Council of European National
Top-Level Domain Registries

Report on
IETF100

Singapore
11-17 November 2017

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

IETF100: “We make the internet better… for humans?”

Toasting to its 100th meeting, the IETF plenary heard a call for engineers to pay attention the societal consequences of technologies they were standardizing for. The time when internet engineers could step away from taking responsibility for potential collateral damage or abuse of their products was over, said Henning Schulzrinne, professor at Columbia University, long-time IETF participant and FCC official during the Obama administration.

During a special panel discussion with Schulzrinne, Monique Morrow, founder of the initiative the Humanized Internet and Jun Murai, father of the internet in Japan, Schulzrinne called on his colleagues to change the self-attributed IETF mandate “we make the internet better” into “we make the internet better for humans”. Perhaps the most surprising revelation of the IETF100 “birthday plenary” was the fact that Schulzrinne’s concerns over some technical trends resonated with at least some of the participants.

From tool of empowerment to tool of suppression

The original idea of networking as a positive tool of empowerment and democratization has not held the test of time. Meanwhile, Schulzrinne said that increasingly, the goal was to restrict communication and the network was enabling authoritarian states and suppressive societies. The job of engineers was no longer limited to “getting to play with the good stuff”, he said.

Considering potential ways their technology could be used and abused was part of the discussion. Morrow also underlined the need for the engineers to be aware of ethics. He pointed to cyber warfare and profiling to add to the list of not so positive technical developments. At the same time, he warned against the much discussed “politization” of the internet.

Engineers could hardly make judgment calls on these topics, said Schulzrinne, and as polis in its original sense meant “community of citizens and its governance in a natural way”, the internet should be part of the political discussion, but should not become a tool for that discussion.

Looking ahead into potential technology trends, Schulzrinne described a rather large spectrum. Quantum transport and/or projects like the Brain Circuits Project (BCP) might change the “transport” paradigm and make TCP/IP obsolete. At the same time, technologies have proven to be enormously long-living, so Schulzrinne expects not only that the network will become the third commodity after water and electricity, but also that in 2047, even Ipv4 packets might still be around.

Coming full circle and back to monopolies

With regard to economics, Schulzrinne portrayed the possible situation in 2047 as having come full circle. In 1986, when the IETF started, there were the telecom incumbents. After having broken these up, the internet was once more on its way to being dominated by a few large network/content integrated platform providers. This thought had been explained recently in more depth by Geoff Huston (see CENTR's RIPE75 report). Contrary to the early IETF days, in a world of a few giant companies, it could become more difficult for the engineers participating in the IETF not to question – as individuals and citizens – the strategies of their employers, Schulzrinne said.

More and more operators of data centres/enterprise networks don’t have any understanding of the specs. With network automation being a big trend, they just bought hardware and software. Soon, nobody will know who is producing the standards and where they have been produced, Schulzrinne said.

The IETF has been affected by these trends in several ways. It had to adjust financially. Having benefited from the diverting of money set free by the change from Telecom switches to cheaper internet technology, money was again re-directed and spent elsewhere, not for the utility “network”. The number of participants could also decline and in fact already has.

Wither IETF?

A decline in the number of IETF participants was briefly discussed in the IAOC part of the plenary. The IETF99 in Prague was attended by 153 people less than expected and budgeted: the meeting was therefore $250,000 USD under budget. Paid attendance at the IETF98 Chicago was short of 105 participants and IETF97 short of 127. The gap between forecast and actual attendance results in funding gaps, which the IAOC answers by calls to ISOC (see also for the IETF Buenos Aires meeting).

New sources of funding for the IETF have been on the leadership's agenda for several years. One approach the current IETF Chair and her predecessor were championing was to bring new groups, beside the classical vendors, and new work to the IETF. During the Singapore meeting, Routing Area Director Alia Atlas organized a meeting on IETF outreach activities, which spelled out the various type of activities. They range from the very successful hackathons (which still struggle to find sponsors) to special remote hubs (in countries like India).

IETF100 panellist Jun Murai asked the IETF engineers to "de-silo" work on future technologies, as many areas like medical technology or agriculture were not aware of the technology developed at the IETF.

New options pursued with regard to closing funding gaps seem to be more important in light of Schulzrinne's comment on the decline being a trend of changing markets.

During the IAOC plenary, IAOC Chair Leslie Daigle announced that for 2018, the IAOC had asked ISOC to pay an additional $900,000 USD to make up for an expected decline in revenues of $1M USD (with a remaining budget of $7M USD, and ISOC's contribution already making up around fifty percent of it, final numbers are to be decided and published after the ISOC Board meeting at Singapore following the IETF meeting). For 2019, IETF participants will face rather steep increases in registration fees for the meetings, with over 10 percent for 2019 and more than 3 percent as of 2020. A question discussed during the IAOC plenary was if in the future, remote participants should pay for their participation. Numbers, but also cost of the remote participation, is on the rise and burdens the IETF meeting budget.

IASA 2.0: IETF is not seeking independence from ISOC, just more money

At its meeting in Singapore, the group considering the potential reforms to the current administrative system of the IETF rather clearly hummed against a full separation of the IETF from ISOC. Instead, two equally large hums favoured either an evolutionary path to an IASA 2.0 structure or a subsidiary/stakeholder organization with more control over budgetary and contractual decisions, yet still as a body of ISOC.

A mere IASA 2.0 could try to solve the top issues like the relation between ISOC and the IETF. A subsidiary, according to Brian Haberman's presentation, could have "its own bank account, bylaws, charter, board, staff, and corporate identity".

A nice illustration by Rich Salz (Akamai) of the three options was shared after the meeting:

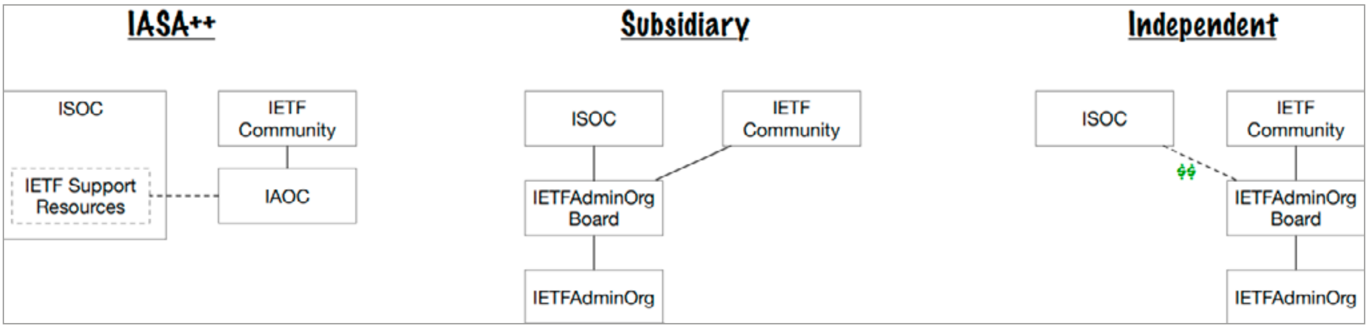

12 Years after IASA 1.0

IASA, which includes IAOC (the IASA oversight body), the single IASA employee, the IAD and the IETF trust were developed and established during a first round of updating the IETF structure and initiated by the then IETF Chair Harald Alvestrand (then Cisco) around 2003. Alvestrand pushed the establishment of the administrative structures to answer growth and professionalization needs. It fell to the second Scandinavian IETF Chair, Jari Arkko, to initiate the current review, 12 years after IASA 1.0.

Opinions differ on the way forward. Former IAB Chair Andrew Sullivan said that the IETF had to make up their mind if it wanted to be "adult" or "grown up". On the other end of the spectrum, John Levine warned that it would be rather difficult for the IETF to raise a similar amount of money as it now received from ISOC, "with so few strings attached".

The funding issues and the changes to the IETF, but also to ISOC on the other hand, are at the core of the discussion. When the current IASA was initiated in 2004/2005, ISOC was still a small organization with around 10 employees, just having received the contract to be the sponsoring organization (and beneficiary) of the .org registry, earlier managed by VeriSign. While the .org contract allowed ISOC to generously sponsor the IETF (and step in whenever the IETF went off the rails budget-wise), ISOC has grown into in size and budget, and has become a 100-employee lobbying organization, which sometimes is more involved in technical discussions, instead of only serving as the home for the unincorporated IETF.

Clarity, budgetary control, new funding sources

The issue list of the IASA 2.0 RFC highlights the lack of clarity (with regard to responsibility, representation of the IETF, authority and oversight over staff and budget) as well as the lack of resources and a lack of transparency of the current IAOC. According to the draft, the goals discussed in Singapore are:

- to protect IETF culture
- improve the working environment for standardization
- define the IETF-ISOC relationship
- support a re-envisioned funding model
- provide clarity about the IETF-ISOC financial arrangements
- clarify the overall roles and responsibilities and also support staff roles and responsibilities
- re-define the role of the IETF community in relation to administrative activities
- define improved transparency requirements
- define a transition plan

ISOC's Kathy Brown underlined ISOC would be supportive of the IETF either way and intended to leave the decision to the IETF. At the same time, Brown noted that where ISOC was paying bills, they certainly had accountability towards their Board. ISOC is the official contractor for the conference hotels and the employer of the IAD, who was Ray Pelletier.

Pelletier has retired as of 31 October. Since 1 October, Portia Wenze-Danley has been hired as Interim IAD.

Some observers think that given the financial links between both organizations, not all the issues can be fully solved.

The fight around encryption

With additional encryption in many places, TLS, Quic and (slowly) DNS, there is a pushback from middle box operators/vendors – or, as those concerned about the pushback say, a few of these operators/vendors.

The heat of the discussion in Singapore concentrated on an individual submission (not going through the IESG) in the Ops Area. Titled "Effect of Encryption on Operators" it is seen by many critics as a potential reference document for operators to ask for limitations (or exceptions) in encryption put into new standards.

The document was presented by Kathleen Moriarty, down-stepping Security AD, who after much critic underlined she only inherited it from her AD colleagues. Moriarty said that in essence, the document was a result of the post-Snowden RFC on pervasive monitoring (RFC 7258) which included the acknowledgement: "Making networks unmanageable to mitigate PM is not an acceptable outcome, but ignoring PM would go against the consensus documented here. An appropriate balance will emerge over time as real instances of this tension are considered."

The just under 50-page document describes how, through the added encryption, operators "lose" options for passive monitoring, traffic optimization and management. It includes a special section on the issues for mobility network optimization. The "response to increased encryption and looking forward"-chapter (chapter 8) once more favours "considerations for protocol design should factor in network management functions to work toward the balance". Without taking the concerns into consideration, there was a risk that as the IETF standardizes encryption for its protocols, it would not be fully implemented, Moriarty warned in Singapore. There was not much discussion at the Ops Area meeting in Singapore: most participants seemed to support the document, even if one participant reported that when presenting the document to members or RIPE, some had asked if encryption should now be considered as "bad".

The main discussion is taking place instead on the mailing list with the most recent, longer analysis by privacy expert Christian Huitema illustrating some essential concerns. Huitema questions consensus on some of the network operating mechanisms – what some defend as performance enhancing proxies are in fact performance decreasing proxies to others. Another mechanism claimed to be lost, HTTP header insertion, he writes, is no networking management tool and should not figure in the document at all. While Huitema acknowledges the considerable re-writes, he still expresses the opinion that more work has to be done.

While the document nevertheless looks like to be in the stretch run to be published as an informational RFC, the discussions on how to balance encryption – network management/troubleshooting/monitoring goes on.

Quic continues to discuss how to address concerns, for example in the "spin-bit" draft. The spin-bit in the Quic header will allow operators to measure end-to-end RTT on QUIC flows.

An attempt to address third party monitoring without risking ossification was just started with a document looking for randomization – "greasing" – in the current Quic header.

In the TLS WG there are at least two proposals laying out how to balance the "more encryption" vs manageability concerns, the most recent being a proposal by Cisco that wants to move TLS one layer up to the application layer. The mechanism presented in Singapore "defines a mechanism for transporting TLS records in HTTP message bodies between clients and services. Reactions in Singapore were very mixed, with privacy advocate Daniel Kahn-Gillmore (ACLU) and HTTP Chair and Quic Co-Chair Mark Nottingham (Akamai) warning against further cat-and-mouse games instead of pushing for middle box compliance with the new protocol. An older proposal, following the discussion of data centre issues with TLS 1.3, is from Russ Housley (Vigil Security, former IETF Chair and NSA contractor). He proposes a "TLS Visibility Extension" to specifically address one of the impacts of (EC)DH "through an opt-in mechanism that allows a TLS client and server to explicitly grant access to the TLS session plaintext." Neither of these documents have been adopted as WG documents, but discussion on the middle box issue in TLS is ongoing.

DNS activities and beyond

With a number of DNS-related activities ongoing outside of the DNS WG, one might nearly ask if there would could be a need to revive good old DNS extension WG (which years ago standardized DNSSEC, for example).

DNS over HTTPS

A rather straightforward activity is the DNS over HTTPS work, which, after being presented in the Dispatch WG during IETF99, is now being pushed ahead in a new, dedicated WG. The "DNS queries over HTTPS" (DOH) WG is chaired by Ben Schwarz (Google – and Google being one of the early users of a DNS over HTTPS solution) and David Lawrence (Akamai), and is working its way over an already pretty short list of issues, one of which is the different ways of caching between DNS and HTTP.

The basic use cases for DNS over HTTPS according to a re-write of the draft text are "to prevent on-path network devices from interfering with DNS operations", with interference including "spoofing DNS responses, blocking DNS requests, and tracking." For this use, clients "will be explicitly configured to use a DOH server as a recursive resolver by its user (or administrator)" for some or all queries. A second use case is "allowing web applications to access DNS information, by using existing APIs in browsers to access it over HTTP in a safe way consistent with Cross Origin Resource Sharing (CORS) [CORS]." For the second use "the browser does not consult the DOH server or use its responses for any DNS lookups outside the scope of the application using them; i.e., there is (currently) no API that allows a Web site to poison DNS for others." Contrary to the DOH draft, the DNS wire format draft was "proxying DNS queries over HTTP instead of over DNS itself", note the authors of the DOH, Paul Hoffman (ICANN) and Patrick McManus (Mozilla).

DNS as a brick for a federated single sign-on system?

The DNS as the basis for a federated single sign-on system is at the heart of a proposal and prototype presented during the Oauth WG session by Marcos Sanz, Denic and Vittorio Bertola, Open-Xchange. While based on OpenID Connect, using the DNS for the "Public ID infrastructure" would bring real interoperability and a better way for multiple providers to offer identities the same way. Portability, too, would be gained, the authors said. According to them, using the DNS would in fact allow "the user, rather than the identity provider to become the sole owners of their identifier by acquiring a personal domain name". The proposal was rejected by several participants of the WG, underlining that similar attempts for a federated ID system had failed already. On the other hand, the authors pointed to an ongoing test implementation of the system and are planning to present the PIDI and the related discovery mechanism outside of the IETF as well (e.g. at the upcoming Domain Pulse in Munich).

Rethinking the DNS (again)?

With all the developments around the DNS, a long-time IETF participant and author, John Klensin, asks a recurring question in an individual draft: is it time to re-consider the DNS or think about a replacement? Klensin explains how from multi-type queries over privacy or the special name discussion, the DNS obviously did not meet the expectations people had about its functioning. Some of the band-aids produced in recent years just illustrated this. The author underlines that the document at least might help "to stimulate thought about how far we want to try to push the existing DNS, to examine whether expectations of it are already exceeding its plausible capabilities, and to start discussion on a redesign or alternatives to one if the time for that decision has come."

DNSOP: Yet another fight around a .internal (.homenet) TLD, defining the DNS and more

In 2017, the DNS WG published 5 instead of 7 RFCs and so has slowed down a little, according to Co-Chair Tim Wizinski. Still with a lot of attempts to (re-)use the DNS (see highlight section above), the WG members had a rather full agenda.

One of the document the WG hopes to get broad review on before going to last call (mid-January) is the update to DNS terminology. These definitions are expected to set a standard and be used broadly as authoritative for DNS concepts and terms. It will be a full standard document (as the predecessor was informational).

.internal instead of .homenet?

The WG entertained another edition of the fight around a special TLD name after Warren Kumari (now IETF Area Director) proposed the IETF should introduce an ".internal" TLD on the basis of the Special Names process. Kumari argued that at least some people now squatting on TLDs for internal use (like .home or similar) could be expected to use such a TLD. The difference with regard to the failed attempt to have .homenet delegated according to Kumari is that .internal was intended for much broader uses – and delegation was not time critical as no special protocol that was worked on was dependent on the delegation. Kumari asked for the .internal TLD to be assigned to the IANA and a DNSSEC insecure delegation be inserted in the root zone: requests to the root shall cog to a delegated blackhole at iana.org, according to the draft.

Kumari's proposal – similar to the earlier .homenet application – met considerable objection. Andrew Sullivan (Oracle) questioned the statement (in the draft) that there was no process for the delegation and pointed once more to the ICANN process. David Conrad (ICANN) on the other hand said that he did not fully understand why, given the special names procedure at the IETF, .internal should be thrown over to ICANN. Stuart Cheshire (Apple) complained the IETF continued to ignore what was happening outside of the IETF. There was no decision in any way on Kumari's draft and proposal, but one might wonder what would happen if .internal got a chance to proceed, given that homenet was sent back to homenet.arpa.

Key-roll issues

Work on the deferred key-roll is also going on at the IETF. In Singapore, Geoff Huston presented a proposal intended to help to get a better hold on the distribution of the new trust anchor. The proposed mechanism puts the measurement on the client side. By using a set of queries with special tags, users shall be able to check if a "special Root Zone KSK is ready to be used as a trusted key within the context of a key-roll by this resolver".

According to the draft, the sentinel process will test with three names:

- a validly signed name so that responses about names in this zone can be authenticated by a validating resolver - name containing the left-most label "_is-ta-<tag-index>."
- another validly-signed name - containing the left-most label "_not-ta-<tag-index>.".
- a name signed with a DNSSEC signature that cannot be validated

The responses of the validating server allow to determine the key state of the resolution environment of the user.

o Vnew: A DNSSEC-Validating resolver that includes this mechanism that has loaded the nominated key into its trusted key stash will respond with an A record response for "_is-ta", SERVFAIL for "_not-ta" and SERVFAIL for the invalid name.

o Vold: A DNSSEC-Validating resolver that includes this mechanism that has not loaded the nominated key into its trusted key stash will respond with an SERVFAIL record for "_is-ta", an A record response for "_not-ta" and SERVFAIL for the invalid name.

o Vleg: A DNSSEC-Validating resolver that does not include this mechanism will respond with an A record response for "_is-ta", an A record response for "_not-ta" and SERVFAIL for the invalid name.

o nonV: A non-DNSSEC-Validating resolver will respond with an A record response for "_is-ta", an A record response for "_not-ta" and an A record response for the invalid name.

While some concerns were raised on the DNSOP mailing list if it was not preferable to put a telemetry interface at the client side – including some concerns about potential privacy issues when testing via end user systems – the majority of experts support fast adoption and implementation. Huston said, it was up to everybody how they performed the tests, while he intends to rely on his well-known classical add-based testing set-up.

David Conrad, ICANN, in Singapore said he preferred fast adoption. For ICANN the mechanism could help to get a clearer picture of the distribution of the DNSSEC Root KSK.

With regard to the security consideration draft for 5011, automatic key-rolls, the WG briefly discussed the critical comments that this work lacked operator input and should not be published (Ed Lewis, ICANN). Nevertheless IETF participants including ICANN representatives supported adoption of the document. A major open question posed during last call is if time for the intervals (when it is safe to revoke old keys and so on) should be based on intervals or wall-time. The complexity of the math in the calculations was said to be an issue.

A new draft document which the WG still has to adopt relates to guidelines for DNSSEC validation.

More information in DNS answers

Several documents on the agenda of the WG are about additional information sent on in DNS answers. On its way through the WG are extensible DNS error codes, that will allow to provide additional information to serve fail answers, hinting for example at the cause of DNS and DNSSEC failures. A registry shall list various – and future – error codes.

Kaznori Fujiwara made another proposal to provide multiple answers in a single DNS response. Authoritative servers should for example add NSEC resource record or A/AAAA resource records of the query name, even if not asked for it. While many in the WG warned that such a mechanism would ease amplification and DDoS attacks and said that pull instead of push mechanisms were preferable, Fujiwara pointed out that the enrichening of answers was already common practice. He provided an overview over the many proposals that have been made to standardize the mechanisms.

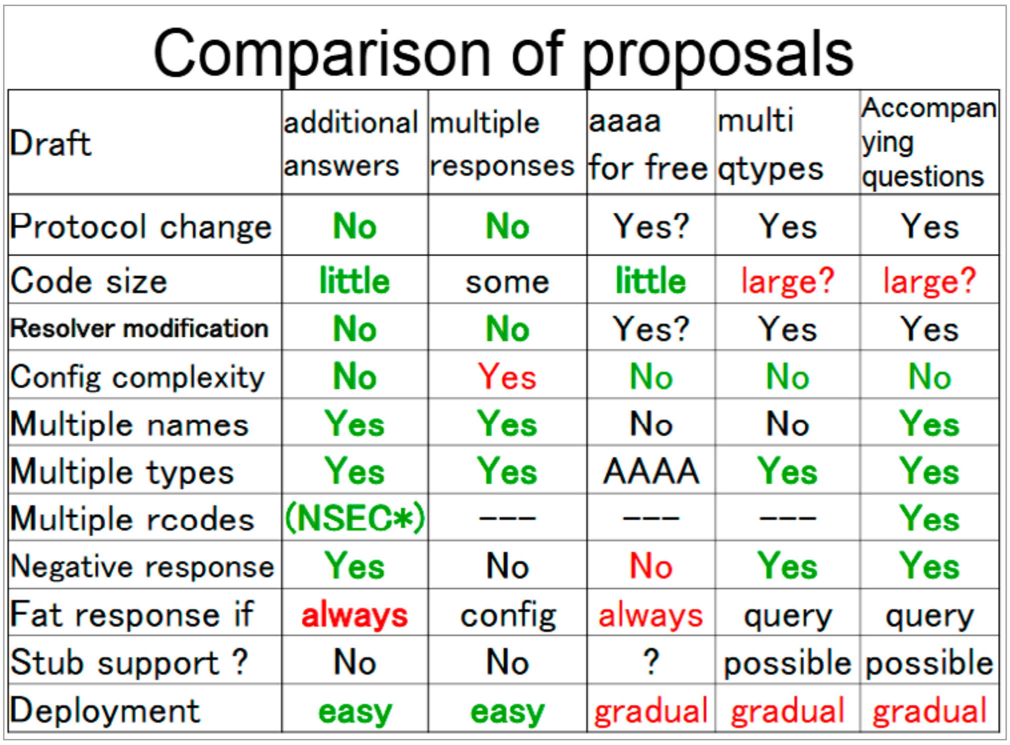

Comparison of proposals

Draft	additional answers	multiple responses	aaaa for free	multi qtypes	Accompanying questions
Protocol change	No	No	Yes?	Yes	Yes
Code size	little	some	little	large?	large?
Resolver modification	No	No	Yes?	Yes	Yes
Config complexity	No	Yes	No	No	No
Multiple names	Yes	Yes	No	No	Yes
Multiple types	Yes	Yes	AAAA	Yes	Yes
Multiple rcodes	(NSEC*)	---	---	---	Yes
Negative response	Yes	No	No	Yes	Yes
Fat response if	always	config	always	query	query
Stub support ?	No	No	?	possible	possible
Deployment	easy	easy	gradual	gradual	gradual

Other documents discussed in Singapore included the much debated DNS proxy document by Ray Bellis. The document does include a privacy consideration section, underlining that if “used incorrectly, this RR could expose internal network information”. As the specification was only intended for use of a server-side proxy that would be under the same administrative control as the DNS servers themselves, “there was no change in the scope within which any private information might be shared”.

Unexpected Dprive side-meeting: Implementation steps in DNS Privacy

Members of the DPRIVE WG gathered in Singapore for a side meeting at the request of WG members to report on progress of implementations, discuss the document on padding and to ask for next steps.

Android DNS over TLS Client Beta-ready

Implementations were presented by Eric Kline and Ben Schwartz for the Android Open Source Project and by Sarah Dickinson (remotely) from the DNS Privacy Project. DNS over TLS can now be put on Android phones, with the code sitting in the Android library. Once downloaded, users can choose from three different options: privacy mode, opportunistic mode and privacy off mode. Once turned on, the client tries to connect to the DNS resolver via DNS over TLS and if the server provides it, goes encrypted. A live test at the Singapore IETF meeting worked in opportunistic mode, as the IETF meeting network had a DNS over TLS enabled resolver by Warren Kumari (Google). With this implementation, DNS over TLS requests can start getting serious numbers.

Microsoft command-line GUI and Android GUI ready for Stubby

A second major effort are the implementations prepared by the DNS Privacy Project. Sara Dickinson, Sinodun, announced the upcoming start of a user-friendly GUI for Mac OS – planned for the week after the IETF meeting. Interest in using DNS over TLS seemed considerable, she reported from the most recent launch of a Microsoft client, which for the time being is still command line based, but would eventually be complemented with a user-friendly GUI.

The differences between the Google and the DNS Privacy Project implementations lie mainly in the pre-set group of DNS over TLS enabled resolvers of the latter, while the Android implementation just tries the recursive resolver at hand. Another difference is that for her implementations, Dickinson chose to already implement padding, while Google/Android so far has not.

Working Group last Call for Passing: Between Latency and Security considerations

Alexander Mayrhofer, nic.at, presented the pending draft on padding, explaining the background for the choice of 128 bits on the client and 426 bits on the server. The choice was based on mathematical analysis by Daniel Kahn-Gillmore (American Civil Liberty Union, ACLU). Mayrhofer asked for additional comments on the resulting 300 to 400 bits of additional load. The rather “generous” choice could make a difference in latency especially for providers that connected devices sitting on edge networks and having a couple of hundred million requests per day (resulting in losing 300 bits a couple of hundred million bits). He said a more conservative choice would also be possible. While additional academic research was called for on the mathematical basis, the WG members present and WG Chair Tim Wicinksi agreed to go ahead with WG last call. Delaying the decision further could result in implementations picking different padding policies. Such differences in padding could result in helping to identify the source of encrypted traffic.

Dprive will meet during the next IETF in London and will then finally start to talk about the re-charter to consider securing the resolver to authoritative server part of the DNS. Participants of the Singapore side meeting also loosely agreed to plan for some kind of interim meeting before the next IETF, either remotely or alongside the second edition of a DNS Privacy Workshop at the Network and Distributed System Symposium (NDSS) on 18 February 2018.

The WG needs a new co-Chair, as Warren Kumari stepped down.

DNSSD: Advancing drafts and talking privacy

After quite some years of standardization, the DNSSD WG has started to consider privacy implications of service offer, service discovery and use of service use. Information leaked includes host names, network parameters and also further description of corresponding service instances. Discovery at public hotspots can result in serious privacy problems, according to a current draft by Christian Huitema and Daniel Kaiser, University of Constance.

Huitema's/Kaiser's draft proposes as a solution that hosts discover Private Discovery Service Instances via DNS-SD using special formats to protect their privacy in a first stage. In the second step hosts directly query these Private Discovery Servers via DNS-SD over TLS, with a pairwise shared secret necessary for connection establishment. A draft on securely pairing "by agreeing on a secret and manually verifying the secret's authenticity using an SAS (short authentication string)" is here, an accompanying draft on pairing issues here.

In Singapore, Stuart Cheshire (Apple), one of the main authors of the WG (also author of a DNSSD overview "roadmap" document), presented a draft that tries to compile the various privacy aspects DNSSD authors might consider, depending on the features of their respective drafts. Goals to be considered (and weighed according to situation and protocol efficiency) according to Cheshire are authenticity and integrity, confidentiality, anonymity and resistance to several kind of attacks (dictionary attacks, tracking, message linking and denial of service).

In Singapore, Cheshire expressed his opinion that a current Huitema/Kaiser draft was not covering the full range of issues. He presented a longer list of work on DNSSD privacy mechanisms, he included several Apple technologies, like the option of contacts only mode for connection establishment in Apple AirDrop, the "finding your home accessories"-mechanism of Apple HomeKit. Similar work could also be found in Google Nest accessories (IEEE 802.15.4 mesh networking) and Zigbee dotdot. Cheshire also spoke of two still confidential projects that are ongoing and pointed to a patent just granted to an Apple project of his own, abandoned five years ago. No IPR disclosure has been made so far, yet it sounded as if Cheshire at least wanted to underline Apple's earlier interest (and/or claim?).

A document soon to be published as an RFC, currently pending before the IESG, is the one on a "Discovery Proxy for Multicast DNS-Based Service Discovery". The planned RFC makes the attempt to combine the ease-of-use approach of Multicast DNS for service discovery in a local network with the efficiency and scalability of the classical Unicast DNS. The new discovery proxy uses Multicast DNS to discover Multicast DNS Records on its local link and makes corresponding DNS records visible in the Unicast DNS namespace. This is where the naming architecture ideas discussed in DNSSD, and much more, Homenet comes in. The approach mitigates issues arising in larger networks with multiple links (between which multicast DNS are not propagated).

According to the future RFC, the basic mechanism for the discovery proxy is:

"In simple terms, a descriptive DNS name is chosen for each link in an organization. Using a DNS NS record, responsibility for that DNS name is delegated to a Discovery Proxy physically attached to that link. Now, when a remote client issues a unicast query for a name falling within the delegated subdomain, the normal DNS delegation mechanism results in the unicast query arriving at the Discovery Proxy, since it has been declared authoritative for those names. Now, instead of consulting a textual zone file on disk to discover the answer to the query, as a traditional DNS server would, a Discovery Proxy consults its local link, using Multicast DNS, to find the answer to the question."

According to several IESG members, changes needed include a more in-depth analysis of the security risks and elimination of the IPR claim of Apple, which is integrated in the document under point 10. IPR claims usually are not included in RFC text.

A document on DNS push notifications is also going to the IESG and to last call. It allows clients to be updated on changes in DNS records on subscription unrelated to a DNS request.
The push notifications draft is dependent on finalization of DNS session signalling, introduced in a draft on stateful DNS. It shall allow to reduce per message session signalling by introducing TLVs to manage timeouts and terminations for DNS sessions (see DNS OP). The notification draft awaits the finalization of the DNS session signalling draft that introduces a standard way to last call is expected in December 2017.

Finally, DNSSD also attracts interest from the WG on "Autonomic Networking Integrated Model and Approach" (ANIMA). According to its charter, the main objective for the WG is "to develop a system of autonomic functions that carry out the intentions of the network operator without the need for detailed low-level management of individual devices". For this kind of network automation, there is also a need for service discovery.

Toerless Eckert (Huawei) in Singapore briefly presented the work on discovery, synchronization and negotiation considered in a draft on a "GeneRic, Autonomic, Signalling Protocol" (Grasp). Grasp shall "enable autonomic nodes and autonomic service agents to dynamically discover peers, to synchronize state with each other, and to negotiate parameter settings with each other", according to the draft. DNSSD is acknowledged in Grasp as a possible discovery mechanism for some parts, especially for application layer services.

Brian Carpenter, one of the authors of Grasp, underlined that in the future there might also be a need for a dedicated IANA name space for Grasp – so possibly, in addition to the home.arpa draft underway in homenet, a similar discussion for a draft might lie ahead in the future.

The link of the DNSD and the Homenet WGs was briefly discussed during the DNSSD session by Ted Lemon, author of a number of Homenet drafts, including the Simple Name Architecture draft. In essence, Homenet was one use case of DNSSD, said Lemon. It only lacked the professional management in place in DNSSD environments, which have been driven by Apple's Stuart Cheshire in the first place.

Cheshire and Lemon were the only participants, according to Cheshire, in the first ever DNSSD slot at the IETF100 Hackathon meeting. Cheshire said he intended to apply for another slot during IETF102.

Homenet – ISP or user model

Work on homenet is progressing rather slowly. While the draft on introducing home.arpa is in the RFC editor queue and Babel for a routing protocol is far advanced, there was considerable discussion on the potential models for implementation and practical gaps in the homenet architecture.

Following a presentation of IPv6 expert Jordi Palet (from Consulintel.es), so far homenet protocols were not yet implemented in home routers, confirmed Hans Liu from Dlink. Ted Lemon acknowledged that there could be a need to build homenet into routers to allow for implementation apart from the exploratory steps made by WG members and a few geeks.

The WG was asked for guidelines if homenet should only be implemented in high end retail routers or also in ISP routers. The WG discussed the pros and cons of either a "friendly ISP homenet router" vs the "my router is my castle" concepts, with several pointing out that there was not much incentive for ISPs to offer homenet functions including, for example bridging between several networks in one homenet. Making the ISP the manager of homenet functions could also result in issues from the regulatory point of view, with Europe's GDPR being mentioned by one participant. If users on the other hand are to be the users of homenet, there was still quite some work ahead as the "my router/home is my castle" approach was currently only for geeks.

On the homenet architecture Andrew Sullivan (Oracle) pointed to gaps in the current document which made a number of features a "must implement", for example secure delegation and DNSSEC, but did not explain in the draft specification how this should be done. Ted Lemon and Stuart Cheshire argued that the WG had hesitated to allow for a full-fledged homenet naming architecture, so the more complicated things were cut out. When talking about homenet security, Lemon underlined that secure delegation, DNSSEC and other things would be easier with a global DNS name.

The homenet simple naming architecture document covers "local publication of names, as well as name resolution service for local and global names for devices connected to the homenet", but not DNSSEC or a global DNS name which has been the topic of an earlier draft.
Security mechanisms and trust establishment are now to be made new topics which was also discussed in the Babel WG.

IETF101 will be held on 17-23 March 2018 in London, UK.

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 54 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