be an invitation for them. What would be even worse would be "if we get creative again", Huston ranted. He pointed to the consequences of the IETF rejecting the development of network address translation standards. For the time being 50 to 60 million devices would be pressed into the dual stack system, he said.

Engineers of Facebook, Apple and Microsoft, speaking as usual in their individual capacity, were more welcoming. By continuing to fix issues of IPv4, the IETF created the expectation that IPv4 would be around for a long time. There was much applause for an announcement like the one from Apple, who decided that IOS9 would be IPv6 only.

Ruediger Volk (Deutsche Telekom) and Fred Baker (Cisco) offered an alternative for the strong signal: instead of declaring IPv4 historic, the IETF could announce the "end of engineering" or "end of life" status. End of engineering equates to still available, but no new features could be expected. So far, interesting new features of IPv6 were often copied into IPv4, preventing the newer protocol from distinguishing itself.

Howard accepted that his "historic" proposal might be premature, but countered arguments that regulators might step in when a provider would not offer IPv4 anymore. It could very well be a competition issue¨: some newer competitors had to spend a lot on transition while others were just comfortably sitting on rich old address reserves.

UTA and secure email – Competing with DANE

"Strict transport security" is an attempt of several large email providers (Google, Yahoo, Comcast, Microsoft, Linkedin and 1&1) to marry TLS with SMTP in a way that would make it a little less "opportunistic". Preventing downgrade attacks mail servers will publish policies allowing the distributing mail server to check if TLS is offered by the receiving mail server. Some providers would even be satisfied with the mere option of reporting for a start, said Google's Mark Risher in Buenos Aires. The policies could be stored as new resource record or simple text file in the DNS.

The approach discussed intensively in the UTA WG in Buenos Aires is competing with DANE and intends to do so in the way that the lack of DNSSEC deployment is to be addressed according to Risher. Instead of waiting for DNSSEC deployment, SMTP STS was proposed, he said. Yet the backing up of the policy checks via DNS (and could we have that secured, please?) calls the don't wait for DANE into question again. As an alternative, the authentication via WebPKI was discussed to allow for a secure mail policy check. While discussion is ongoing, it looks like something that will proceed quickly in the draft.

Yet another approach discussed looks at Mail User Agent Strict Transparency which shall allow users to make informed decisions about the level of security they are willing to accept. MUA STS as it shall be named instead of the former "deep" shall provide for the most secure variant (TLS between MUA client and mail server) as the default. Coloring would show if the security level is not kept up with for a given connection. The user will have to touch the default for lowering the security level. A mere "send anyway" from the secure default would not be accepted according to the proposed standard. Only a "work offline" or "try again later" could be offered, author Chris Newman said.

Instead of StartTLS MUA STS shall use implicit TLS which allows to negotiate TLS for the first connection (by signaling over a special port).

Controversial Venue Selection

Should the IETF consider additional selection criteria from the human rights realm or should the IETF community just be given more of a say in venue selection in the first place? IETF leadership bashed over venue selection.

Selection of IETF venues became an issue of a major debate in Buenos Aires after Google engineer and IETF long-time contributor Ted Hardie took the mic for a fierce rant over the venue chosen for IETF100. Singapore, said Hardie, was still criminalizing Gay relations which would make bringing families to the meeting impossible for LGBT IETF participants.

Addressing members of the IAOC meeting committee (who often do bring their families) Hardie requested them to "have the decency not to bring their families to IETF100. Hardie was supported by many community members in the plenary discussion that followed. Several community members were enraged when Tobias Gondrom, IETF Trust Chair and IAOC member, tried to explain that there were so many different criteria to be balanced.

Even before the plenary, there had been considerable unrest about the selection processes for several reasons. A dedicated BoF on venue selection offered lengthy explanations of criteria currently addressed in the IAOC selection procedure. The criteria, besides the sheer space needs (recent IETF meetings had up to 1,300 attendees) and the preference for the "one roof" policy (i.e. at least a third of the participants can find hotel rooms at the conference venue) include sponsorship questions and other issues. Even issues like "supermarket close by" to allow for cheaper self-catering are on the list. A draft informational RFC can be found here.

From a political point of view the existence of openness (for the meeting and the network) are the core principles. Problems with visa issuing had resulted in a clear decline to choose US venues in recent years, a fact that from time to time is questioned by some US participants (who still are the majority).

The Buenos Aires meeting, pushed for by IETF Chair Jari Arkko, while implementing the IETF diversity plans, was reported to be a very expensive meeting (with a minus budget-wise).

Next steps on venue selection include a decision by the IAOC on renegotiating the IETF100 venue, more debate over the selection mechanism and more debate over the selection criteria.

In an effort to address the issue of transparency of venue selection, the IAOC asked for comments from the community about potential problems regarding the following cities under consideration: Paris, Montreal and Copenhagen.

Next IETF: Berlin, 17-22 July 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 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

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