Council of European National Top-Level Domain Registries

Report on IETF106

Singapore
16-22 November 2019

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

Battle over PIR

Ethos Capital has agreed to buy Public Interest Registry (PIR) with all its assets from the Internet Society (ISOC). The news was announced on 13 November 2019, just three days before the start of the IETF meeting in Singapore. What might look like a normal business transaction from the outside has put some parts of the internet community on alert, who consequently got back to ISOC with questions.

Selling a Cash-Cow?

The most important question was certainly about the motivation for selling what can be seen as ISOC's cash-cow for the last one-and-a-half decades. In 2018, it was agreed that PIR would only could give ISOC 43 million US Dollars, earned from the registry business with .org, the newly started .ngo and the Cyrillic version of .org. In 2017, under an exceptional wave of domain registrations, ISOC received nearly double that amount, close to 80 million US Dollars. Without counting an outlier year which partly resulted from the hording of domains, PIR has still secured a relatively stable income over the years, allowing the organisation to grow from a two-person office to a fifty-people organisation. The question therefore was, has the organisation sold its cash-cow?

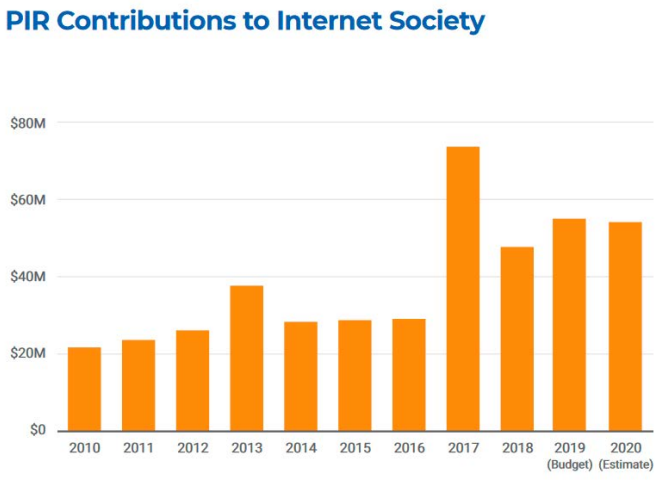

Amounts shown in US Dollars. Source: Public Interest Registry[2]

ISOC CEO Andrew Sullivan explained that a core motive for the deal was to diversify the funding sources for the organisation. As it is only dependant on one industry, the domain name business, the organisation is more subject to the fluctuations of this industry. With the amount received from the sale ISOC would be free, according to Sullivan, to make more diversified investments. The ISOC CEO acknowledged that to match the current income, the sale must have resulted in a large amount of money.

At the same time he pointed out that large and medium-sized foundations were regularly able to earn 8 or 9 percent from their endowments. It was certainly necessary to have the right advisors in place to make good investment decisions, he said, adding that currently he could only say that the ISOC Board had done its job. The actual amount that Ethos will pay will not be announced for now. This was a request from Ethos Capital, according to Sullivan. However, as soon as ISOC does its own annual reporting, the amount will become public.

Calculations vary about how much ISOC needs to make from the deal to match its current income stream, some think 500 to 600 Million US dollars will be enough, whilst others say a billion is needed.

Secrecy

The amount of secrecy around the deal was another question at least some observers, ISOC members and chapters have raised. The ISOC Board, PIR and Ethos Capital seem to have done a perfect job at keeping negotiations under wraps. While again, this might have been a request from the new investor, it has certainly had the effect of feeding the suspicion and making the promise "Following the close of the transaction, PIR will continue to meet the highest standards of public transparency, accountability, and social performance in line with its long-standing purpose-driven mission, and will consider seeking B Corporation certification" sound at least a little hollow.

Who is Ethos Capital?

Many observers have also asked who Ethos Capital is, as it is no household name in the domain industry. In fact the company looks brand-new, the website lacks the details you would ask from a transparent organisation, and in some countries, it even lacks a lawful company web presence. The link to the domain name industry is that its founder, Erik Brooks, who was working for Abry Partners at the time, organised Abry's acquisition of Donuts. Through that deal, in which former ICANN CEO Fadi Chehade acted as

a consultant, Abry got into contact with Chehade and ICANN. The only other person mentioned on the ethoscapital.org website is a former ICANN employee. Chehade in fact registered the domain name ethoscapital.org in May 2019, just around the time when ICANN announced that it would lift the price caps for .org (alongside other TLDs).

Due to intense criticism and a number of highly critical news pieces (see for example Kieren McCarthy's article) after its board meeting in Singapore, ISOC published an FAQ tackling the major concerns of chapters and ISOC members, of which some examples can be found below.

"Is Abry Partners involved in this transaction?

Abry Partners is not involved in this transaction. Abry Partners is a private equity firm where Erik Brooks worked for 20 years, prior to leaving and starting Ethos Capital.

Is Fadi Chehade involved in this transaction?

Fadi Chehade's company, Chehade & Company, is an adviser to Ethos. Chehade & Company is an advisory company with clients across the technology, education and creative sectors.

Mr Chehade is a board member of Sentry Data Systems and Interactions LLC and serves as an advisory board member of the World Economic Forum's Center for the Fourth Industrial Revolution. Previously he was the President and CEO of ICANN, a member of the UN Secretary-General's High-Level Panel on Digital Cooperation, and a Senior Advisor to the Executive Chairman of the World Economic Forum."

What does this mean for the org-community and the IETF?

For the IETF, one question of the deal was important. Would ISOC continue sponsoring the IETF? Or would it become a task for the newly-established organisation to sponsor the IETF? Despite the IETF's plan to become more independent, organisationally and also financially, ISOC is still a major source for funding or, at least, a safety net for the IETF.

Sullivan clarified that funding of the IETF would not come from Ethos Capital's "Community Enablement Fund to support the financing of current and additional initiatives undertaken by key internet organisations" (one of three self-obligations Ethos announced according to a blog post with ISOC). The funding for the IETF would continue to come from ISOC, Sullivan underlined.

Opponents and proponents agree on one thing, namely that .org registrants can expect prices to go up under Ethos, especially if it bought PIR for a large sum. Where there is disagreement again is how much this will hurt individuals and not-for-profit organisations with more than one .org domain. The proponents claim that only hoarders will be affected, but some think NGOs in developing or least developed countries will also suffer.

Will regulators approve the deal?

There are altogether three different "authorities" that have to approve the deal signed by the parties: ICANN, the State Attorney of Pennsylvania (where PIR is incorporated) and the Fund of Orphans and Widows. The State Attorney has to sign off because through this deal, PIR becomes a for-profit private company.

Some of the critics hope that these regulatory steps will stop the deal; one group started a petition on change.org which gained traction rather slowly, with around 420 signatures after one week. ICANN reacted to press requests by simply acknowledging that it had received the request and was checking the details.

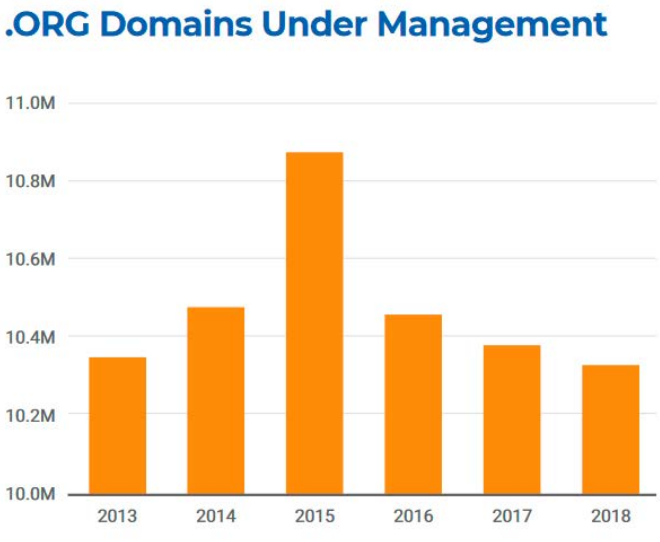

Source: Public Interest Registry[3]

A new DNS = ADNS, ODNS

The quest for solving the DoH dispute continues. Neither Google, nor Cloudflare or Mozilla presented the newest proposal to fix DoH implementation. Instead, a team of Apple engineers (plus a fast, new 'acquisition', Patrick McManus, previously Mozilla) presented the "Adaptive DNS" (ADNS). According to Tommy Pauly (Apple), the declared goal of the proposed new specification was to improve privacy without having a fixed public resolver.

An ADNS architecture

Instead of sending all DNS queries to one fixed resolver – like in Mozilla's DoH implementation – ADNS works through a list of options to seek name resolution. While it does not require user intervention to make the choice, ADNS is dependent on a number of new elements, developed in other Working Groups (WGs). The necessary elements of the ADNS architecture are:

1. a DNS record that indicates a designated DoH server associated with a name (draft in DNSOP);
2. an extension to DoH that allows client IP addresses to be disassociated from queries via proxying (draft in DPRIVE I-D.pauly-dprive-oblivious-doh);
3. a DoH server that responds to queries directly and supports proxying;
4. and client behaviour rules on how to resolve names using a combination of designated DoH resolvers, proxied queries, and local resolvers.

The core concept of ADNS is to allow requests to be "adaptively" handled either locally (according to the respective policies, either based on filtering or providing internal name servers) or with a "designated" DoH server responsible for a queried domain which the client knows is offering DoH resolution for the respective domain.

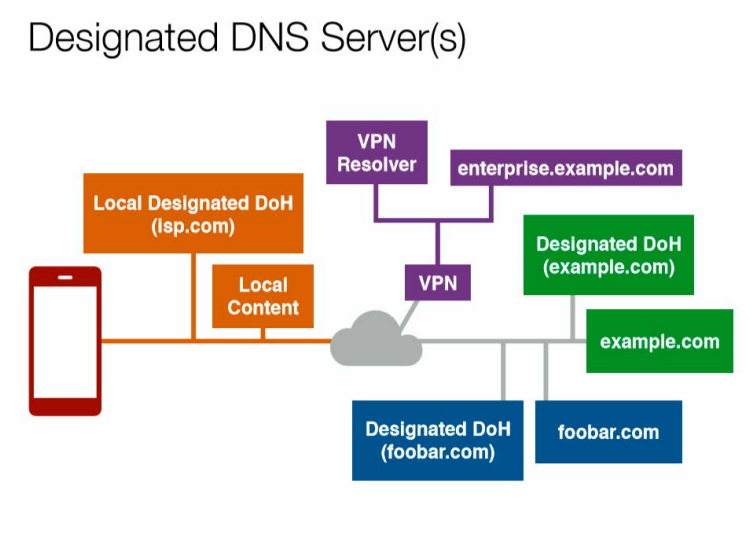

A new record type for designated DoH server

For bootstrapping a client must have knowledge of at least one or two domains with their own DoH resolvers, which either have to be queried over classical DNS, or the known DoH resolvers are made a default. Further on, the draft includes whitelisting or Dynamic Host Configuration Protocol (DHCP) as possible options.

The DoH servers for a given zone indicate their resolver role for the domain through new service binding records, HTTPSSVC, SVCP. According to news from DNSOP, HTTPSSVC/SVCP records will be queried alongside the A/AAAA records. A dedicated draft for this new record was presented in the DNS WG.

Once the relation is established, the designated server not only serves answers for the respective zone, but also acts as a proxy to resolve domains outside of its zone for the querying client.

Oblivious DNS

For highly sensitive content which a user wants to hide from all but the authoritative server, yet another type to resolve queries is proposed. Called "oblivious DNS" (ODNS?) the client sends encrypted requests to an "oblivious proxy" which does not decrypt and answer, but sends them on to another server, the "oblivious target" who decrypts and does the resolving. The concept, which was proposed as an extension to DoH and described in a separate draft, splits knowledge about IP address and query data. A known attack is when oblivious proxy and oblivious target collide.

Deciding which resolver to use

Given there are now a number of different resolving choices, the ADNS draft clearly lists in which order the different resolving modes should be used, depending on the specific hostname:

1. Exclusive Direct Resolver (resolver provisioned by VPN with domain rules for hostname resolved). If the resolution fails, the connection will fail.
2. Direct Resolver, such as local router, with domain rules known to be authoritative for the domain which contains the hostname. If the resolution fails, the connection tries the next resolver configuration based on this list.
3. The most specific Designated DoH Server that has been whitelisted. For example, given two Designated DoH Servers, one for "foo.example.com" and another "example.com", clients connecting to "bar.foo.example.com" should use the former. (privacy sensitive clients should not skip)
4. Oblivious DoH queries using multiple DoH Servers. If this resolution fails, Privacy-Sensitive Connections should not resolve.
5. The default Direct Resolver, generally the resolver provisioned by the local router, is used as a last resort for any connection that is not explicit.

Why DoH and not DoT?

During his presentation, Pauly answered what he said were the FAQ so far. He explained that the choice for DoH instead of DoT resulted from the potential of connection reuse, the option of multiplexing and also the easier migration to the new transport protocol, QUIC. However Pauly pointed out that ADNS could also designate DoT servers. The second issue was the expectation that designated resolvers had to be DNSSEC-signed, otherwise attackers could lure traffic their way. This might be a barrier to entry, given the adoption rate.

WG reaction

While there was positive feedback during the ABCD (DoH follow-up) BoF and during DPRIVE on the attempt to create a decentralized DoH implementation, some observers had some more or less fundamental questions. Alex Mayrhofer (nic.at) underlined that ADNS would create a completely new world on how the DNS is treated, especially by lifting the barrier between resolver and authoritative DNS server, as DoH-designated servers would be authoritative for a domain/some domains as well as resolving others. Ben Schwartz (Google), author of the SVCP draft, spoke of a "mode switching resolver".

Mayrhofer instead called it the hosts.txt file for the 21st Century. Hosts.txt is the list of hosts on the internet that was maintained manually before the DNS was standardized in 1983/84.

Stephen Farrell (IAB, and former Security Area AD) welcomed the proposal, but warned against being too optimistic with regards to DNSSEC deployment. With regards to privacy, Vittorio Bertola (OpenNet) said that while the decentralization was a good step forward, the distribution / spreading of DNS data to various parties was no progress. One can expect DPRIVE to take on the work as a WG item, together with the related oblivious draft - which some said was conceptually close to TOR. If broadly implemented the proposal could change the face of the DNS considerably.

ABCD – a failed BoF

Pauly also presented his draft on the much awaited ABCD BoF, a follow-up to the disputes over Mozilla's DoH implementation. The BoF failed to agree on forming a working group in Singapore, due to an artificially blown-up charter text, for which the BoF chairs are mainly responsible.

Mozilla's canary proposal and more

ABCD saw presentations of Pauly's discovery proposal as one possible mechanism to avoid the enforcement of one DNS resolution mechanism onto everybody using a given application.

Andy Grover presented Mozilla's quick fix to this issue with the so-called "canary" domain proposal. Using a canary-domain test, the browser company will check if clients have set some sort of parental control mechanisms. As part of making DoH the default, Firefox will attempt to resolve the canary domain using the local DNS configuration. If the canary domain is blocked, Mozilla takes this as a signal that DNS parental software is in place and

will not proceed to make DoH the default for the respective client.

Technically, operators have to put the canary domain . use-application-dns.net, on their blocking list to allow for NXDomain or Servfail answers or for the return of a NOERROR code that comes without A or AAAA records. Grover said that the company could not wait for standardization as it needed a quick solution. While he did not elaborate further, the company has come under scrutiny by US legislators and has obviously been pressed for a solution. At the same time Grover said that Mozilla was very interested in getting the solution standardized, to avoid multiple canary solutions for other potential DoH implementers. Mozilla's response to political pressure has resulted in people questioning the declared motivation for DoH in the first place. Mozilla had argued that protecting the network against censorship was one of the motives. The question is how one could stop malicious (state/state network) actors from preventing the encryption of DNS traffic given the option to stop it via the canary.

Debate about the notion of "full consensus" instead of ABCD charter debate

While Pauly's proposal would fit in the DPRIVE WG work (and presumably will end up just there), the canary domain proposal quite certainly does not fit in the DPRIVE WG's charter. Furthermore, the ABCD BoF Chairs listed a number of additional drafts on client configuration and on systematic considerations with regard to centralization and operator's issues with regards to DoH:

2019: Drafts related to client configuration

- DNS Resolver Information Self-publication (adopted in DNSOP)
- DNS Resolver Information: "DoH"
- DNS Resolver-Based Policy Detection Domain (presented in DPRIVE and APCD BoF)
- Adaptive DNS: Improving Privacy of Name Resolution (presented in DPRIVE and ABCD BoF)
- A Bootstrapping Procedure to Discover and Authenticate DNS-over-(D)TLS and DNS-over-HTTPS Servers
- Selecting Resolvers from a Set of Distributed DNS Resolvers
- DNS over HTTP resolver announcement Using DHCP or Router Advertisements
- Indication of Local DNS Privacy Service During User Access
- Client DNS Filtering Profile Request

2019: Drafts on relevant systemic considerations

- DNS over HTTPS (DoH) Considerations for Operator Networks
- A privacy analysis on DoH deployment
- Centralised DNS over HTTPS (DoH) Implementation Issues and Risks
- Centralised Architectures in Internet Infrastructure

The chartering discussion in Singapore nevertheless deteriorated into a quarrel about text the co-chairs had added to the more lightweight and narrowly focussed original text, the concept of "full consensus" in particular resulted in considerable debate. A section that was added shortly before IETF 106 had listed a number of contentious topics (end-user privacy and pervasive surveillance, detection and suppression of malware, use of records from untrusted sources, policy enforcement and control of the stub resolver configuration, use and impacts of large recursive resolution services) and declared them to be non-topics, saying: "the working group will not attempt to resolve disagreements on these topics, and will require full consensus on any statements regarding these areas".

Consensus, or rough consensus nevertheless, is one of the more delicate concepts of the IETF (see also RFC 7282). "Full consensus" was an alien concept to the IETF and many people complained. Furthermore, the list of topics in the scope was blown up in the new charter version, with nobody expressing consent to the extended list during the session. While there was one notable rough consensus, in that the extended list was too long, there was a lack of moderation during the session which prevented any progress. Immediately after the BoF, former IETF Chair Jari Arkko, set it straight offering the following narrow scope proposal:

** write a specification that allows the discovery of and the use of DNS servers, with something like adaptive DNS as a starting point*

- including general security analysis, privacy impacts analysis, and resistance to pervasive surveillance analysis regarding this proposal

** use standard IETF WG process*

Notable changes (1)

ORIGINAL	LATEST DRAFT
Specific initial areas of focus include: • Resolver discovery • Expression of resolver policy • Query routing in the presence of resolver choice	• Communicating configuration between the network, operating system, and applications • Discovery of resolvers and their capabilities and behaviors • Query routing in a multi-resolver environment • Multiple non-equivalent query paths, such as split-horizon DNS or geo-sensitive answers • Local DNS caches (e.g. partitioning, use of stale records) • Resilience and fault-tolerance (e.g. single points of failure) • Support for debugging and analysis • DNS Push (accepting responses to queries that have not yet been issued) • Ossification and evolvability

Given the failure of the BoF, the next start of a potential WG might be IETF 107. IETF Chair Alissa Cooper remarked during the session that the BoF had not managed to get beyond a situation in which two camps were fighting each other.

The pro-DoH/Web camp argued for example that there was no need for a new WG as there was the DNS WG (David Schinazi, Google, formerly Apple). Patrick McManus (Fastly, formerly Mozilla) said the charter lacked specificity. Martin Thomson (Mozilla) warned against an "octopus-like" WG Charter. From the "other" camp, Chris Box (BT) said the narrow list might be a little bit too narrow, but the long list was too long. Some, like Ralf Weber (Akamai), called for a more structured discussion in a potential WG. Given the BoF was a failure, the interested parties have one more attempt.

Hitchhiker's Guide to QUIC?

Just as much as DoH and possibly ADNS (ODNS) will change the DNS, QUIC is going to change transport and take a bite of the traditional transport by TCP. QUIC (Quick UDP Internet Connections) is a new internet transport protocol, encrypted-by-default, that tries to make transport faster, more secure and aims to replace TCP and TLS on the web according to some.

Currently the numbers reported on QUIC usage are between 2.6 and 9 percent. Come December the QUIC Working Group will proceed to Working Group last call for two of its core documents:

draft-ietf-quic-tls-24 Using TLS to Secure QUIC

draft-ietf-quic-transport-24 QUIC: A UDP-Based Multiplexed and Secure Transport

During a timely talk at the Transport Area Open Meeting, QUIC Co-Chair Mark Nottingham said that the group would allow for an extended phase to comment on the new transport protocol.

Though it has lasted longer than proponents originally expected, the QUIC WG has perhaps been one of the most intense WGs, with three regular WG meetings during IETFs – each having two sessions – as well as meeting between IETFs at three annual side meetings. Since interoperability tests are getting better and the drafts are stabilising, the current plan is to bring these proposals to the IESG in mid-2020, Nottingham said. Other documents will follow more or less quickly on the heels of the core documents, Nottingham said, including:

- draft-ietf-quic-recovery-24 QUIC Loss Detection and Congestion Control
- draft-ietf-quic-qpack-11 QPACK: Header Compression for HTTP/3

and documents on operational issues like

- draft-ietf-quic-invariants-07 Version-Independent Properties of QUIC.

The WG is also already working on version 2 of QUIC, but Nottingham said that the focus of the group for now was to ship the core protocol.

Like DoH, QUIC seems to underline that a re-design of the net is being driven by what might be called the web companies. The reaction from many of these companies to the QUIC development illustrates this shift. Nottingham brought long lists of extensions and applications, either already taken up or waiting in the wings to make it to the WG. QUIC, he said, will be the new hot topic.

Extensions considered by the WG are:

- QUIC Load Balancers (duke-quic-load-balancers
- QUIC Version Negotiation (schinazi-quic-version-negotiation)
- QUIC Datagrams (pauly-quic-datagram)
- Loss Bits (ferrieuxhamchaoui-tsvwg-lossbits) (future document)

Nottingham also reported about a growing number of applications that have already expressed interest to use QUIC (like WebTransport, vvv-webtransport-quic, proxy/tunnelling, e.g., draft-schinazi-masque), as well as DNS and Netconf). The respective work will be done in other WGs, according to Nottingham. People are already working on "pluginised QUIC" as well as QUIC for Satcom.

With the considerable number of proposals related to QUIC, two ideas were discussed during the meeting. One was to allow for a dedicated QUIC Dispatch Group, that would hear all QUIC-related drafts and send them off to the responsible WG.

Another proposal to prepare for the QUIC deployment was made by IAB Chair Ted Hardie, who said it might be time to prepare a "Hitchhiker's Guide to QUIC" to allow implementers to get it right from the start – something that for older protocols was only done after standardization. Hardie pointed to the Hitchhiker's Guide to SIP as a model.

Working groups

DNS - Yet another edition of .internal and a final solution for aname, bname, cname

Two drafts in particular caught the attention of DNS experts during the two DNS WG sessions.

One is the attempt to solve the "ANAME, DNAME, CNAME"-issue in one go. Ben Schwartz (Google) presented a draft, whose declared goal is to allow a client to query a name and get the "full set of information" needed for connecting to a service. According to the author, the new record will provide a whole bundle of information instead of an IP address only. It acts like CNAME, but could sit at the APEX as an alias.

Schwartz's draft is another, according to the author, more complete answer to the request to have SRV or a functional equivalent implemented for HTTP and attempts for delegation using ALTSVC, ANAME and ESNIKEYs. The problem with the many earlier approaches had always been that they resulted in incompatibilities while at the same time only solving one part of the functions respectively.

The WG mostly welcomed the draft (e.g, David Schinazi from Google, Tommy Pauly from Apple, Brian Dickson from GoDaddy, Ondrej Sury from ISC), with additional questions to be discussed. Schwartz himself asked for comments on two questions, namely how to balance ESNI strictness against reliability and misconfiguration. Schwartz explained that the current requirements prevent fallback from ESNI to non-ESNI unless the server specifically indicated that it was allowed. Another question was whether there was a need to limit the chain length.

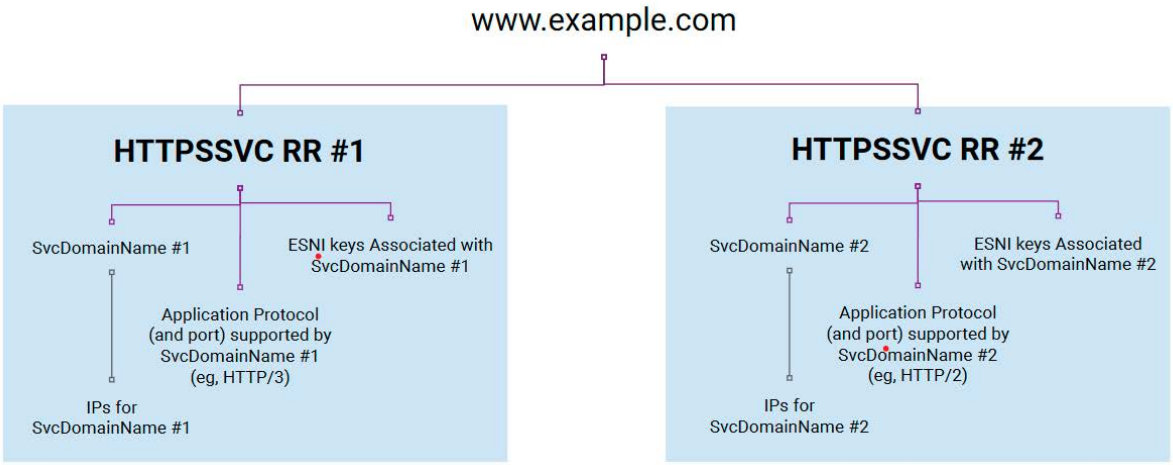

After first presenting an HTTPS solution, HTTPSSVC, the authors now also provide a generic solution, SVCB. The new records will allow for the delegation of an operational authority for an origin within the DNS to an alternate name.

According to the draft text SVCB and HTTPSSVC will allow for the provision of authoritative service endpoints, along with parameters associated with each of these endpoints "while acknowledging different responses to the record request from different hosting environments or CDNs (multi-homing) and while enabling CNAME-like functionality at the zone apex (example.com) for participating protocols". In essence, the proposal, as explained by Schwartz, will allow multi-CDN hosting with encrypted ESNI.

Schwartz asked for further recommendations from the server operators on the graph about server behaviour, authoritative as well as recursive.

The current server behaviour is described as follows in the draft:

When processing an SVCB response from an authoritative server, add it to the Additional section (unless it is the Answer).

If all records are in ServiceForm, resolve A and AAAA records for each SvcDomainName (or for the owner name if the SvcDomainName is "."), and include all the results in the Additional section.

Otherwise, select an AliasForm record at random, and resolve A, AAAA, and SVCB records for the SvcDomainName. If the SVCB record does not exist,

add the A and AAAA records to the Additional section. Otherwise, go to step 1, subject to loop detection heuristics.

All DNS servers SHOULD treat the SvcParam portion of the SVCB RR as opaque and SHOULD NOT try to alter their behavior based on its contents.

When responding to a query that includes the DNSSEC OK bit ([RFC3225]), DNSSEC-capable recursive and authoritative DNS servers MUST accompany each RRSet in the Additional section with the same DNSSEC-related records that it would send when providing that RRSet as an Answer.

Before the new record types can be requested at IANA, the draft has to be stabilised, the WG concluded.

Another much-discussed proposal is another go for an "internal"-zone, which failed to receive support when the IETF tried to get .internal or .home. Interestingly, it is two authors from ICANN, Roy Arends and Ed Lewis, who put the new proposal for a non-ICANN-delegated internal address zone on the IETF table. To avoid the need for delegation, the IETF could chose an unassigned alpha-2 code from ISO list 3166-1, which lists country codes. According to Arends, from all potential alpha-2 codes there were a number that were neither assigned nor could be expected to be assigned in the future. From the list (see graph) Arends and Lewis propose to select .zz.

The short form and lack of semantics was an advantage of the label according Arends. Not everybody agreed. .internal draft author Warren Kumari said that the lack of semantic meaning could result in confusion of users. Petr Spacek (CZ.NIC) pointed out that collisions would still happen over time, with companies merging, etc. Nevertheless, what feels like a majority of participants in the DNS WG was supportive of the idea. The draft still has to be discussed before a decision can be made to take it up as a WG document.

Other drafts currently worked on in the DNSOP are:

- **Message Digest for DNS Zones, draft-ietf-dnsop-dns-zone-digest, Duane Wessels**
- **Extended DNS Errors, draft-ietf-dnsop-extended-error, Wes Hardaker**
- **DNS Transport over TCP - Operational Requirements, draft-ietf-dnsop-dns-tcp-requirements, Duane Wessels**
- **Interoperable Domain Name System (DNS) Server Cookies, draft-ietf-dnsop-server-cookies, Willem Toorop**
- **Related Domains By DNS, draft-brotman-rdbd, Stephen Farrell**
- **Operational recommendations for management of DNSSEC Validator, draft-mglt-dnsop-dnssec-validator-requirements, Daniel Migault**
- **Avoid IP fragmentation in DNS, draft-fujiwara-dnsop-avoid-fragmentation, Kazunori Fujiwara**

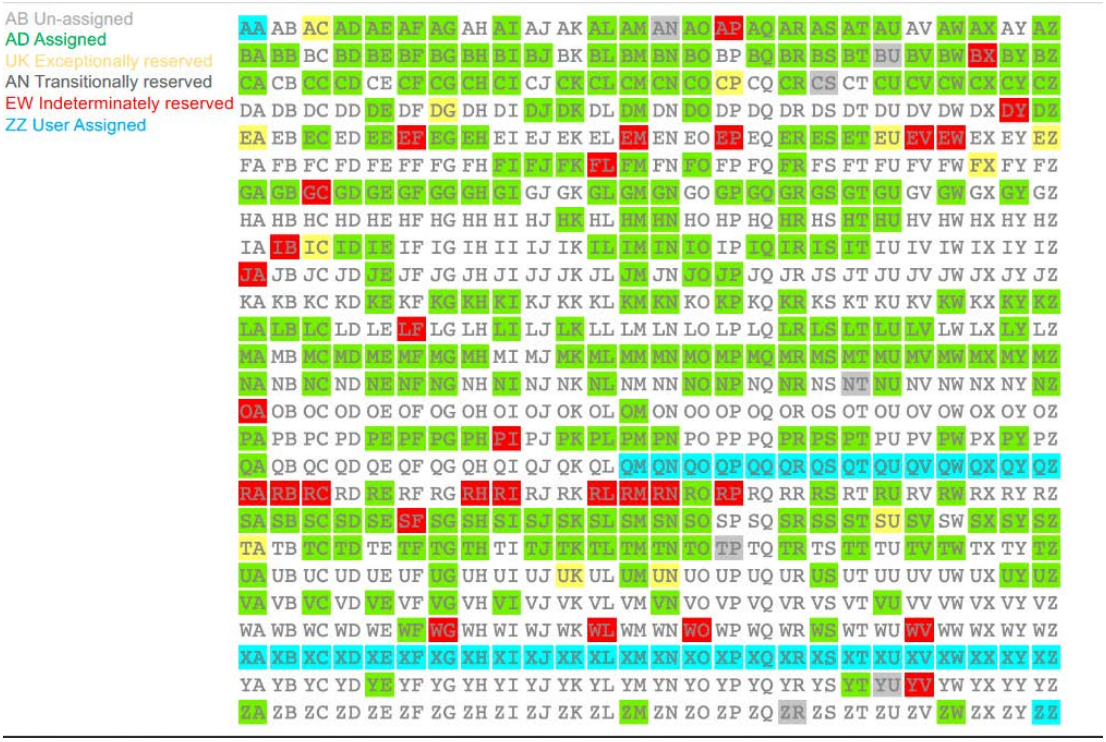

RegEXT – Rubber-stamping registry-registrar documents

The RegEXT Working Group once more revived well-known discussions over its work. One constant concern is that the business practices of some companies/organisations are given the "IETF standard" seal. During the Singapore session, the draft on bundling registrations, that has been pursued by CNNIC authors for many years, got another push-back as the IESG had obviously recommended to make it an informational document only. Against the arguments by Ning Kong, a consultant for CNNIC, that the authors did not want to settle for a merely informational status, the RegEXT Co-Chair Jim Galvin said that the document could not proceed further if the authors did not accept.

Yet another problem was once more highlighted when the RegExt WG chairs said that since the Registry Data Escrow Specification had only received two responses, there were not enough comments to send the document to the IESG. The WG has experienced considerable problems in garnering enough interest from people to review the documents for quite some time. The reason is certainly that those that follow the specific standardization efforts are only a very small group of registry operators, as well as a small number of registrars who can afford to follow the work.

Furthermore, as Galvin mentioned when talking about the escrow specification and also the domain name registration data objects mapping document, the respective practices are obligations for ICANN-contracted parties. Therefore, when standardizing the practices this must not make the current approach incompatible. This clearly demonstrates that a deviation from contractual clauses is not welcome.

Consequently there was considerable push-back against a proposal by Galvin himself to bring more than a dozen different practices used in ICANN registry-registrar reporting to the WG for standardization.

Alex Mayrhofer warned that the practices were mere practices ruling B2B relationships. There was a lack of public interest for the internet as a whole. Therefore he thought this effort would be the ultimate rubber-stamp action and an abuse of the IETF. Richard Wilhelm (Verisign) also warned that the TechOps community at ICANN had no agreement on some of the practices. Bringing it to the IETF without really involving the relevant community in the debate would override the TechOps processes.

Finally Mario Loffredo (Registro .it) presented the progress of three RDAP related drafts and had to face questions about a proper privacy consideration section, particular in the draft about the RDAP reverse search.

Two proposals for new work were briefly discussed briefly. One is an older proposal by ICANN, which wants to see the Trademark Clearing House operations standardized.

The other is a proposal by Mayrhofer to standardize a feature allowing for domain suggestions to registrants, which was said to be unnecessary, since big registrars already had their private solutions.

GenART Dispatch: organisational issues

The newly-established GenArt Dispatch was established to deal with a number of existing proposals that deal with the very organisation of work in the IETF. In typical "dispatch"-style, the group will weigh the proposals and decide how they should be dealt with. The documents presented in Singapore were all decided to be best dealt in an AD sponsored draft document.

In Singapore GenART discussed a straightforward proposal from Joel Halpern, which fervently rejects the growing practice that the IESG would pass documents in the IETF workstream without the documents having reached consensus. The practice was a door-opener for abuse, one participant claimed. Halpern argued that the original RFC had not envisaged the various streams that had been established (IAB, IRTF, besides IETF). The document "proposes that the IETF never publish any IETF stream RFCs without IETF rough consensus." The WG seems to be fine with this, and IETF Chair Alissa Cooper was asked to take this up.

Another RFC document-related proposal was Martin Thompson's (Mozilla) call to say goodby to the expiration of draft documents.

A bigger discussion for the group will be the question of equal participation of remote participants in starting to recall an Area Director. In reviewing that document, the barriers for starting such a recall will

also be lowered, according to the current draft by Subramaniam Moonsamey and John Klensin.

Blurred lines: the relation of IETF and her research sister IRTF

Colin Perkins, new head of the Internet Research Task Force (IRTF), research sister body to the IETF, included an in-depth discussion about the relation of the two bodies on the agenda in Singapore. Five years ago, RFC 7418 tried to explain the IRTF's role to participants of the IETF who brought work there. This time the focus was more on how the lines between the two organisations had been blurred by the IRTF, which is increasingly gearing up to IETF processes.

Research going on in the Research Group is not requested to be documented in RFCs, and the IRTF does not need to follow the IETF rough consensus concept for adopting documents either. Instead, research papers are published as they are (some good research papers every year receive the applied network research prize, see below), and some like Stephen Farrell (Trinity College Dublin and member of the IAB) said that a lack of consensus was healthy in research.

Possible changes for a more independant and research-focussed IRTF mentioned during the session were the co-location of IRTF meetings with other research conferences and the establishment of relations with other organisations (including, as one participant offered, the ITU for example).

Former IRTF Chair Aaron Falk described several types of relations between the two sister bodies observed "in the wild":

1. Sometimes work was brought from an IRTF WG into a dedicated WG, as a "one sho"", like the IETF Anima WG which was spun off from the Network Management Research Group (NMRG). Another nice example here is the work around IoT, as the IoT RG has contributed to several working groups.
2. Another type of relationship has evolved between the Crypto Forum RG which stepped up as an expert body to select secure ciphersuites for IETF protocol after it had become clear that NIST had been compromised by the NSA. Now the CFRG is the standing body to further advise the IETF on ciphersuites. Another example of this type is the Internet Congestion Control Research Group (ICCRG) which has become the expert body and for congestion control proposals for the Transport Area WG.
3. The third variant, mentioned by Falk, are research groups that do basic research in new technology areas, like the recently established Quantum Internet Research Group.

According to Falk, a question for the current deliberations is if there is a need to document the criteria and conditions for a successful transfer (and if more transfer was a stated goal). Just counting RFCs by IRTF contributors (for scientists this was sometimes non-gratifying as RFCs in some research institutions are not accepted as a regular publication), or RFCs spun-off to the IETF from IRTF work, was too narrow. "There are more metrics for success than the relationship with the IETF", said Melinda Shore Principal Security Architect at Fastly and Chair of IETF and IRTF groups.

Barriers and incentives for researchers to contribute to the IRTF (and IETF) was also discussed in relation to a presentation by Marie-José Montpetit.

Nice Work at the Privacy Enhancements and Assessments Group (PEARG)

Quite a nice example of how the IRTF research-oriented work can support the development work of IETF participants was showcased in Singapore by the still relatively new PEARG. Both a proposal on documenting evolving fingerprinting practices as well as a privacy framework for logging in networks can inform actual protocol development (as well as operational) work by IETF engineers. Another document being considered is based on observations about the de-anonymization risks which stem from inference avalanches enabled from machine learning training data sets. See more here, here and here.

IETF107 will be held in Vancouver from 21-27 March 2020

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 8 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. Full membership is open to organisations, corporate bodies or individuals that operate a country code top level domain registry.

Rate this CENTR Report on IETF106

(Thank you for your feedback!)

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