Council of European National Top-Level Domain Registries

Report on RIPE79

Rotterdam

14-18 October 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

RIPE NCC: At the Crossroads

Thirty years after the community first came together to establish the Réseaux IP Européens (RIPE NCC), the organisation has arrived at a juncture marked not only by the final running-out of IPv4 addresses, but also by the stepping down by one of its key figures, Axel Pawlik, who had served as Managing Director of the RIPE NCC for 20 years.

A sudden withdrawal

Pawlik's withdrawal, which had not even been communicated to the extended RIPE staff until just before the start of the RIPE 79 meeting in Amsterdam, was only announced during the regular RIPE NCC Service session. Members were largely taken by surprise. Pawlik stepped down with immediate effect, handing over to senior management to step in and take on his functions until a new CEO has been chosen by the Executive Board.

According to the official press release, "the Executive Board will quickly start the process of recruiting a new Managing Director. In the meantime, the Board has appointed a new Management Team of Gwen van Berne (CFO), Kaveh Ranjbar (CIO) and Felipe Victolla Silveira (COO), who will lead the RIPE NCC and ensure there is continuity of service and operations during the interim period".

It seems that no preparations were made for a coordinated handover which feeds into the general feeling during the meeting that Pawlik's withdrawal had not been well planned or unanimous. While not commenting on the reasons beyond his intention to make space for the next generation, Pawlik said to this reporter that he was still considering what to do next. As he said farewell to the community, he underlined that he would emulate IPv4, and run out without really leaving.

Pawlik's withdrawal comes 20 years to the day after he started his position. During his tenure he has seen the community grow from 1,600 to over 20,000, staff grow from 60 to 160 and the budget rise from 4 million € to 34 million €. He steered the organisation during controversies with the ITU and ICANN, founded the Numbers Resource Organisation (and served in various roles for it). He also oversaw the frenzied run of members, newcomers and speculative companies for the ever rarer IPv4 bits.

End of IPv4 (well, really)

RIPE 79 will also go down in history for being the last RIPE meeting before the final waiting list policy for IPv4 addresses kicks in. Contrary to forecasts earlier in the year that predicted the run-out for spring 2020, over the summer the run for allocations of /22 address blocks accelerated.

The policy to hand out /22 blocks to every member and every newcomer started once RIPE NCC was down to the last /8 received from IANA. This block (185/8) was already exhausted last year, but due to returned addresses RIPE NCC was able to stretch resources further. With 3 million IPv4 addresses still in the pool over summer RIPE NCC experienced record numbers of requests. In July alone the RIPE NCC registration desk received 788 requests for v4 space.

While currently there are still around one million IPv4 numbers in the pool, according to Silveira, there are no continuous /22 blocks. Members and newcomers who are now on the waiting list will, he assured the RIPE 79 participants, receive equivalents of /22. Due to the backlog there are currently so many applicants in the queue that the available space might not suffice to satisfy all requests.

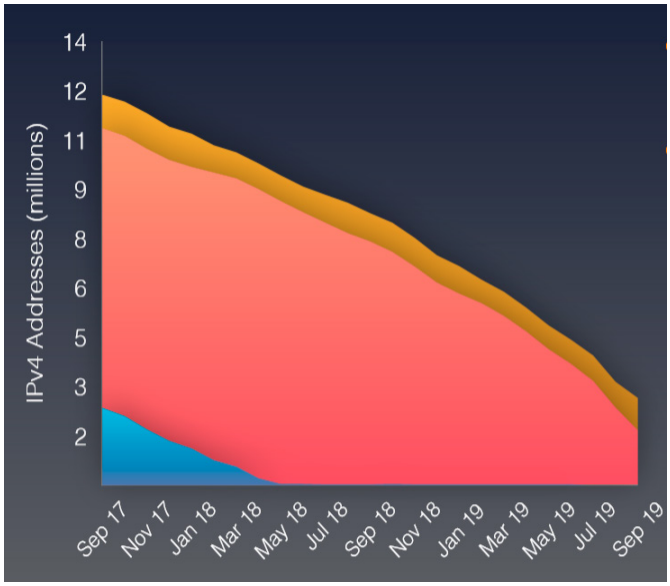

As soon as the numbers in the pool are below the equivalent of a /22 – (1024 single IPv4 addresses), the final IPv4 policy (2019/2) will kick in. It is a waiting list policy according to which only newcomers who have not been awarded any space are eligible to

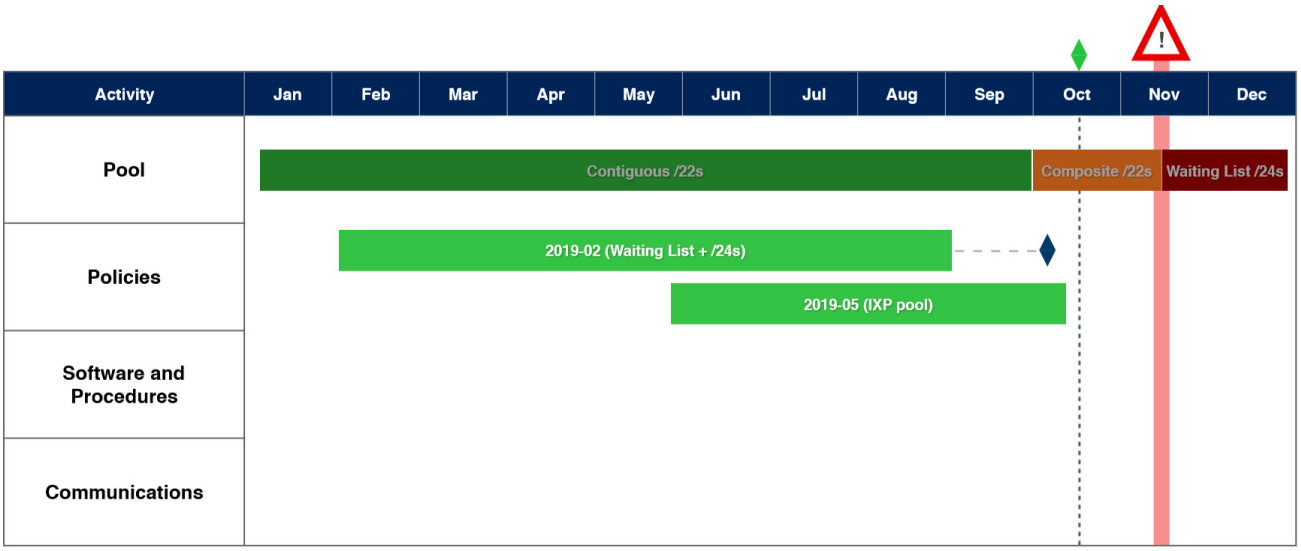

receive addresses, and it will only be a /24, 256 single addresses.

According to Silveira the final phase will be reached in November. After this there is one remaining source of addresses to be allocated. According to Silveira the recovery rate is around 1300/24 per year. This is space that RIPE NCC can allocate to newcomers, after it had the space cleared in a six-month quarantine.

Questions were once more posed on how the technical community can push IPv6 implementation. The proposal of Eric Bais, Member of the RIPE NCC Executive Board, to have the DNS Root servers stop serving IPv4 answers from 2026, was opposed to by the majority of participants in Rotterdam. “The stone age also did not end because it ran out of stones”, joked Jen Linkova, Co-Chair of the IPv6 Working Group (WG). People have to make conscious decisions.

The current rate of IPv6 announcements is around 25-30 percent.

Need to restructure RIPE NCC

The switch to the IPv4 waiting list policy marks a considerable change for the work of RIPE NCC. After years of being pressed for extensive due diligence and constant processing of small block assignment, now they have to hand out IPv6 blocks, that satisfy the needs of users for longer periods of time. For IPv4 allocation, there is still a waiting list with small numbers of space to be available. Other recipients for IPv4 numbers include future IXPs (according to 2019/4).

The second decisive change will be a consolidation and shrinking of the membership (and thereby membership fees). Since the start of the last mile policy in 2011 the number of members was blown up due to more and more companies securing their piece of the cake. Furthermore, some members started to open additional Local Internet Registries (LIR) to be able to not only receive one, but several /22 blocks. Briefly banned, RIPE NCC decided to allow it, as the respective companies just started to set up additional (fake) companies to acquire address space.

According to RIPE stats there are 25,000 LIRs, but only around 20,000 members, which is illustrative of the several-LIR policy. As IPv4 is essentially gone and only very small sets of addresses are available via the waiting list, RIPE NCC expects that by the end of 2020 the RIPE membership will decline by 1,500. More consolidation will follow.

While mentioned many times in Pawlik’s reports over recent years, following these developments, RIPE NCC has had to consider how to react with regards to financial and structural changes, possibly considering letting go of some of its staff and/or looking into other activities.

Challenges II (Governments and the re-purposing of the RIPE database)

A development which is accelerating according to speakers in various sessions at the RIPE NCC meeting is the growing attention that governments are giving to the internet’s self-regulatory bodies. During the Cooperation WG, RIPE NCC’s Head of External Relations said that national governments and supra-national legislators like the EU, the G7 or G20 “have really been starting to flex their muscles in the last few years”. Outgoing Managing Director, Axel Pawlik, underlined the hugely increased focus on what address registries

are doing, adding that one of the challenges the community has to take on is how to address this.

How this muscle-stretching takes place was on show in a panel debate in the Cooperation WG about the evolving controversy around the "purpose" or "re-purposing" of the RIPE database. The discussion was co-hosted by Europol and had two speakers presenting the requests of law enforcement agencies.

Please Re-purpose!

Cathrin Bauer-Bulst, Deputy Head of Unit for the fight against cybercrime in the European Commission's Directorate-General for Migration and Home Affairs (DG HOME), listed five points, with what she called the "lack of accuracy" of the RIPE database on top. She demanded that RIPE members should cascade policies down to those using the resources. RIPE also has to acknowledge that some resource users do not respect the rules and address the lack of enforcement tools against them.

Underlining that law enforcement and regulators were committed to cooperating with the RIPE membership, both Bauer-Bulst and Chris Lewis-Evans, Manager of Internet and Infrastructure Investigations, National Crime Agency, demanded that RIPE update the purpose of the database to include law enforcement investigatory interests.

Bauer-Bulst said: "(…) the RIPE database is just an essential piece in the very first step which is getting one step closer to the actual user of the resource. And if it does not do that properly, then there is an issue". According to Bauer-Bulst "that is the principal aim that we are pursuing when asking for accuracy, and we are open to finding a better word for that". Lewis-Evans added: "I think the purpose really needs to be updated on what we use the database for because there is so much more public interest around what is going on, so that has changed considerably to when the database was first here, that better informs what we mean by access. I think once we have the purpose we can talk about accuracy".

Tatiana Tropina, Assistant Professor in Cybersecurity governance, ISGA, Leiden University, recommended that RIPE members and the law enforcement agency (LEA) representatives had to compromise, as otherwise regulation would kick in and potentially in a messy way.

Peter Koch, Senior Policy Advisor, DENIC eG, warned against misunderstandings between the different communities. He underlined the fact that LEA representatives and the RIPE community had different concepts of accuracy. From the point of view of the registry – similar to other identifier registries in Internet Governance - accuracy was needed only to make sure that in the case of contested resources, the registry could decide who was the legitimate owner/holder of the resources.

The idea that the database necessarily allows for the identification of those using resources is a fallacy, this was much more a task for those moving the packets (instead of those registering the IP addresses). Koch also compared the problems of using registry data for the identification of a user to other classes of identifiers, such as domains in emails – with spoofing for example being widespread.

He also reminded participants that neither the original design nor the purposes intended by the design necessarily matched the use cases law enforcement had in mind. That law enforcement had used the database for a number of years - and the results delivered had not always been what law enforcement expected – in the end resulted from the fact that this special use was not reflected in the initial purpose. Koch drew a parallel to the Whois debate at ICANN where similar arguments had been exchanged.

The fact that re-purposing the database to address new use cases was not in the original design was questioned by several participants, including the RIPE Chair, Hans Petter Holen. Holen said during the debate that the time for the RIPE database might be over. Instead the RIPE membership could consider solely running the registry, which is essentially a list of "all the phone companies" and putting them in a public database.

Spencer Payton, Senior Internet Resource Analyst, RIPE NCC, and Daniel Karrenberg, one of the authors of the first version of the RIPE database, pointed to ongoing work on accuracy and due diligence. Karrenberg was also one of the proponents of the recently-established taskforce on database requirements. The taskforce is clearly a proactive move by the community to attack the questions around database accuracy, purpose and use cases.

The taskforce intends to produce a first draft by December 2019 and deliver a document for last-call at RIPE 81 in October 2020. The members of the taskforce are Peter Koch, Shane Kerr, Nick Hilliard, Bijal Sanghani, Sara Marcolla and James Kennedy. Since Europol is represented through Sara Marcolla from the E3C, the taskforce will have to take a position with regards to the re-purposing requests.

Challenges III (Re-Structuring RIPE's policy and organisational processes)

The database requirements taskforce is only one of several taskforces to ponder over RIPE's traditional processes and mechanisms. A dedicated community plenary hosted by the RIPE Chair, Hans Petter Holen, discussed several other taskforces.

RIPE accountability fallout/RIPE Chair selection

After the IANA reform was completed, an Accountability Taskforce set out to check the accountability of RIPE NCC and its various bodies. Their task was to "review existing RIPE community structures, documentation and processes to ensure they were accountable and in alignment with RIPE NCC values". Since RIPE78 the final document, RIPE 723, has been published. Besides reasserting how RIPE NCC understands its own organisation and function, the document makes 15 recommendations (see full list below) which RIPE Chair Hans Petter Holen addressed in Rotterdam in the "Community Plenary".

1. *Consider whether any formalised commitments are needed from the RIPE NCC (that it will implement policy, follow relevant community directions, etc).*
2. *Consider whether the RIPE Chair should be asked to disclose financial details associated with performing RIPE Chair duties and who covers these.*
3. *Consider reviewing whether current informal safeguards are enough to prevent bad actors from passing a policy proposal without the wider community having an opportunity to comment (not a great risk in the taskforce's view).*
4. *Consider including an explanation at the top of obsoleted RIPE Documents when there is no replacement document that it refers to. Possibly create a new "Archived" status for documents that are no longer current, but not exactly obsolete.*
5. *The community should consider whether more can be done to distinguish between the different types of RIPE Documents and whether consistency can be applied to the metadata for these documents moving forward.*
6. *The taskforce believes that the community needs to make progress on finalising the RIPE Chair replacement procedure.*
7. *Consider whether the RIPE Chair should report back to the community after representing RIPE in other forums.*
8. *Consider aligning the process for selecting working group chairs across the community.*
9. *Consider having more of a standardised process for informing new WG Chairs about relevant RIPE Documents and their responsibilities.*
10. *Consider developing a "crash course" for new chairs that covers things like how to effectively chair a session or determine consensus.*
11. *Consider developing general information for newcomers to explain how to participate in working groups, taskforces and BoFs and how the community functions more generally (the RIPE NCC could be tasked to produce this content)*
12. *Consider providing an overview of what the Working Group Chair Collective does and what it is responsible for.*
13. *The RIPE Document that defines taskforces is obsolete, and the working description on ripe.net no longer seems fit for purpose. Consider updating this with the description provided in ripe-464, which has been accepted by the RIPE community.*
14. *Develop documentation around the plenary and what its powers are. Also consider doing more to record closing plenary decisions which are not minuted currently.*
15. *Consider putting in place some kind of semi-regular review of the RIPE community's accountability.*

One issue which has just been finalized is a formal procedure to select a RIPE Chair, which has been discussed over recent meetings. RIPE documents 727 and 728 passed in August 2019 now determine

the RIPE Chair selection process and the Nominating Committee Details. After 30 years of having a very informal process – with only the founding Chair, the late Rob Blokzijl, and the successor he chose, Hans Petter Holen, serving in this role – now there is a formal process, with clear terms and timelines for the RIPE Chair.

In essence, a Chair can only serve two five-year terms, he will be assisted by a Vice Chair and selected by a randomly-designated nominating committee. The selection of a new chair will be discussed over three meetings; the nominations meeting, the consultations meeting and the concluding transitions meeting. A clarification of the role has been described in RIPE 714.

RIPE started the Chair selection process with RIPE 79 as the nominations meeting. The Chair of the NomCom 2019/2020, Karrenberg, who was selected by the RIPE Chair according to the new procedure, published both the call for the RIPE Chair and Vice-Chair, as well as the call for the 10 voting members of the NomCom. (Self-) Nominations for the NomCom members were due by 10 November. Nominations for the RIPE Chair and Vice-Chair are due on 15 December. RIPE 80 will serve as the consultation meeting during which the different candidates will be presented to the community. RIPE 81 will presumably be the so-called transition meeting, when the new Chair and Vice-Chair take their seats.

According to several sources, Holen intends to run the new procedure himself and serve for one 5-year term, together with a Vice-Chair. While traditionally the community is always eager to keep its leadership as long as no change is warranted, there might be more candidates stepping up, not least for the Vice-Chair mandate.

Conditions for candidates (both Chairpersons and NomCom members) include physical presence at several RIPE meetings (remote attendance does not count). To be eligible for the NomCom, members must have attended at least three out the five most recent meetings. According to the document: "Volunteers must provide their full name, email address, and primary company or organization affiliation (if any) when volunteering. Volunteers are expected to be familiar with the RIPE processes and procedures, which are readily learned by active participation in a working group and especially by serving as a document editor or working group chair."

One issue the accountability also laid their finger on is that a unified procedure to select WG Chairs is still lacking. For several years, Holen has tried to push WG Chairs to come up with a unified procedure, but term limits and selection processes are still not in place.

Checks and balances for a community-driven organisation

Touching on recommendation 7, Holen explained how he had represented the membership in other fora. He also briefly addressed the question of financial details associated with performing the RIPE Chair duties (Recommendation 2), mainly noticing that so far, the RIPE Chair is not reimbursed (apart from his travel costs being covered by RIPE). It might be worth considering changing that, Holen said.

Another recommendation includes a potential formalisation of the relation between RIPE NCC and the RIPE community (recommendation 1), especially when it comes to the implementation of policies passed through RIPE's Policy Development Process (PDP).

Potential PDP Taskforce

What could turn into a complicated and big issue for RIPE is the re-consideration of its PDP as such. One core question addressed during the Community Plenary was how numbers ("counting heads") should be weighed when deciding about consensus in PDPs. A comparison presented by Petrit Hasani (RIPE NCC) illustrated that the number of participants at RIPE meetings was higher than the number of participants who discuss policy proposals on the WG mailing lists. The practice of taking mailing list consensus as decisive over the discussions of meeting attendants could be reconsidered, Holen concluded.

A more qualitative approach could also be recommended, following lots of "+1"-style support for policy proposals on the list. Anonymity in mailing list discussions in the worst case could open the door to targeted campaigns and paid-for trolling. What constitutes consensus for the RIPE community was defined by the Accountability Taskforce.

Proposals made during the session by participants were to concentrate policy discussions in one mailing list to make it easier to follow these, instead of having PDP discussions on the various WG mailing lists (Sascha Luck, remotely); and monthly webinars in which policy proponents explained their proposals (Abdukarim Oloyede, AFRINIC).

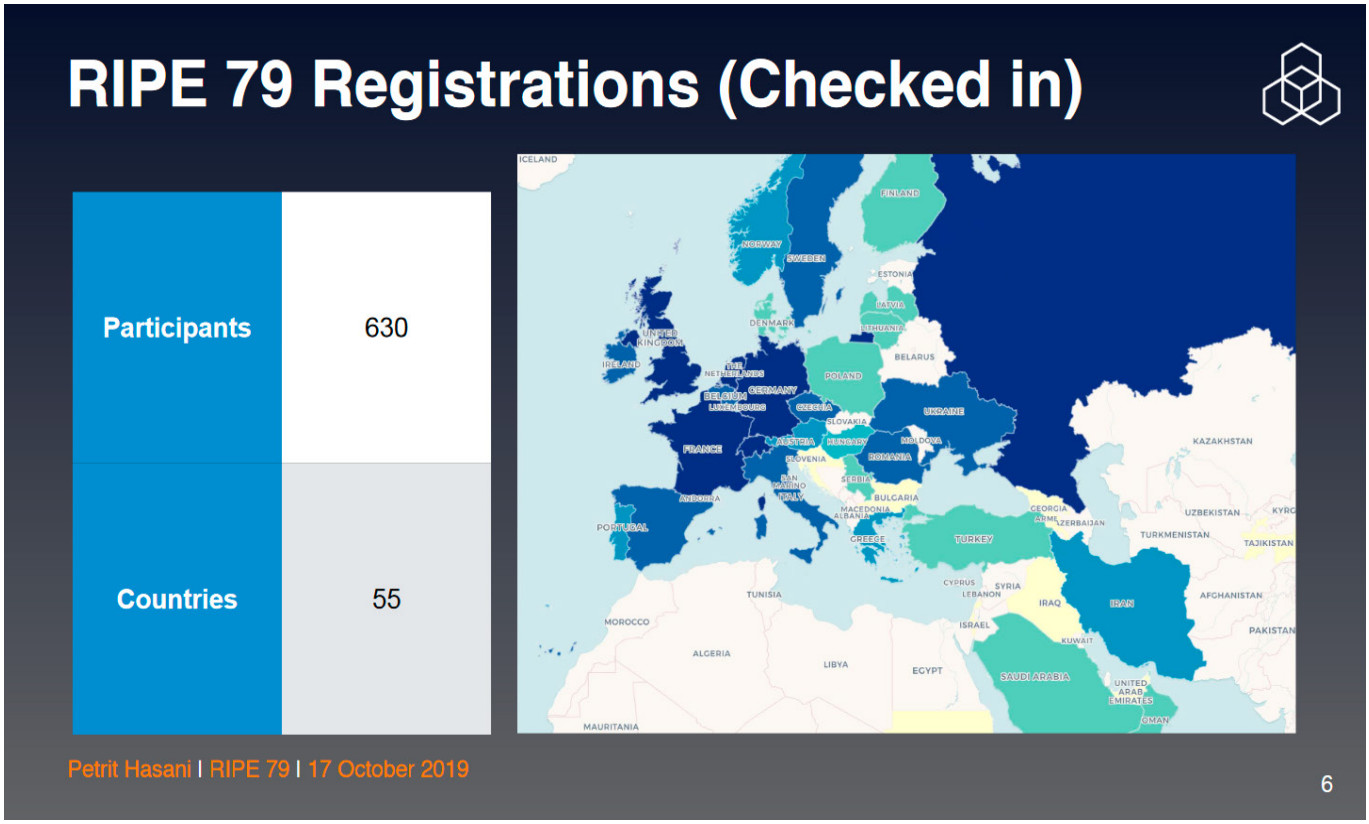

The RIPE Chair concluded that he might set up yet another Taskforce to focus on the problems of the PDP.

Pushing Diversity?

Diversity has been one challenge the RIPE community, assisted by staff, has tried to address in various ways. Care was organised for parents with young children (which was fully booked at Rotterdam for example); women attending tech-lunch meetings and panels have tried to focus on the problem of female under-representation in the industry in general and at RIPE meetings. The organisation gave itself a first Code of Conduct, and WG Chairs have received training to improve how they moderate discussions, according to RIPE staff.

Nevertheless, according to the RIPE Diversity Taskforce, RIPE NCC has to step up its anti-harassment policy. During the community plenary Brian Nisbet (original member of the TF) and Sacha Romijn (first-timer at the RIPE meeting) presented an updated and much "teethier" Code of Conduct. The original version, according to the initiators, did not spell out procedures and sanctions for violations well enough.

The draft includes an additional structure like a 4-6 member CoC Team (nominated by the RIPE Chair in consultation with the Diversity Taskforce) that will receive reports for victims or witnesses of inappropriate behaviour, as well as including a list of sanctions an offender might have to face. The list starts from private or public warnings or reprimands, to obligations for public apologies, to a ban from approaching the victim or even a ban from attending (any) future RIPE meeting.

Initiators rejected comments against heavy sanctions and called for a slower pace, with Nisbet underlining that the first CoC had made no change and this was already slow motion. Pointers were also made to other organizations who have such CoCs and the risk that people would stop coming to RIPE meetings. To support the argument Nisbet and Romijn pointed to the results of a quickly done survey in which 38 out of the 68 respondents said they had felt harrassed in some form during RIPE meetings.

Of those objecting, Malcolm Hutty's written comments seem to describe the counter arguments most clearly. Hutty warned that given the gravity of the sanctions foreseen, the procedure lacked fairness and due process. For example, the Code neither required nor suggested that the CoC Team should attempt to speak to the accused party, and "indeed, there is no requirement that the accused person is even informed of the details of the allegation against them". Instead a public reprimand could be the first the subject hears of the matter. The accuser (victim or witness) on the other

hand might stay anonymous, plus CoC Team members could act as a "Code of Conduct patrol, making reports of violations and then sitting in judgement on their own allegations". Hutty called the proposed Code "as a whole is riddled with bias" and recommended that the document should be abandoned.

While the document seems not to be able to win a majority of members, some questioned whether there was even a need to have majority support.

The RIPE Chair will now have to consider how to proceed with this rather sensible issue.

Where have all the networks gone? The robustness of BGP and the power of concentration

The concentration and consolidation of network services and network traffic have caught the eye of the technical community over recent years. In a BoF organized by Hisham Ibrahim, External Relations Officer and Technical Advisor (Middle East Regional Program Manager) at the RIPE NCC, Ibrahim elaborated on what he said the internet had evolved to, that is, interoperable vertical silos. Large Content Delivery Networks (CND) formed their silos as did large companies like big platform operators or sovereign countries with the latter starting to do their own scrubbing and filtering of traffic in and out of their silo.

On the downside this meant that the internet as was, with low barriers of entry for everybody interested in reaching and serving everybody on the network, is shrinking. Also the silos in essence were able to create their own network policies, technologies, protocols, independent from the need to interoperate with other silos or the old-style internet. Traffic stats and observations in a number of presentations and reports illustrate the development.

Nokia CTO and scholar Craig Labovitz has concluded in his research on traffic patterns over the last decade that while the internet is getting bigger in terms of traffic volume, at the same time it is getting smaller by the concentration of content sources. 90 percent of consumer traffic today is processed by CDNs, and IPv6 represents 20 percent of traffic but is stagnating, an effect by concentration and v4 content. Furthermore, while the number of routes was over 800000, fewer than 500 routes were used for 90 percent of the traffic.

Additional evidence can be found in earlier presentations by Geoff Huston, Chief Scientist at APNIC (see Death of Transit) and a recent longer report by ISOC. The most troubling observations in the ISOC report are certainly the deep dependencies, which means that applications and services become dependant on a small number of platform providers, and the domino effect this could have with regards to other parts of the global economy. One example given during the BoF session was ID management. When trying to establish their ID management system, CZ.NIC experienced that it was impossible to go against Google, according to Petr Špaček (CZ.NIC). The ID management put in place by Dutch banks (iDIN), which can now also be used for other applications, illustrated that it was still possible to do things in the open internet, argued Ilijtsch van Benim. Another ID management system that is trying to be federated is the one supported by a number of ccTLDs, ID4me. Brian Trammell, a member of the IAB (who has recently moved to Google, though he underlined he was not speaking for them), also said that the positive side was that the silos were all interoperating and interconnected.

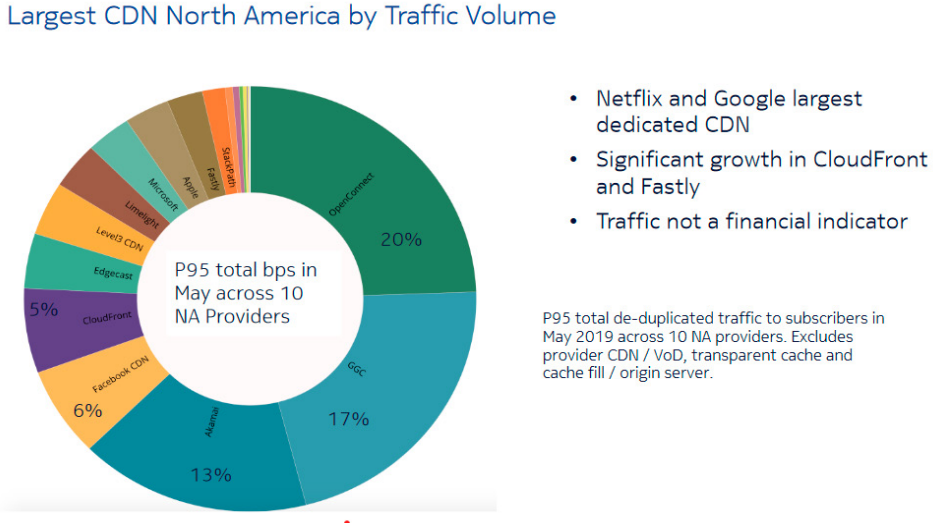

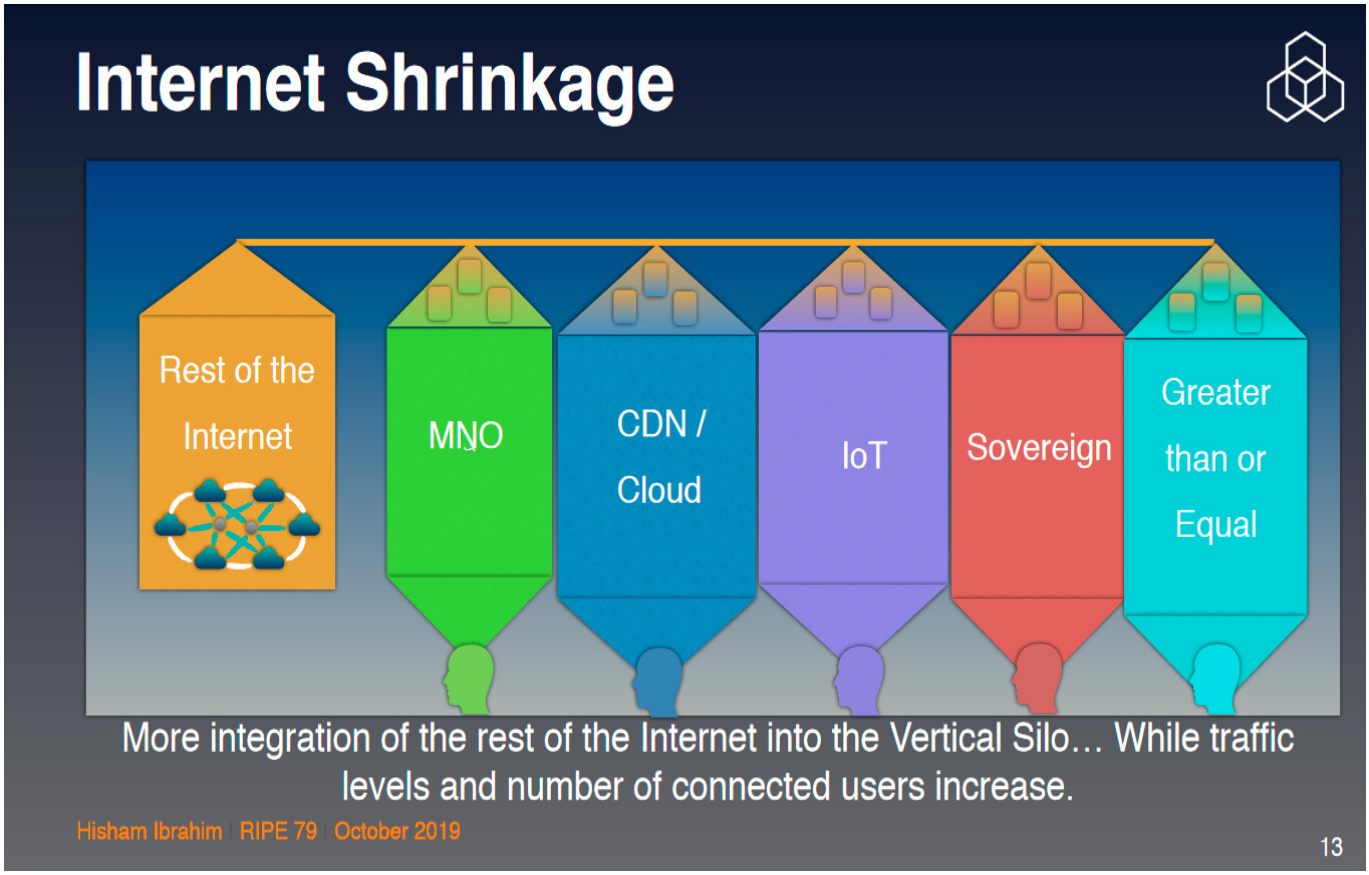

The most pessimistic view came from Geoff Huston who argued that with the death of competition (resulting from the network effect and the development of monopolies) on one hand, and the lack of regulatory mechanisms on the other (due to deregulation) it was game over for a hundred years.

However, in a plenary talk Huston called the Internet routing protocol BGP underestimated and still fit for purpose. According to Huston the BGP has lasted due to its simplicity, its preparedness to reuse functionality instead of duplication, a focus on what is necessary and its flexibility with regards to business models and policies. In his opinion one should not expect a change to this basic internet protocol anytime soon, because the community of operators has learned to use its strengths and tolerate its innate weaknesses (like security weaknesses, now on the agenda with RPKI and additional protocols). The levels of abuse, Huston said, were tolerable, and the protocol and business model have come to terms with each other. Other inter-domain routing protocols would only be a good option if there were a uni-provider internet, he said. Given the concentration debate, scenarios for this might nevertheless become possible.

Working Groups

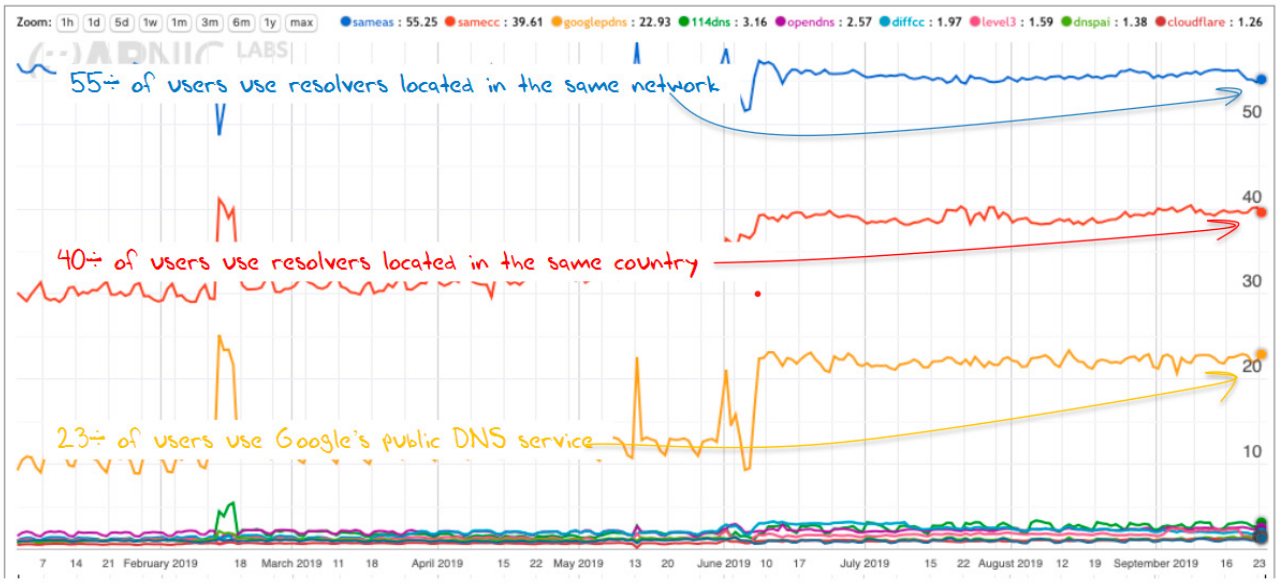

DNS WG

Geoff Huston (APNIC Chief Scientist) broke down his stats on how centralised the DNS is. According to his findings, around 23 percent of users have Google in their full resolver set (9 percent use Google as a first resolver). Google's public resolvers are used for routing around national/local filters or used by ISPs directly because of cost reasons. Other open resolvers have much smaller percentages of DNS traffic (all figures are available at https://stats.labs.apnic.net/rvrs).

In the overall picture, Huston noted, it was still mostly ISP settings that decided over the resolvers used. Users rarely change the default setting. Despite Huston declaring that concentration is in its early stages, he sees a trend towards the concentration of traffic in a few large resolver farms (three resolver farms are responsible for 30 percent of queries, 450 visible DNS resolver sets handle 90 percent). In his longer written piece, Huston stated "out of some 15 million experiments on unique end points, some 592 grouped resolvers out of a total pool of 23,092 such resolver sets completely serve 90% of these 15 million end points, and these users direct all their queries to resolvers in these 592 resolver sets".

Huston's question on whether there is pressure for the DNS to aggregate to ever larger resolver farms therefore seems to be answered to some extent.

There are still some open questions, he noted, such as what the economic model of name resolution in a highly aggregated environment will be and if data mining will be used to generate revenue streams. The most difficult questions Huston posed to RIPE participants were if it was possible to reduce information exposure while still using common resolver caches and what the nature of the trade-off between resolution performance and information leakage in DNS resolution was. To both, he said, he had no answers. With regards to the question of changing from the current model of the DNS as an infrastructure to DNS being decided anew every time a user runs an application, Huston predicts that this could become reality. Huston in his blog post announced further work on the level and layers of DNS centralisation.

Huston's study was one of the contributions related to the DoH discussion. The RIPE DNS WG also received a presentation of the study on how DoH, DoT and the classical Do53 compare speed-wise. A group of researchers from the University of Chicago and Princeton University continued measurement campaigns to check how encryption influences the speed of DNS answers and page loads. The work which was already presented at IETF 105 (Applied Network Research Prize) so far has some interesting results showing that DoH and DoT from different operators might be faster than normal DNS. DoT at the same time beats DoH when it comes to page load times,

from which the researchers concluded that TCP was preferable. Work is ongoing, and the influence of where the queries are performed from still has to be better understood despite the fact that researchers this time used Google, Quad9 and Cloudflare servers in Frankfurt for their tests.

In other work, Petr Špaček (CZ.NIC) presented a new tool to benchmark DNS resolver software, avoiding shortcomings of benchmarks focussing on resolution performance. The latter over-focussed on the measure of queries per second, instead of looking at the number of parallel clients a resolver can handle. The DNS Shotgun, an open source tool based on software developed by DNSOARC (dnsjit), according to Špaček helps to simulate real clients from captured traffic.

In a first phase traffic is captured from a particular deployment which, in phase two is then replayed to the setup chosen by the researchers. By compressing the used PCAP data the researchers can simulate more traffic over the same time. Results showed differences of the various DNS server software abilities to process parallel queries and, in one case (Bind) helped to find a DNSSEC related bug. For the Shotgun software, go to https://gitlab.labs.nic.cz/knot/shotgun

For the results of the benchmarking study (Power DNS, BIND, KnotDNS and Unbound) see here.

For a nice overview over the software tools see an OARC report.

More DNS work was presented on the DNS measurement project, the search for the perfect TTL for caching and using DANE in HTTPS validating on Linux.

Willem Toorop (NLnet Labs) asked RIPE members and interested parties to give feedback on a project to make the Atlas-based DNS measurement initiative publicly available. Measurements undertaken since the start include what DNSSEC algorithms are in use and how resolvers performed during the KSK roll-over. Now the group, which includes RIPE Atlas experts, is considering inviting others to use the DNS Daemon for their own measurements, but also allowing interested parties to rate the resolvers they use with regards to security, privacy, performance, based on such measurements. “For example, for security, you would get 60% of the five stars already if you are doing DNSSEC validation, which is the main security feature of DNS, and then 3% for each additional algorithm that resolver supports and then 10% for trust anchors and 10% for not doing NX domain”, Toorop explained. The group hopes to receive feedback on the idea to have a statusofthedns.org platform.

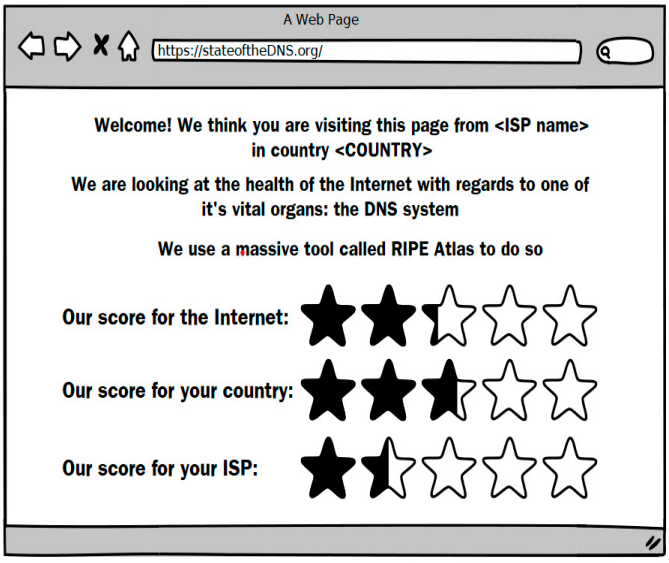

Danish, a Linux daemon for validating HTTPS DANE, is still rather experimental, Andrew McConachie reported, and for the time being it was very much an exercise in the generation of NXDOMAIN responses. Some TLSA records nevertheless could be found in the wild. At the moment the daemon can inspect TLS handshake traffic with lib-pcap and install ACLs to potentially deny traffic when the validation fails. Much more work is necessary for the daemon to be able to run on firewalls or end-hosts.

Giovanni Moura, SIDN Labs, once again presented his plea for longer TTLs, pointing to the advantages in performance. Caching near the client, he said, would beat even great infrastructures. According to measurements taken median response times from anycast without caching were up to 29.96 ms, while cached Unicast answers would only need 7.38 ms. Furthermore, query load would go down, so TTL would matter more than performance. More details and results from the authors that also show how they were able to reduce latency in one country-code TLD from 183 ms to 28.7 ms (.uy, with the TTL raised from 300 s to a day, at the same time making the difference between authoritative servers and root default smaller) can be found here.

Possible recommendations to have longer TTLs have been proposed to the IETF as an informational (instead of the earlier planned authoritative) document.

Cooperation WG

Michiel Steltman, Director of the Digital Infrastructure Foundation Netherlands (Stichting Digitale Infrastructuur Nederland) promoted the approach by the organisation for a responsible vulnerability disclosure process and the distribution of information to companies. Given that patching rates were low - most of the over 20,000 vulnerabilities found by various experts and initiatives in 2018 were not patched and also not exploited – the organisation started abuseplatform.nl. Abuseplatform.nl is part of the Abuse 2.0 project, an initiative of AbuseIO, ECP, DHPA, DINL, ISPConnect, NBIP, the Ministry of Economic Affairs and Climate Policy, the Ministry of Justice and Security, SIDN and many others.

The idea was originally developed for Dutch hosters to draw on many sources that crawl and find vulnerabilities (Spamhouse database for open relays to US, Stop Badware, additional, new sources) and on sources of those aggregating found bugs (another Dutch Foundation, Abuse.io, the Dutch National Cybersecurity Center, the Technical University of Delft and others). According to Steltman the TU Delft recently received government funding to expand their analysis of performance losses in networks that result from "infections".

With the RIPE talk the organisation wanted to expand the initiative further. Infrastructure providers like RIPE NCC and its members, the LIRs, as well as ISPs or hosters are called on by the Digital Infrastructure Foundation to monitor "badness" in their networks, to subscribe to the feed (and help to pep up the aggregated feed the foundation can provide) and to forward information to their customers. Infrastructure providers also, according to Steltman, should motivate users or customers to fix abuse issues in their services or act themselves. The Foundation wants those who cooperate to sign a Code of Conduct.

The effort is clearly trying to create intermediary responsibility up to the point where they not only assist customers, but monitor their networks and, in case of doubt, even act proactively. Researcher Tatjana Tropina warned against mixing different sorts of abuse in the initiative, as the mentioned child abuse material was much more a question of content crime, that needed different tools than decisions over DDoS, malware or botnet infections.

There were also questions regarding potential GDPR violations through proactive crawling and collecting of information. Here, Steltman argued that the organisation was talking to DPAs with regards to the exemptions available for cybersecurity/public interest data collecting. On the question of how LIRs who only act as registries for customers actually moving data could take a proactive role, Steltman acknowledged the diversity of operators. Hosting companies could be LIRs, but at the same time not all LIRs might have infrastructures or be "infrastructure providers". The initiative for example did not expect an internet exchange to act, "but we do expect from somebody who runs infrastructure to be more proactive".

Steltman said that legislators were certainly prepared to take legislative steps, if operators did not come up with their own solutions in time.

In his presentation on RIPE NCC's work on Internet Governance, Chris Buckridge also reminded participants that legislators were becoming more and more active. He pointed to RIPE NCC's answer to the UN High Level Group of Expert Panel's Recommendation on Digital Cooperation and noted that RIPE NCC's experts felt a certain "urgency and severity to the Internet governance issues". The feeling that "something must be done", but the lack of clarity on what this "something" means makes working in IG difficult, not least because new issues, fora and legislatorial efforts are popping up everywhere.

One ongoing effort the RIPE NCC is observing is the consultation by BEREC about a practical definition of the Network Termination Point (NTP). In the new European Electronic Communications Code (2018/72), article 2(9) of the NTP is defined as "the physical point at which an end-user is provided with access to a public communications network; in the case of networks involving switching or routing, the NTP is identified by means of a specific network address, which may be linked to an end-user's number or name".

During the Connect WG session Marco Hogewoning (RIPE NCC) explained that if the regulator were to decide that the modem was part of the network, the network operator "owned" the modem (and could use this endpoint for activities with regards to engagement for IPv6 or for IoT security). Otherwise each user could decide what modem he wanted (choosing your own modem is currently the prerogative of end users in several EU countries). Hogewoning encouraged

feedback from RIPE members to RIPE NCC to prepare a possible RIPE answer. The discussion is ongoing on the Connect WG mailing list (deadline of the consultation is 21 November 2019).

Abuse WG

Two weeks before RIPE 79 the most controversial policy proposal, "BGP Hijacking is an attack", was withdrawn, following concerns expressed by not only RIPE members via the mailing list (and during RIPE 78), but also by the Executive Board. During the regular impact analysis of the policy RIPE NCC has stated that implementing a process to judge route hijacking events (reported by victims), while in general covered by RIPE NCC's mandate would "expand the scope of the RIPE NCC service portfolio with the introduction of a reporting, evaluation and arbitration process for the purpose of validating claims concerning 'BGP Hijacks'". It would, the implementation reads, "add a routing regulator role to the RIPE NCC in addition to the established role as registry." While those in charge of making decisions would be a pool of experts, RIPE NCC still would be in charge of retributive actions and other parts of the policy.

The "routing regulator" role seems to be a euphemism for something that many members in the community are rejecting. The Executive Board seemed to be more concerned about the possible legal risk for the organisation, as sanctions have the potential to make RIPE NCC the target of legal claims. Another calculation made in the implementation report is that the measures would be costly, with one to two potential reports per day for the pool of external experts, and RIPE NCC would be on the receiving end of a rather large number of reports. Hijacks on the other hand might not be massively undercut, the calculation goes on.

The added cost of abuse policies were also documented by RIPE NCC in its report on implementing the regular validation of abuse-mailbox attributes. According to the report presented in the Abuse WG, between February 2019 and October 2019 RIPE NCC checked 77,200 abuse-mailbox attributes, including 18,200 LIR Org objects, and 45,500 LIR resources. 5,457 of 77,168 mailboxes failed automatic validation, according to Marco Schmidt (RIPE NCC). This is about 7 percent. Schmidt reported that altogether 8,000 abuse-mailboxes had been updated during the process, which meant that some members realised issues and proactively fixed things, despite having passed the test. The process was now complete and has been integrated in the regular workload of RIPE NCC. The costs RIPE members have to bear for this effort are considerable. As 20-25 percent of tickets needed manual follow-up work, three additional FTEs were hired temporarily, he said.

While this policy was a good first step, it was still unclear how many of the 93 percent that passed the automatic validation test were for real, Jordi Palet Martinez told the Abuse WG when presenting his policy on requiring validation checks to be answered by people instead than by automatic checks only. It was unclear, how many of the validated boxes were "fake", Palet noted. The proposal notes that emails sent to the abuse-mailboxes essentially have to "require intervention by the recipient, the abuse-mailbox host must not require from reporters to complete a form and must guarantee that abuse reports and related logs, examples, or email headers are received". After an initial validation of no longer than 15 days, the validation request will be escalated to other contacts of the LIR in question with a delay of no more than 15 days again. After this suspension, a follow-up procedure is required (see graph below).

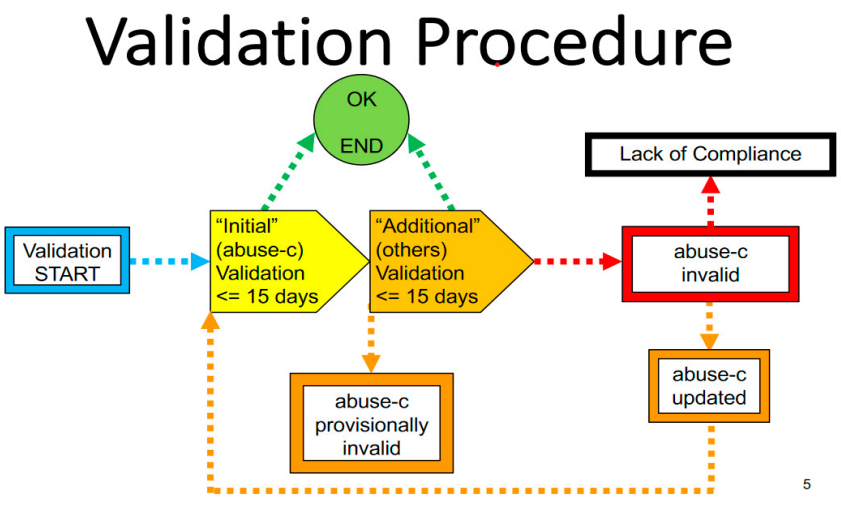

The deadline for further comments on the proposal was 30 October and so far, the numerous opponents remain unconvinced of the usefulness of the proposal. One major critique is that the policy would amount to prescribing businesses how to run their operations, something many declare not to be a task for RIPE NCC.

One of those who objected to the proposal is Peter Koch, DENIC eG, who criticized the authors for describing a test instead of a policy in the document, with the motivation being "weaponizing the registry by compliance cases". Without clear-cut real-world problem descriptions, these kinds of policy proposals had to be described as an "abuse of the policy process". Koch asked for a "moratorium" on such proposals.

The abuse WG heard two additional non-policy related presentations on potential tools to support anti-abuse work. In the first, Carlos Friacas from FCCN asked how RIPE members valued the possible drop of Autonomous Systems using the ASN Drop list curated by Spamhouse. The second presentation was on LACNIC's Warning Advice and Reporting Centre (WARP).

The next RIPE meeting will take place in Berlin, Germany, on 11-15 May 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 55 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.

Rate this CENTR Report on RIPE79

(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