Council of European National Top Level Domain Registries

Report on IETF 91 Honolulu

10 - 15 November 2014

Table of Contents

Discussion the IANA transition: No changes necessary for now

Procedurally the IETF is close to finalize its contributions to the IANA Stewardship Transition Coordination Group (ICG). The contribution will come in regular format – as an RFC document. The Honolulu meeting allowed for a last face to face discussion about the IETF answer to the NTIA call before the IETF files its RFC-answer to the ICG.

The document for which the WG last call ended during IETF week meanwhile was once more vetted on the IETF mailing list after the meeting. It has now been sent on to the IESG, according to regular RFC procedures and the IESG asks for final statements before December, 15th.

Once the IESG passed and the RFC editor edited it and assigned an RFC number to it, it will be passed to the ICG which expects contributions from the various IANA „customers" by January, 15th.

Services provided by the ICANN operated IANA to the IETF include the managing of the protocol number registry and the .arpa TLD.

Negotiating wish lists?

There obviously was still some uncertainty about how the final IANA stewardship proposal will be further developed once the different communities have decided about their preferred IANA models.

Participants in Honolulu considered the possibility that the IETF Administrative Oversight Committee (IAOC) might have to further „negotiate" a joint proposal with the other IANA customers (RIRs, ICANN). Bob Hinden, IAOC member, warned that „negotiating" the community's position would be difficult at best.

Long time participant Scott Bradner recommended to have the next steps seen through by the IETF and IAB chairs who were in charge to represent IETF interests, with the IOC as a backup. The RFC, once formally passed, should be the cornerstone of the IETF position. At the same time IOC member Tobias Gondrom found a „wish list" that could allow the IAOC members (and or IETF/IAB chairs) to give up what is least, and defend what is most important.

Russ Mundy (Sparta, and IGC member) explained at Honolulu that he did not expect the IGC to merge the different proposals or make them into one compromise text in case of differences. But IGC member Alissa Cooper, also IGC member, disagreed saying that the ICG would send only one, and not three proposals to the NTIA.

Lack of trust in ICANN

While the IETF/IAB in a similar way as the RIRs support minimal changes in how IANA is currently operated, questions raised by several participants during the IANA Plan WG meeting touched back-up plans for a potential failure of ICANN as the IANA operator. ICANN might „go berserk" after losing its former controllers at the US administration, several IETF 91 participants were concerned.

Problems mentioned during the discussion:

- ICANN could ask for payment for providing IANA services to the IETF
- ICANN could assume ownership of IANA.org (creating a problem once there is a change of operator for the protocol number registry)
- ICANN could question the IETF role as authority for .arpa (or its authority over special domains)

Randy Bush, Internet Initiative Japan warned that ICANN was one of the „least transparent and least accountable organisation" and had „more lawyers that we have protocols".

The possibility to switch operators, a certain planning for cost in case of switching and finally the question of ownership of iana.org were dressed. Brian Dickson, (working for Twitter) recommended to consider the IETF's possibility to control ICANN and retain some core tasks at the IETF (root zone).

Detaching ICANN from politics of names?

Given the level of distrust in ICANN, and more so the distrust in the political stability of the names part, there was at least one call to consider detaching the IETF's tasks more from ICANN. ICANN could end up „at the centre of a cyber world war by necessity because the name is a control point
for the Internet", Phil Hallam-Baker, Comodo said. It would be good if the IETF would be completely decoupled from such arguments. Instead of owning IANA, stability of the IETF registries were paramount.

Separation of functions or a break away from the current IANA model for the time being is not discussed, but A separation of function is also not part of the draft contribution of the IETF.

„Draft Response to the Internet Coordination Group Request for Proposals on the IANA protocol parameters registries" - link

The IETF lists its tasks related to IANA as

- special use registries with regard to domain names [RFC6761]
- specification of DNS protocol
- specification of minimum requirement for root servers [RFC2870], currently under review
- consultations with the RIRs over evolution of the routing architecture
- policy role with regard to IP address space and AS number space [RFC7020],[RFC7249], unique local addresses (ULAs) [RFC4193]
- maintenance of sub-registries for special Ipv4- and v6 assignments [RFC3307], [RFC5771], [RFC6890]
- developing standards that might influence RIRs and services providers

The obligation to allow for a potential switch to a new IANA operator is addressed in the final document. With regard to potential IPR issues over IANA, the document solely underlines the need of the protocol registry to be in the public domain.
Core aspects in the RFC are:

No major changes are required. Over the years since the creation of ICANN, the IETF, ICANN, and IAB have together created a system of agreements, policies, and oversight mechanisms that already cover what is needed. This system has worked well without any operational involvement from the NTIA. Therefore, no new organizations or structures are needed.

(..)

The IETF community is quite satisfied with the current arrangement with ICANN. RFC 2860 remains in force and has served the IETF community very well. RFC 6220 has laid out an appropriate service description and requirements.

However in the absence of the NTIA contract a few new arrangements may be needed in order to ensure the IETF community's expectations are met. Those expectations are the following:

The protocol parameters registries are in the public domain. It is the preference of the IETF community that all relevant parties acknowledge that fact as part of the transition.

It is possible in the future that the operation of the protocol parameters registries may be transitioned from ICANN to subsequent operator(s). It is the preference of the IETF community that, as part of the NTIA transition, ICANN acknowledge that it will carry out the obligations established under C.7.3 and I.61 of the current IANA functions contract between ICANN and the NTIA [NTIA-Contract] to achieve a smooth transition to subsequent operator(s), should the need arise.

Furthermore, in the event of a transition it is the expectation of the IETF community that ICANN, the IETF, and subsequent operator(s) will work together to minimize disruption in the use the protocol parameters registries or other resources currently located at iana.org.

On the question of jurisdiction, the IETF draft states: „This mechanism is global in nature. The current agreement does not specify a jurisdiction."

DNS privacy solutions - solutions shaped by different interests

New Working Group dprive started

The work on DNS privacy has gotten off with a new dedicated working group dprive officially started just in time for a first meeting in Hawai. While the pace of the dns privacy advances seemed to have slowed down during the IETF in Toronto the dprive WG now is sifting through the list of alternative proposals that have been put on that table by various parties. The charter adopted for dprive limits the first work to tackling privacy issues between the stub server and the resolver.

Basically, there are two different strings of work, as Stephane Bortzmeyer, AFNIC and author of the problem statement, explains. One designs for data minimization, the other is introducing encryption for DNS traffic.

Strip requests of unnecessary data

Data minimization for which Bortzmeyer has presented a draft in DNS OP wants to send only what is necessary in a query going up the DNS tree. Full requests (www.example.com) does not have to be sent on from a recursive resolver to the the root, as the answer anyway will result in a pointer to .com.

While sending „less than before" and only what is „absolutely necessary to receive an answer" is one trend in the discussion, there is another that goes in the opposite direction. According to the proposal by several authors from Google and Akamai is to send more information in fact to allow for localized answers. The draft that would introduce „an EDNS0 option to allow Recursive Resolvers, if they are willing, to forward details about the origin network from which a query is coming when talking to Authoritative Nameserver". The draft was not presented in Hawai, though.

It is clear that minimization is not favoured by some business models and companies who are interested in the analysis and further (commercial) use of DNS „big data".

Encrypting DNS traffic

With regard to encrypting DNS traffic there is a still growing list of proposals, one of them just being put to running code test by a group of researchers from the University of Southern California together with VeriSign Labs: DNS over TLS (which will include a shift form UDP traffic to TCP traffic for the DNS, a proposal on this has been presented during DNS, see below).

For the DNS over TLS variant presented by the VeriSign Labs/USC group clients and servers will signal their preferences in the header (see draft proposal):

„Clients and servers indicate their support for, and desire to use, DNS-over-TLS by setting a bit in the Flags field of the EDNS0 [RFC6891] OPT meta-RR. The "TLS OK" (TO) bit is defined as the second bit of the third and fourth bytes of the "extended RCODE and flags" portion of the EDNS0 OPT meta-RR, immediately adjacent to the "DNSSEC OK" (DO) bit [RFC4033] (…)"

In order to „minimize" data exposure the authors recommend to send a standard dummy query to initialize the TLS-communication in the form of:

„(RD=0, QNAME="STARTTLS", QCLASS=CH, and QTYPE=TXT ("STARTTLS/CH/TXT")"

After successful TLS negotiation, the connections can be encrypted and would be protected from eavesdropping. Given

that only months ago, the public nature of the DNS has been said to be a given and difficult to change, makes this quite a development.

Time costs tested

But there are concerns, not the least by large DNS services providers who warn especially against added latency (due to additional round trips) and more bandwidth hunger for running DNS. One large DNS provider said, costs would be going up significantly, additional hardware would be necessary.

VeriSign Labs and USC have tried to calculate time „cost" and said they were confident these were reasonable. Tests and results are explained at length in the research paper which states that „TCP TCP set-up and DNS resolution would be less than 1 ms, TLS set-up is more expensive 8 or 26 ms, varying by implementation. Resumption would be ten times faster than full TLS set-up by the way. Summing up, VeriSign Labs/USC calculate that „end-to-end latency from TLS to the recursive resolver is only about 9% slower when UDP is used to the authoritative server, and 22% slower with TCP to the authoritative."

They also reject the notion of „new hardware necessary" in their research paper, saying that „a large recursive resolver may have 24k active connections requiring about 3.6 GB additional RAM." The Group nevertheless underlined that overhead needed to be minimized through during implementation through „query pipelining, out-of-order responses, TLS connection resumption, and plausible time-out".

Solution space

DNS over TLS is by far not the only solution and has, by the way, been around for some time before dprive. It has already been an option in Nlnet Unbound servers, using a different port, though.

Two other presentations were given in Hawai. Paul Hoffmann (Internet Mail Consortium), briefly introduced two three potential approaches for encrypted DNS:

- wrapping of DNS queries in http requests (and answers as http responses, eg https://8.8.8.8/.well-known/dns-in-https/TN4AAAABAAAAAAAAB2V4YW1wbGUDY29tAAABAAE=) as a request uri
- plain DNS over TCP using ALPN for transport (port 443 instead of 53)
- or using a completely new port.

Hoffmann at the same time said DNS over TLS could be a rational choice. The potential advantage of his proposals might be that middleboxes inclined to reject TLS over port 53 could be outsmarted.

Another written proposal is to be prepared further by Philip Hallam-Baker who seems to favour UDP, but not DTLS on top of it. He declared 100 percent connectivity, performance equal to the existing, stateless transaction with no public key, bypass interference, elimination of amplification and relay attacks, low footprint and low complexity, confidentiality and the enabling of what he called „curated DNS" the cornerstones of his proposal.

Much older, and in fact pre-Snowden, is the suite of DNSCrypt and DNSCurve which encrypt DNS traffic from the stub resolver to the recursive and from the recursive to the authoritative using Elliptic Curve algorithms and do not use a centralized, hierarchical key distribution management. Originally designed by Dan Bernstein since 2008 as an answer to concerns over DNSSEC the suite has not been adopted by DNS operators.

VeriSign Labs/USC while acknowledging that the protocols address the same privacy concerns as DNS over TLS criticize that it would not address DdoS-attacks, something the DNS over TCP and TLS wants to solve alongside the privacy issue.

Finally there seems to be more proposals to come, one has been posted as a draft document by Hosnieh Rafiee, engineer at Huawei Technologies Düsseldorf in Munich: „CGA-TSIG/e: Algorithms for Secure DNS Authentication and Optional DNS Confidentiality. Rafiee had been unable to get a visum, but might be presenting next time.

Next steps – requirements or matrix document?

A decision on the way forward with regard to the different proposals is still open, the WG for now started considering another general document, potentially a requirements draft or matrix to allow for a comparison of pros and cons of the various solutions. Certainly VeriSign Labs/USC are advanced as they can point to (rough) running code. Allison Mankin from the Labs and John Heidemann from USC also presented the broader research document, that goes some way to elaborate on requirements and how the different proposals addressed them – the USC/VeriSign group has a slight bias toward their own proposal.

IAB: Let's abandon clear text on the net, please!

The Internet Architecture Board at the end of the IETF week published a statement on confidentiality in the net that recommends to protocol designers and implementers to use encryption across all layers of the protocol stack, from transport to applications. The statement adds to the earlier draft document by Security Area Director Stephen Farrell (Pervasive Monitoring is an Attack), but is shorter and a little more straightforward in its recommendation: „The IAB now believes it is important for protocol designers, developers, and operators to make encryption the norm for Internet traffic. Encryption should be authenticated where possible, but even protocols providing confidentiality without authentication are useful in the face of pervasive surveillance as described in RFC 7258."

Asked why the IAB decided for „should" authenticate and allowed for „exceptions" from the encrypt-norm, IETF Chair Russ Housley said that there had been a long debate at the IAB. The reason to allow for flexibility was that some protocols (like secure neighbour discovery, send) would otherwise stop to work. Housley said the IAB would publish guidelines following the short statement to explain the rationale and expectations. How much effect will the statement have? At least it is interesting to note that Housley clearly said, he hoped the HTTPBIS WG would reconsider their stance on unencrypted HTTP as an option for HTTPBIS.

Third-Party Management of DS keys

In a bar BoF organized by Olafur Gudmundson (Cloudflare) a potential EPP extension (or some restful – weird – solution) for third party DNS providersas explored. The Participants who were all favouring the third-party DS management option came from registries (Afilias, VeriSign) and Service Providers (Cloudflare, Akamai, Nominum).

In order to provide DNSSEC they want to be able to manage the DS records for their customers, instead of having to go through the ICANN accredited registrars who are the first stop for signing a domain for an end user. While there was an appreciation that some registrars might see the initiative as bypassing them (and taking away business from them), there was also some expectation that some registrars might not be interested in managing the DNSSEC protection for customers themselves as it included to be up to date with key roll overs. A problem acknowledged briefly was customer support in case of failure and liability.

The main technical avenue discussed was an extension in EPP to be drafted and tabled at the DNSOP WG (not the EPP extension WG that had heavy discussions in Hawai). Third party providers recommended to keep the technical solution simple (some third party providers argued to fully implement EPP was an unnecessary burden, and they would prefer to just present some „token" to the registry to legitimize them as DNSSEC manager for the customer.

ICANN would be presented with a technical solution when it was ready, next steps will be preparation of a draft, A non-WG mailing list can be found here.

China's answer to NetMundial

Two Internet related conferences in China and Hong Kong have spurred side discussions about a possible Chinese attempt to „grab" their own root server.

The first event was a high level political conference in Wuzhen, China, on November 19-21. The „World Internet Conference (Wuzhen Summit)" had been organized by by the recently established Cyberspace Administration of China

and the People's Government of Zhejiang. It featured Chinese officials and representatives of large companies like Tencent, Alibaba. ICANN CEO Fadi Chehade faced critical comments over his participation.

The title of the conference, the issues it addressed („unprecedented challenges of unbalanced development, increasing threats to cyber security, uneven distribution of critical Internet resources", its declared goals including „to promote the development of Internet to be the global shared resources for human solidarity and economic progress" and „open the new historical chapters in which the development of the Internet") explains at least some of the uneasiness in Western countries. The workshops titles included: „security and cooperation in Cyberspace", „countering cyber terrorism with enhanced international cooperation", and „reform into the future: building global internet governance ecosystem".

Reservations about the conference were indirectly illustrated by absent speakers (for the high level dialogue on „an interconnected world shared and governed by all" defectors included: Lu Wei, Minister of Cyberspace Administration of China, Kevin Rudd, former Prime Minister of Australia, Sato'Sri Ahmad Shabery Cheek, Minister of Communication and Multimedia, Malaysia and Daniel Sepulveda, Deputy Assistant of Secretary of State of the US.)

There is also a high-level dialogue on „an interconnected world shared and governed by all", a celebrities dialogue and an invitation only session on „constructing a peaceful, safe, open and cooperative cyberspace". The program could be interpreted like a Chinese follow-up to the ITU Internet related events, a Chinese version of the IGF or even a Chinese answer to the NetMundial.

A Wuzhen Declaration was prepared by the hosts, but not published in the end. According to observer Izumi Aizu a Chinese official told a US official, the hosts would not push to call it a declaration because the US had not agreed. The conference organizers have mentioned in their program that they want to establish the Wuzhen Summit as a regular event.

Not seven more root servers, but millions

The second event is a much smaller workshop co-hosted by ZDNS/BII and the Chinese Network Information Center (CNNIC) addressing the issue of „root zone availability". Following up to presentations about spreading the root zone at the DNS Working Group meeting in London the workshop now is dedicated to two draft documents and chaired by Warren Kumari (Google) and Paul Vixie (former BIND CEO). The proposals are the one of Kumari and Paul Hoffmann on the concept to reduce latency for root zone requests through a local cached version of the root zone file http://tools.ietf.org/html/draft-wkumari-dnsop-root-loopback-00. It was presented in Honolulu in the DNSOP session and was very well received this time for its additional feature of keeping some nasty traffic from going up to the root. One request made at the DNSOP session was to support the rationale by figures on how long connection times to „local" (anycast) root servers currently were.

The second proposal is from Paul Vixie, Xiaodong Lee and Ziwei Yan on „How to scale the DNS root system". The workshop according to the program will explore the differences and commonalities of the two proposals „with an eye towards revising both drafts".

Vixie in a blogpost meanwhile has apologized for his first proposal to add 7 or so additional root servers. He was „starkly opposed to adding more traditional root name servers". Instead the modified proposal by him and his co-authors would based on the A112 concept and would allow everybody to run his own root server either locally or using the AS112 unknown anycast equivalent for the DNS Rootzone. Vixie wrote:

„The problem with the current root name server system is not that there are twelve server operators, but rather, that there are not millions of server operators. I am working toward an Internet with millions of root name servers and name server operators, not an insignificant change from thirteen to twenty."

IANA would need to make minor adjustments:

- to set aside two Ipv4 and Ipv6 blocks to propagate an additional version of the root zone
- to create an otherwise-identical copy of the DNS root zone, having different apex NS records, but signed with the same root zone key
- to create some name server names (X.ROOT.IANA.NET and Y.ROOT.IANA.NET, for examples) each having one address in an IPv4 prefix and one address in an IPv6 prefix
- to operate publication servers capable of serving millions of "stealth secondary" root name servers

- operate a subscription service whereby these servers can ask for and receive NOTIFY messages concerning root zone changes

Working Groups

EPP Extensions: If you standardize extensions, please standardize mine!

The EPP Extension WG saw an interesting discussion about whether or not to make the extension proposed in the WG standards track documents, instead of just informational documents or individual submissions. The discussion also shed some light on the delicate relationship between IETF standards and ICANN contractual arrangements.

WG Chair Jim Galvin (Afilias) recapped that the WG originally had expected to make the „Extension Registry for the EPP"-document a standards track document, while the status of the documents on individual extensions were open to discussion. Each could be either standard track, informational RFC or an individual submission. The difference was that the first two needed a consensus in the WG (making the resulting document a recommendation to the operative community). Individual submissions were an expression of parties documenting their practice.

Galvin acknowledged that judging consensus of the EPP extensions in the WG was complicated by the fact that only few users of the specifications did attend the IETF meetings, namely the domain registries while the registrar community was not represented. Yet the Area Directors put forward, that lack of review in the IETF WG could be made up by review in the ICANN community. The authors of the respective drafts only had to document that review.

That certainly leads to one fundamental question: can documents developed and/or reviewed elsewhere get an IETF standards track stamp? Work has indeed been brought by other communities to the IETF (see for example oauth, or the major audio-code for WebRTC). According to the area directors and IAB member Andrew Sullivan a standard or informational RFC were not too different. Having passed such an IETF RFC would not block anybody from bringing new work (a new solution) to the same problem to the IETF.

Chris Wright, AusRegistry, on the other hand objected to upgrading what had been planned as individual submissions or informational documents to standards track. As his company had their own solutions to some of the issues addressed in the draft EPP Ext documents, he was afraid putting them to standards track could fire back at him and his registrars because ICANN would force them to adopt the „IETF standard" solutions in future contracts.

Would he have known that the proposals would be made standard documents, he would have objected much more in the first place, Wright said. It was not fair, to use a consensus developed outside of an IETF WG. Instead discussions had to be started from scratch – and different solutions had to be compared to make a choice. His answer to the potential upgrading of the launch phase mapping extension was that he would make the effort of having his own solution passed as an IETF standards document, too.

There are three active working documents:
1. EPP Extension registry by VeriSign, it introduces a procedure for the registration and management of EPP extensions and it specifies a format for an IANA registry to record those extensions. The document planned as a standards track document will only be informational – as it does not specify a protocol, it has been submitted to the IESG for publication.
2. Launch Phase Mapping of EPP by VeriSign, Cloud and CentralNIC, proposing an extension for the special needs of a registry when launching a new TLD – The document has been developed in fact outside of the IETF until version 12, and then changed to a WG document; according to the WG Chair is is now on the standards track.

Three documents were briefly touched:
3. TMCH-SMD (Trademark Clearing House Signed Mark Data) standards track, also confirmed by Galvin as possible on standards track (with launch phase) because of ICANN review
4. Internationalized Domain Name Mapping Extension for EPP , it has been expired, but could go to standards track, too, according to Galvin
5. Key Relay Mapping for the EPP allowing for the transfers of DNSSEC keys, from SIDN Labs –
Scott Hollenbeck, VeriSign recommended using the existing transfer option for this
New drafts presented were
6. Bundling of domain names (Chinese simplified, traditional), proposed by Ning Kong, CNNIC, who was asked how much the extension could be used independently from CNNICs bundle policy.
7. EPP Service Messages Extension, according to author Alexander Mayrhofer (nic.at) allows to convey additional messages from the registry to the customer, the registrar. Situations listed in the draft include for example, notifying a registrar about the transfers of additional objects attached to a domain or sending warnings to registrars about open invoices. As a big third party back-end provider (for .at, .berlin, .brussels, .hamburg, .reise, .tirol, .versicherung, .vlaanderen, .voting, .wien)
said he had considered an individual submission or informational document only.

The WG anyway would have to re-charter to take on new work, Galvin said. It will be interesting to watch how standardization@IETF and implementation/mandating by contract@ICANN will be connected – or used by parties.

DNSOP - DNS over TCP, Qname synchronization, Cookie

The biggest issue the DNSOP currently is working on is the upgrade of TCP as transport protocol for DNS. The main aim of the draft presented by John Dickinson (Sinodun Internet Technologies) was to „put TCP on the same footing as UDP“ and „make it an equal transport protocol“. The motives mentioned were privacy and DDoS prevention (see Dprive above). Dickinson listed what was necessary to make up for loss of speed when using TCP instead of UDP: connection reuse, pipelining and fast TCP. Pipelining allows to send queries without waiting for outstanding answers.

Fast TCP

Fast TCP which for now is only available on Linux allows to send data in the syn packet protected by a server cookie. The first response then can be sent back before the three-way handshake is fully completed. This allows to gain speed. Fast TCP makes necessary to change code and kernel. The document will go forward as a WG document, interest was considerable. Potential pitfalls mentioned were that servers might chose differently in what they will accept.

Lorenzo Colitti, Google, said, Google had to stop pipelining because of the high failure rate. Therefore language that servers should expect and accept pipelined queries was not strong enough, Colitti recommended a „must“ and was supported in the call for stronger language by Olafur Gudmundsson (Cloudflare). Without such mandatory language there was a need that servers signal, if they support the pipelining and „fast open“.

On discussion related to the TCP/UDP debate was about the potential need for a pre-negotiation over when a DNS server will close a TCP connection. The potential attack scenario was that a lot of TCP connections can be opened with the original querier going away. While there has been a patch for the problem originally addressed by Paul Wouters with the EDNSO option „edns-tcp-keepalive“), he said, there had been expressions to using the proposal for managing TCP connections. Ray Bellis (Nominet) with his new draft wanted to allow for a less heavy-weight solution by having servers send a single „connection-closed“.bit. Discussion is ongoing.

QNAME

On QNAME minimization (for QNAME see also DNS privacy above) on the way to last call there was still need for running code and testing, document author Stephane Bortzmeyer said. Bortzmeyer said he would be willing to test, but not to make the necessary changes in Bind. He was not convinced that more theoretical analysis was necessary. Review work by a long list of people was also still missing, according to the WG secretary, Paul Hoffmann.

DNS Cookies, DNSSEC negative trust anchors

Beside TCP and QNAME the WG worked through a considerable list of very old, old and newer issues. One of the recurring documents was the one on DNS Cookies (Donald Eastlake) intended to detect common denial-of-service and amplification, forgery or cache poisoning attacks by off-path attackers. Clients using the Cookie system would inlcude a DNS Cookie option in in every DNS request. Depending on whether a server cookie has been cached before, the clients have to use different variants of the Cookie OPT. There was no opposition to adopt the document which dates back to 2006 as a WG document.

WG participants were divided about the introduction of negative trust anchors. The idea presented by Jason Livingood (Comcast) is to prevent the failures of domains where the failure was the result of DNSSEC misconfiguration. The failure reason must be checked by DNSSEC experts of a provider/recursive resolver. Only when it is established that misconfiguration (and not malicious behaviour) is the reason, the negative trust anchor should be set instructing the recursive resolver to not perform DNSSEC validation for the respective domain. Users can access the respective side. Concerns mentioned were that this could result in keeping up misconfiguration permanently.

DBOUND BoF - New WG to be expected on ways to signal cuts between public and private domains or broader policy assertions

DBOUND, Domain Boundaries, met for a „Bar BoF“ with around 30 participants to discuss a potential working group to develop a potential alternative to the Public Suffix List (PSL). The PSL, currently operated by Mozilla, has known shortcomings (for example slow updates) and has been said to fall short as a tool with the rapidly growing new TLD space. After Andrew Sullivan (Dyn) developed a first memo in Asserting DNS Administrative Boundaries Within DNS Zone in March, Casey Deccio, VeriSign, and John Levine now pushed to develop the problem statement and charter for a new WG. Points to be addressed according to Deccio's BoF presentation are evaluation of PSL effectiveness, its maintenance and distribution, determine the „demand for domain relationship identification other than provided by PSL“ and, finally, find solutions to improve, but also extend the PSL. Use cases listed now in the updated Sullivan document include

- HTTP cookie state management (and other cookie applications),
- user interface indicators
- setting the document.domain property („the DOM same-origin policy
- might be helped by being able to identify a common policy realm“),
- email authentication mechanisms (including the DMARC, which is just under way at the IETF),
- SSL and TLS certificates,
- HSTS
- Public Key Pinning
- linking domains together for reporting purposes.

The list could grow and potentially make it difficult to find one solution.

There are, according to the ongoing mailing list discussion, existing alternatives to the PSL. Besides PSL-like lists at IANA and the registry operators, there was a pointer on the mailing list to the so called Bailiwick reconstruction that „approximates the location of an Rrset within the DNS hierarchy“, described in an ISC paper by Robert Edmonds on ISC Passive DNS.

DANE: Looking for implementation and IPR troubles

Implementation and a dispute about an IPR announcement by VeriSign were the two top issues in the Dane WG. The WG proceeds rather smoothly, it is expected to finalized all items on its Charter by November next year, how to promote implementation is now of much interest to the WG.

According to WG Co-Chair Warren Kumari:
SMTP and SRV are in WG last call or close to it

Raw keys are expected to come next
OpenPGP keys to be WGLC early next year

DANE Security model
DANE IPSEC (Paul Wouters said LibreSwan was working on this)
Reverse Binding server to client
Dane 6698 and SRV advanced to Internet Standard
(For the latter four Dane Chairs are looking for editors, all documents here.)

Despite DANE being quicker than many other WGs, there are some timing issues. The authors of the SMIME document would wait for the input following VeriSign's implementation, Jacob Schlyter, Kirei, said.
Possibly putting the Dane base spec and SRV on the standards track by the end of next year might be too early given that DNSSEC which was a condition for DANE still was waiting for more deployment.

VeriSign Labs has S/MIME using DANE for Thunderbird (more VeriSign Labs Dane here) and will open source the resulting code. Signing and encrypting are separated, users still have to invoke the actions Dane Sign, Dane encrypt (click!). One of the interesting features is that clear text will be stored nowhere (encryption only when displayed). One issue discussed was that Thunderbird's existing S/MIME function was not used. Eric Osterweil, VeriSign Labs, who presented the test run, said, it was just much easier and faster to put DANE beside instead of cracking open Thunderbirds SMIME. The test run will also inform the SMIME document, Schlyter said. More testing is going on here (Enigmail).

While VeriSign Labs received applause for the implementation work, the company was frowned upon for the IPR notice filed with the IETF. The claim is based on a November 2013 filed application with the US Patent Office containing a list of 20 claims, the basic two claims being:

> *What is claimed is:*
> *1. A method of conducting domain name system operations, comprising:*
> *accessing a set of policies for operation of a domain name system (DNS) using a domain name system with security extensions (DNSSEC);*
> *generating a set of answers to questions associated with a set of domain names of a zone, based on the set of policies;*
> *generating a set of signed answers from the set of answers and a set of key data;*
> *storing the set of signed answers in a zone file;*
> *receiving a question from a resolver; and*
> *retrieving a signed answer based on the question received from the resolver and the set of policies to transmit to the resolver.*
>
> (*….)*
> *11. A system, comprising:*
> *a network interface to a resolver transmitting a question associated with a set of domain names for a zone, the domain name operating under a domain system with security extensions (DNSSEC); and*
> *a processor, communicating with the resolver via the network interface, the processor being configured to access a set of policies for operation of the domain name system (DNS),*
> *generate a set of answers to questions associated with a set of zones of a domain name, based on the set of policies,*
> *generate a set of signed answers from the set of answers and a set of key data,*
> *store the set of signed answers in a zone file,*
> *receiving the question from the resolver; and*
> *retrieve a signed answer based on the question and the set of policies to transmit to the resolver.*

It is unclear what the potential patent would cover, once granted, especially as it seems to be related to use cases of the technology.

Given the claim Paul Wouters asked to dismiss VeriSign's proposed text on enterprise use cases. Instead, a new document should be developed as long as the IPR claim was open.
Reactions in the WG questioned the effect a patent claim could have for use cases. At the same time WG members

Enterprises wanted to be able to associate keys with functions (instead of personal names), communicate across domains and leverage existing identity management systems, according to VeriSign's use case presentation.

Implementation of DANE in general was discussed following a presentation of Dan York („Most people are unaware of Dane“) who listed questions to answer. While in some countries (Germany) there were positive steps towards deployment and one government - New Zealand - decided to go for DANE (according to Sebastian Castro, Internet.NZ), development in other countries was slow given the perceived potential. Proposals made during the session, more publicity for Dane, an additional easy to read cook book on top of the operational document (Livingood, Comcast) , talking to create a push for adoption (following the Internet.NZ). To check out DANE and DNSSEC support see here.

Security Area

Two presentations in the security area triggered a highly political discussions, which could have made an interesting plenary debate. Representatives of the article 19 NGO promoted to check IETF RFCs with regard to rights issues. While privacy meanwhile was an obvious and well covered aspect checked out when RFCs are passed human rights did not stop at privacy, argued Niels ten Oever, Head Digital at Article 19 and Joana Varon from FreePress. Based on a proposal to do research for the human rights impact of protocol designs by Internet Governance expert Avri Doria the group wants to check existing RFCs for what rights they implicitly embody. According to their understanding RFC 1958 was supporting right to access to information. RFC 2369 and 2919 on aspects of mailing lists equal a digital right to association. Non-english domain names, standardized in RFC 5890 to 5892 secure cultural diversity and access rights.

The idea of the researchers is to establish a research group in the IRTF to advance the work, they want to also interview authors to understand their motivations when including the rights implicitly in the protocols. While the Security Area meeting participants listened politely, there was some push-back against „reading things into the documents“ (Richard Barnes, Mozilla). John Hall, CTO of the Center for Democracy and Technology also received some critical comments for his draft on censorship technology. While Hall said he expected it to be a reference document for developers to allow them to make informed choices when deciding on their standards, IETF participants said, it well could serve as an instructive document for the censors themselves.

All in all it was a surprising amount of policy content in the Security Area meeting. While IETF participants pushed the „policy guys“ back (with an unprecedented level of politeness one should note), politics has very much arrived at the IETF. Not only has the HTTPBIS Chair Mark Nottingham published his own draft on how to deal with „stakeholders“ in the so far geek-reserved technical standardization. Nottingham (and other IETF participants) clearly acknowledged that technical standardization is political.

„Any time you make a decision you are making a political statement and we need to keep that in mind“, Nottingham said, adding that the idyllic sterilized computer environment that did not have to deal with the messiness of the real world was a fallacy. Asked by this author what the difference was in the approach of his „stakeholder draft“ and the NGOs bringing their requests to the IETF in person, Nottingham said, that his draft was arguing for consideration of stakeholder interests in standardization, as interpret by the engineers.

The next IETF meeting will take place in Dallas, March 22-27, 2015