Council of European National Top-Level Domain Registries

Report on RIPE75

Dubai
22-26 October 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

RIPE has its own IoT Working Group

RIPE now officially has an Internet of Things Working Group (IoT WG). Being allocated a full slot during the meeting in Dubai to discuss a charter in addition to several IoT issues, there was broad consensus for the existence of the new WG. During the closing plenary, RIPE Chair Hans Petter Holen initiated the necessary procedure for the official work to begin. As no objection was raised, the new WG was declared operative.

The draft mandate lists the following action items for the IoT WG:

- *To discuss challenges and opportunities of IoT for the RIPE community*
- *To serve as a focal point for the RIPE NCC regarding community input to their IoT activities, including liaisons with other organisations*
- *To invite IoT communities for dialogue on matters of operational relevance, including security, the numbering system, and applicability of standards*
- *To develop positions of the RIPE community on IoT matters*

Jim Reid, Chair of the DNS WG in RIPE for many years, stepped up as Interim Chair to prepare the next steps and WG co-chair election, as Marco Hogewoning, RIPE NCC, withdrew from its current role as lead organizer. WG (policy) work falls under the remit of the RIPE community, while RIPE NCC should stick to operative tasks.

Before the Dubai WG had started, RIPE NCC had organized a full-day meeting on IoT in Leeds. A group of 30 experts and members had discussed causes for the low level of security in IoT. They identified "a race to be the first to market coupled with low margins, as well as a lack of understanding or experience dealing with security and privacy considerations", plus "a tragedy of the commons issues". Negative externalities affected a wider group than simply the manufacturers or users of compromised IoT devices. Regulation, while a possible remedy, was seen as too slow and potentially harmful for smaller players due to the cost of compliance. "Trusted IoT label" regimes were seen as running into similar issues.

Just another IoT body?

During the Leeds meeting, the question of duplication of work was raised with regard to a new IoT body inside RIPE. After all, there is already a large number of IoT fora and working groups. Kevin Meynell, Internet Society, spoke of over 40 such bodies. Meynell presented one of these bodies, the Online Trust Alliance (OTA), an industry body founded in 2007 (DigiCert, Symantec, VeriSign, Microsoft, Twitter, plus 60 more members), which merged into ISOC in April 2017. Over the years, OTA has developed a framework of principles it hopes to get adopted by industry players. The framework has around 40 principles in the key areas of "security, user access and credentials, privacy, disclosure and transparency, and notifications".

Beside ISOC, the other I* organizations are also active on IoT: for example, the upcoming interim meeting of the IRTF things2things research group, which has been created several years ago, and the IETF IoT Directorate, even though the IRTF and IETF are more focused on standardization.

Other competitors in the field of standardization include ETSI and the International Telecommunication Union (ITU). At the same time, the ITU harbours a number of parallel initiatives. One of these is pushing for the Digital Object Architecture handle system, proposed by Bob Kahn, as a potential alternative naming system for the IoT. Interestingly the office of the ICANN CTO just took it upon itself to take a closer look into the handle system and found it risky for lack of transparency, lack of decentralization and closed-source nature, but in the context of a project with the University of La Plata, showed the potential of using the handles on top of the DNS.

The other ITU initiative followed more closely by the RIPE NCC is the use of IPv6 for IoT, and the Reference model of IPv6 subnet addressing plan for Internet of things deployment. ITU's work on IoT is part of the Study Group 20 work (Question 3/20 – Architectures, management, protocols and Quality of Service; Question 4/20 – IoT studies related to network infrastructural architecture; IoT signalling and protocols will be developed in collaboration with ITU-T SG11). A full list of Study Group 20 working groups (called "questions") is available here.

RIPE contribution – Operational expertise and the S in IoT

Positioning the RIPE in the field of IoT, for one part, is a reaction to the initiatives mentioned above. Point 3 and 4 of the charter speak to that. RIPE NCC also felt it had to create the IoT space for the community before going out to other organizations that call on RIPE for expertise, Hogewoning said. The ITU and ISOC have already requested liaisons.

The RIPE and the RIPE NCC see their core contribution to the cacophony of IoT talk through operational expertise and the fact that RIPE gathers the operators that will have to deal with the growing number of sensors and smart devices in their networks, especially with regard to security issues.

To underline this point, RIPE NCC points to the RIPE Atlas network, which was, as Robert Kisteleki described it in a blog post recently, an IoT network long before other such networks came online. Since 2010, the Atlas' tiny probes (version three is a TP Link Mini Router) have been distributed globally and helped members and external researchers alike to feel the temperature of the internet, or simply to check on local or global censors or attacks. Security measures include:

- *Trust anchors (i.e. the starting points of verification of components from the probes' perspective) are pre-installed on all probes before they are distributed.*
- *New firmwares are distributed to the probes inside the existing communication infrastructure.*
- *All probe firmware updates are signed and each probe has pre-installed public keys to independently verify the firmware signature before upgrading.*
- *We have mechanisms in place that try to detect unexpected behaviour such as outliers or violations of the internal protocols.*
- *The probes don't provide direct services to the host or to the world, reducing the network-based attack surface against them considerably.*

RIPE NCC also supports a "responsible disclosure" approach. More on the security concept will be shared soon, Kisteleki said to this reporter.

Death of transit, death of the public network

For many years, there was a fight between the big platforms and the eyeball/access networks over who should pay the other party: network providers for the content, content for the transport to the end user/customer. Geoff Huston in his plenary talk at RIPE75 said that the failure to reach an agreement had resulted in a wrong answer being given: content providers and content delivery networks (CDN) now brought content directly to users, sometimes as far as into the edge networks. The network providers' role – and business model – was eclipsed by the strategy, and the switch from a public to a private network was underway, albeit unnoticed and slowly.

With the big content providers and content delivery networks (Google, Facebook, Microsoft, Amazon, Akamai and Cloudflare) using their own private networks, sometimes even their own undersea cables, transit was no longer necessary and even global addressing, global naming or migration to IPv6 seemed to become obsolete, Huston warned. Open standards and common protocols are not used inside the CDNs: they are running proprietary protocols. And while outside of the "bunker" of the well-maintained network content delivery was protected, everybody outside of this "bunker" faced ever worse pollution with DDoS attacks.

"So here is the new architecture. And the new architecture is private CDN feeds and they have now CDN service cones. Clients don't talk to clients, clients don't even reach out from their service cone anymore. This is not a global network, because users don't send packets to users. Everything comes down with CDNs. It's a private network, so what is the universal service obligation?"

Consequences of network privatization is that these networks are not regulated. Universal service obligations, net neutrality, rights of access, none have meaning for the private networks.

With only a few large CDNs being made available to allow people to really reach worldwide audiences, there is no competition anymore. According to Huston, the new incumbents can decide "what can and cannot be put online". During the discussion, Lee

Howard, Comcast, said that (High-priced) offers for those who can afford to have their content put into an edge network equated to privileged transport and was contrary to network neutrality.

Another result of the privatization of what had been the internet was that only "economically viable" places would be served by bunkers; continents like Africa were not on the map of the new private net.

During the discussion, participants pointed to overcoming of the incumbent operators' by ISP by putting IP traffic on top of the networks, but Huston pointed out that it took around 80 years (1920-2000) for the telecom operators' model to finally break apart: "80 years of suppressive technology, 80 years of monopoly, 80 years of abuse before it finally got too much and no doubt we will all get sick of the incumbents, but not this year, not this generation, certainly not this decade: you will need to be patient."

Power shift

Interestingly, a presentation by Falk Bornstaedt, Deutsche Telekom AG (DTAG), the Internet Exchange Working Group, illustrated the shift of power to the ever-growing large private networks, the OTT and the CDN. Bornstaedt offered to share information from the DTAG's flow control analysis with the OTT and CDN providers. It would allow to route around over-streched links to the DTAG network and distribute traffic more evenly and effectively. While Bornstaedt announced the opening up would be DTAG's new strategy, at least one of the guests, Cloudflare, rebuffed the offer during the discussion.

Diversity, "Women in Tech" and IPv6 for toddlers

Along with the new WG on IoT came also the Diversity Task Force. Having held several BoFs at preceding meetings, the Task Force in now officially formed and according to its charter, will monitor and report about the RIPE community's diversity and try to push for more diverse participation. The TF will continuously update its work plan. One action item on the list for the first RIPE meeting in 2018 in Marseille is to reach out to Afnic and its members, to France-IX and local hacker spaces. These activities aim at bringing more new people to the RIPE meetings.

RIPE already has a number of activities (even pre-existing to the TF) that target new people from different backgrounds, for example the RIPE fellowship and the RIPE Academic Cooperation Initiative (RACI), which brings a number of scientists to each RIPE meeting.

Women attendees are one focus of the diversity issue, since the numbers are still small. A test measurement using an opt-in gender question for RIPE75 was answered by 54 participants (out of 460). Of the 54 respondents, 70 percent were male, 22 were female, 8 were binary and 2 preferred not to say. Given that women could be expected to be more open to answer the gender question, 22 percent looks like a rather optimistic figure.

RIPE NCC currently has a staff of 99 men and 67 women, with some of the technical departments being all-male (research) or nearly all male (information technology and security). The senior management is all male.

For the first time, the RIPE NCC/Diversity TF organized a "Women in Tech" lunch during the RIPE week. In Dubai, two software engineers from Lebanon described their experiences. Maya Kodeih, Operational Assistant Director for IT at Ogero Telecom looking back at her 20 years' experience in Telecom called IT a boys' club in need to become a boys's and girls' club. Only a big restructuring of Ogero resulted in her finally making a career step, she reported. At the same time, the number of women managers of top IT companies has grown, said Kodeih, listing Sheryl Sandberg (Facebook), Lucy Peng (Alibaba), and Lebanese game developer guru Reine Abbas.

Kodeih's younger colleague Zeina Daghlas, working at the largest Lebanese mobile operator Jawwal, had a more positive story to tell, as she has been benefiting from several of the nextgen programs of the I* organizations and attended IETF, RIPE and ICANN meetings. Interestingly, a recent Unesco study placed the percentage of women in IT in the Arab countries higher than in many western countries.

One of the more practical steps the RIPE Diversity Task Force is working on is child-care during the RIPE meetings. A test to have on-site child care shall be made for the first RIPE meeting in 2018 in Marseille. Issues to be dealt with were to find a place, an operator that offered the temporary child care, insurance and good planning, as babies would need different care than toddlers or older children. With RIPE NCC being in charge of preparing the program, a fun question is: will there be IPv6 courses for young beginners?

Europol comes forward with a policy proposal on abuse

With only a small policy proposal for IPv6 PI space on the table of the RIPE Address Policy WG, the most interesting current policy proposal is certainly the one presented by Gregory Mounier, Europol, with the support of soon-to-be (1 January 2018) representative to the NRO Number Council Hervé Clément, Orange. The proposal 2017-02 wants to tighten the screws on the anti-abuse activities of the RIPE NCC, by prompting them to regularly check for the validity and actual function of the abuse contact address of RIPE members.

While the Abuse WG introduced an obligation (ripe-563) for members to provide an abuse-c contact, the policy did not provide for the validation of "abuse-c", explained Mounier during the session in Dubai. Information could be inaccurate and the RIPE did receive hundreds of reports of invalid information per year. The policy proposal wants to mandate the RIPE NCC to check annually if members' abuse contacts can actually be reached. While there is a process in place to make checks of members' registry data with the "assisted registry checks (ARC)", at the moment these checks are not performed annually and they focus on Registry Consistency, Resource Consistency and Route/rDNS consistency. Basically, ARCs identify "inconsistencies between the Routing Registry and BGP announcements" and detect "lame reverse DNS delegations".

The abuse-c validity was not part of these checks. Therefore, a new policy to bolster ripe-563 was necessary, argues Mounier.

The Europol officer has become a regular participant in RIPE meetings (as well as other Internet Governance meetings). During recent RIPE meetings, Mounier has called for richer WHOIS information in the RIPE database, namely the information about the holders of sub-assignments of address blocks and potential country tags. He also mentioned that better enforcement with regard to inaccurate data in the database was on Europol's wish-list. Given the longer wish-list, 2017-02 looks like a rather cautious start to get Europol into the business of policy-making.

Nevertheless, reactions were mainly negative. Only one positive comment was made from Jordi Palet, consulintel.es, during the Anti-Abuse WG meeting: he welcomed the policy initiative, arguing it was beneficial to fight victim abuse. The overwhelming majority of comments was opposed to the idea, giving a number of reasons: ineffectiveness of the policy with regard to the declared aim of fighting abuse, burdening once more those that already fight abuse and follow the policies, while also being an additional burden (cost) for the RIPE NCC. The idea that was the most vehemently rejected was that RIPE NCC should also escalate the process up to closing the LIR and deregistering its resources, according to Mounier.

Despite the rather grim reception of the proposal in the community, it has been advanced from "discussion phase" to "review phase" by the Abuse Policy WG Chair, Brian Nisbet. As soon as a new version of the proposal has been sent to the mailing list, RIPE NCC will perform a so-called impact assessment and based on it make an estimate of potential costs for the organization.

It remains to be seen if the community can be convinced and if 2017-02 will only become the first of a number of policy proposals from Europol. The Europol-RIPE NCC MoU concluded earlier this year contains, in section 3.j, the respective joint plan to "work to enhance Europol's involvement in the RIPE community, particularly through participation in RIPE Working Groups and the RIPE policy development process".

Working Groups

DNS Working Group: A more secure, but also more complex DNS

DNS has been simple, and fairly easy to understand. Not anymore, it seems, thanks to the efforts to secure it ever more tightly. The need to reinforce security was illustrated very aptly during the RIPE meeting during a plenary presentation from Iranian engineer Babak Farrokhian, who through investigating subtle DNS attacks, developed a new set of nice tools for DNS attack detectives.

Securing the paths from the stub upwards

Presenters in the DNS WG touched upon the same general issue, DNS security, albeit approaching it from different angles: Benno Overeinder, Nlnet Labs, explained the various steps taken to secure the last (DNS) mile and making the case for a TLS extension for "transport of a DNS record set serialized with the DNSSEC signatures [RFC4034] needed to authenticate that record set." The pending IETF draft wants to allow TLS clients to perform DANE Authentication of a TLS server without performing additional DNS record lookups (see also an older draft by Adam Langley from 2011).

The idea is to establish a mechanism that will give proof to a verifier that the DNS record is authentic without performing DNS queries itself, which is good when there are outdated middleboxes, and also prevents added latency. Overeinder's presentation very much pushes the idea of empowering the edge – something currently on the agenda of the DNS Privacy project working on the DNS over TLS stub resolver software "Stubby". A Microsoft version for Stubby was just released, a MacOS version with a much more user-friendly graphical user interface (GUI) is expected around the upcoming IETF100.

Making DNS plug-and-play, despite DNSSEC

Whilst presenting new features of the Knot DNS Server and Fred Registry system, Jaromir Talir underlined the need to simplify third-party management for DNSSEC signed domains. Knot Server 2.6 and Fred Registry 2.32, according to Talir, already fully support the respective standards that shall allow for more automation and a plug-and-play DNSSEC signed domain for end-users. There are three related IETF documents:

- Automating DNSSEC Delegation Trust Maintenance (RFC 7344)
- Managing DS Records from the Parent via CDS/CDNSKEY (RFC 8078)
- Draft underway in the IETF RegEXT WG on "Third Party DNS operator to Registrars/Registries Protocol".

Some years ago, there was a push by Cloudflare and Google engineers to automate maintenance of delegation information for third-party operators in the triangle between Registrar – Registrant – DNS Operator, as changing a domain's DNSSEC status can be difficult because of the various parties involved.

RFC 7344 specifies how trust can be maintained between parent and child in a DNSSEC key roll-over. The parent also "periodically (or upon request) polls its signed children and automatically publishes new DS records", according to the RFC. RFC 8078 added "a method for initial trust setup and also for the removal of a secure entry point." These two features had been lacking from 7344.

CDS resource records are supposed to enable DNSSEC validation, i.e. place an initial DS Resource Record Set (RRset) in the parent, roll over the KSK (updating the DS records in the parent to reflect the new set of KSKs at the child) and finally turn off DNSSEC validation (delete all the DS records), if necessary.

The newest attempt to standardize and simplify signed domains' management by a third party "describes a simple protocol that allows a third-party DNS operator to: establish the initial chain of trust (bootstrap DNSSEC) for a delegation; update DS records for a delegation; and remove DS records from a secure delegation." All these operations may be performed by the DNS operator in a trusted manner and without involving the registrant. The draft is currently on the agenda of REGExt.

Talir explained the implementation of these drafts' features in Knot (automated KSK rollover) and that Fred was to allow further implementation of DNSSEC in .cz, where 51 percent of 1,3 million .cz-names were currently DNSSEC signed. With Fred's implementation of RFC7244 and RFC8078, the registry would be able to take responsibility for managing KeySet when domain publishes CDNSKEY. The resource records

would be collected by the registry and the key-set could be linked to a domain to generate the DS in the zone file. To maintain and update the keys, the registry performed daily scans of all domains in the zone file for CDNSKEY records. Currently, this takes three hours for .cz. “By implementing all these things, we could bring the DNS back to the good old days: set-up and forget”, he said. The management by the parent could also boost DNSSEC globally.

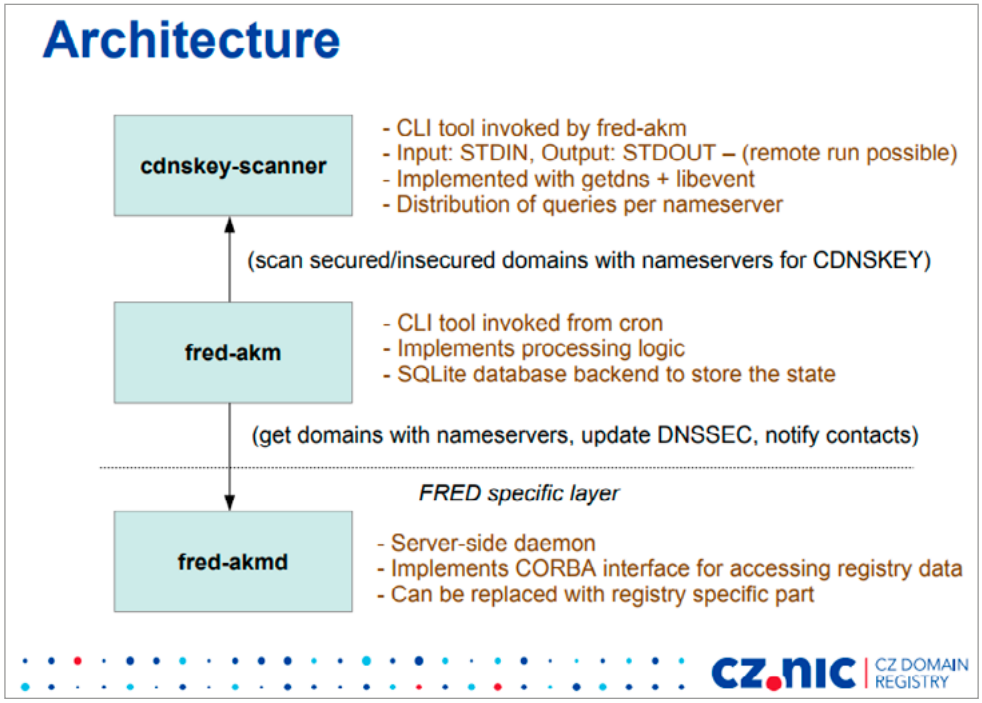

DNSPing and DNSTraceroute for DNS Detectives

A most intriguing presentation on DNS filtering was given by Babak Farrokhi, an independent Iranian software developer. When trying to identify the source of failed requests for mx records, Farrokhi started an in-depth investigation into the matter, developing two new tools for DNS troubleshooting and is now preparing an RFC for resolver transparency (meaning that everybody will be able to check which DNS resolver he/she is using).

Farrokhi’s investigation started with a failure to reach Twitter’s mx records, and what made the expert curious was the bad format of the answers he received from Google Public DNS: when digging for mx Twitter.com @8.8.8.8, the answer included “Got bad packet: bad label type”. This couldn’t be expected. A comparison of response times for Twitter and RIPE domains revealed a strange difference.

He used a ping adapted for DNS requests (DNSPing), programmed in Python and available via github. The DNSPing revealed the huge time difference between different domains queried via Google Public DNS. The DNSPing for Twitter showed an amazing/unusual/ suspiciously short delay:

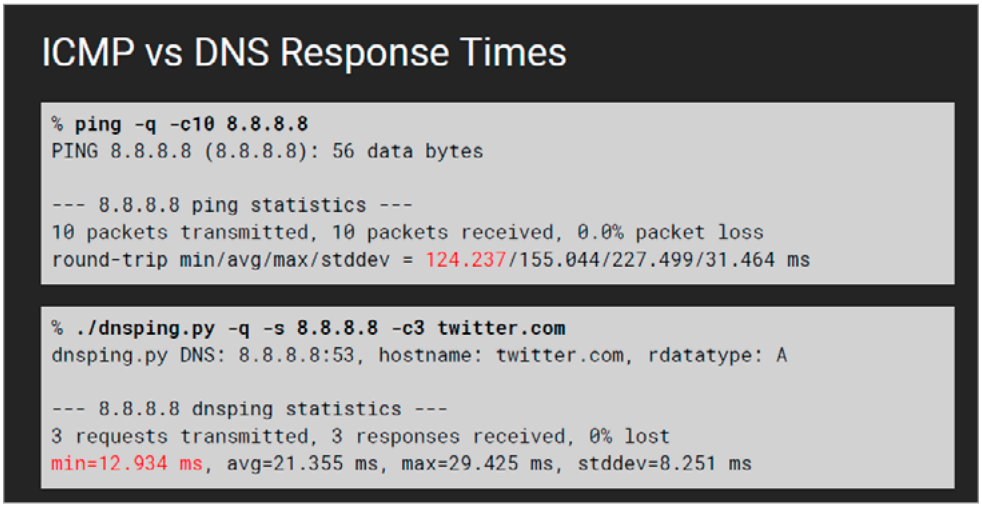

With these first results – different domains treated differently, rogue name server, close to the requester, impersonating as Google Resolver – Farrokhi went on in an effort to identify the server. Another new tool that he hoped would help getting closer to a solution was a special Traceroute for DNS queries. The results illustrated the considerable difference in delay for DNS traceroutes to Twitter or RIPE.net for pinging the Google server and checking for the DNS query for Twitter over Google:

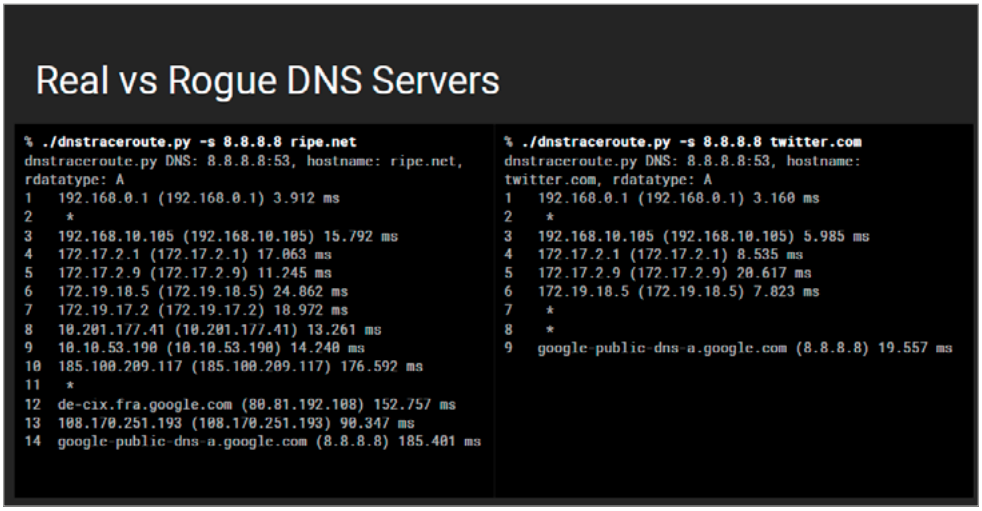

While the queries for ripe.net and other websites took the regular 120-130 ms (or a little more), when querying for Twitter, the answer was a lightning fast 15 ms. By using the so-called TTL trick – setting the TTL to one the rogue servers IP address could be revealed as using private IP addresses not even fit for public routing. Farrokhi checked around 10,000 domains on his system and received 139 broken responses.

Using RIPE Atlas to check the status globally, Farrokhi found that 2 percent of the 500 probes (DNS UDP IPv4) saw the same effect, and 1 percent showed the filtering under TCP traffic.

To uncover the IP address of the rogue server, Farrokhi relied on querying TXT record from maxmind.test-ipv.com. This reveals the public address of the resolver used, which brought Farrokhi to his next tool idea. He will propose a standard version for querying the resolvers that are used, based on the maxmind.test.IPv6.com. With such a tool, "resolver transparency" could be reached.

The resolver transparency is one of the practical results from the research. Farrokhi's other conclusions are the following:

- Don't trust a public DNS resolver, use your own

- There's no such thing as a free lunch (TANSTAAFL)
- Stub resolvers are easy to setup and use (e.g. Stubby)
- Don't trust your upstream, encrypt as much as possible

Rolling, rolling, rolling – or not?

During one of the plenary meetings, ICANN's Principal Research Scientist Roy Arends gave a brief update on the delay for the KSK roll. Immediately after the delay of the roll was announced, several members of the CTO David Conrad's tech team envisaged the publication of the list of affected servers to be published shortly. Around 4 percent of validating resolvers had shown to not taking up the new KSK, causing ICANN to step on the brakes and delay the long prepared KSK roll. The decision not to publish the IP addresses, according to Arends, was motivated by a reluctance to have this seen as a "name and shame" action. Instead, a consultant was hired to contact the resolver operators from the list one by one. This may take a month or two, according to the CTO office. For the moment, the technical department cautiously says it was still trying to "understand the signal" before considering setting a new date.

Reactions at the RIPE meeting were mostly positive. However, a DENIC representative asked if at one point, ICANN shouldn't risk crashing, as at least it would draw plenty of attention to the topic and ultimately support the adoption of DNSSEC and DNSSEC maintenance concepts. Arends called that a fair opinion, but there was a difference between people making conscious decisions to do manual key updates – and forgetting about it – and software bugs that operators and ICANN would not know about.

Cooperation WG

The Cooperation WG had a more fundamental discussion on its agenda regarding the future of the WG, but with one of the Co-Chairs not attending, the remaining Co-chair decided to delay this discussion. For several years, the Cooperation WG seems to have stumbled along, not really achieving its original mandate to improve community-government interaction during the plenary meetings.

One of the reasons certainly has been the competition the Cooperation WG faces from government roundtables organized by RIPE NCC, stemming from joint projects laid out in negotiated MoUs between RIPE NCC and governments (for example, a recent dedicated online tutorial for law enforcement officers, which are in high demand). Governments spending resources and time on these projects, and the much cosier roundtable meetings, might be hesitant to spend additional resources to attend RIPE meetings.

Another effect might be that a number of political topics have migrated to other places during the RIPE meeting, the IoT WG being a good example of this trend.

During the Cooperation WG meeting in Dubai, the WG heard Karen McCabe, Senior Director Technology Policy and International Affairs at the IEEE talk on what the large industry standardization body does with regard to growing ethical questions for technologists. The IEEE, which partnered with IETF and W3C for the OpenStands initiative, has in fact initiated a number of programs and standards efforts that cover the societal and political effects of new technologies. The 130-year-old organization views

itself as a forum to bring together technologists, politicians/regulators and NGOs. Initiatives working on ethical questions in the IEEE include the "IEEE Global Initiative for Ethical Considerations of AI/AS". The initiative has 13 WGs that have produced an impressive list of standard documents in the IEEE P7000™ Series Standards Projects:

§ IEEE P7000™: Model Process for Addressing Ethical Concerns During System Design

§ IEEE P7001™: Transparency of Autonomous Systems

§ IEEE P7002™: Data Privacy Process

§ IEEE P7003™: Algorithmic Bias Considerations

§ IEEE P7004™: Standard on Child and Student Data Governance

§ IEEE P7005™: Standard on Employer Data Governance

§ IEEE P7006™: Standard on Personal Data AI Agent Working Group

§ IEEE P7007™: Ontological Standard for Ethically driven Robotics and Automation Systems

§ IEEE P7008™: Standard for Ethically Driven Nudging for Robotic, Intelligent and Autonomous Systems

§ IEEE P7009™: Standard for Fail-Safe Design of Autonomous and Semi-Autonomous Systems

§ IEEE P7010™: Wellbeing Metrics Standard for Ethical Artificial Intelligence and Autonomous Systems

While these documents supposedly guide the work of engineers in the IEEE standards process, other initiatives lean more toward the convening, lobbying side, like the IEEE Internet Initiative.

Alain Durand, ICANN's CTO office, presented the ongoing work on the Identifier technologies health indicator (ITHI) project that was also promoted during two more sessions at the RIPE meeting and during the consecutive ICANN meeting. Durand listed DNS data accuracy, DNS abuse, overhead in DNS Root traffic, DNS leakage and DNS resolver misbehaviour as main areas to be analysed, or the names community and the numbers community are expected to come up with their own benchmarks/measurement candidates.

Two additional talks geared towards the policy questions were given during the RIPE Dubai plenary sessions.

Uta Maier-Hahn, PhD candidate at the Alexander von Humboldt-Institute for Internet and Society, presented conclusions of several years of research into the peering culture. In her research, she tried to find structures in the rather informal culture of peering and transit. Her main results include that it is uncertainties that make the coordination between the network operators, who are also competitors, necessary in the first place. Both trust and distrust were working in tandem, she said. Transit followed market order, but peering was a new form of barter.

The controversial "jurisdiction in the internet" issue was addressed by British researcher Sara Solmone. Solmone explained the concepts and issues of a) access-based jurisdiction and b) data location as criteria for jurisdiction. The first one allows countries to declare themselves competent through the mere fact that content, put online somewhere in the world, is consumed by users inside the jurisdiction. The result of a strict access concept would make every prosecutor in the world competent to prosecute those who publish the content. The dangerous side effect would clearly violate freedom of expression rights and give extraterritorial jurisdiction over the publishers. A compromise concept is the test of how much of the published content is targeting users in the jurisdiction that seeks to prosecute. Another concept is jurisdiction based on location of data, as exemplified in the Microsoft vs. US case. Mobility of data and the user's lack of knowledge (and lack of influence on) where the data is stored, also makes this controversial, Solmone said.

Address Policy on Sub-allocations

Currently, the Address Policy WG has a thin agenda. With IPv4 practically done and the controversial proposals from both sides – those wanting to get more resources from the dwindling last pool and those wanting to stretch the rest as far as possible for the sake of newcomers – have both withdrawn their policy proposals. One proposal that's on the table is on IPv6 PI space for Wifis. Following in a strict way the current RIPE policy regarding eligibility for IPv6 PI space will not allow organisations to be provided with PI space when this is the case.

There are several consequences to all this: people may opt to postpone IPv6 migration; they might be less specific or less honest when requesting IPv6 PI and answering related questions from the RIPE NCC; some might ask to become an LIR, or ask an existing LIR for less flexible PA space. The policy proposal wants to do away with the restriction.

The next RIPE meeting will take place in Marseille on 14-18 May 2018

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