Council of European National Top-Level Domain Registries

Report on RIPE76

Marseille
14-18 May 2018

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

Speeding up the net and moving control: Bandwidth, bottleneck and RTT (BBR) algorithm

It might be the way to go if the net was moving towards terabit/s connections, yet Google's new algorithm for congestion control – Bandwidth, Bottleneck and RTT (BBR) – pushes flows using competing algorithms aside. Geoff Huston, APNIC, reported that the clash of BBR with flows managed by Cubic was "brutal". BBR was highly aggressive and took up most of the bandwidth provided to Huston during his test with a stable 400 Mbit/s when connecting from his server to a server in Germany.

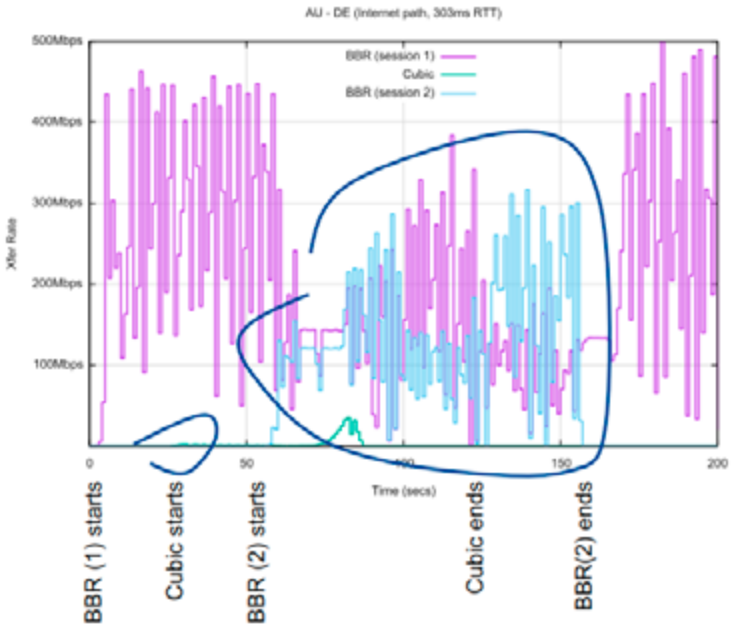

The trick used by BBR lies in the way congestion is asserted. While older congestion control algorithms like RENO or Cubic measure packet loss, BBR is neglecting package loss, bases its assumptions on round trip times and uses probing. By sending more packets for a short time and looking at how RTT changes, BBR flows are much more aggressive. In contrast, Cubic or Reno only send half the packets once they experience packet loss. The result is that the packet based flows loses out to BBR flows.

Earlier measurements done by academics already illustrate that BBR has problems as soon as multiple BBR streams compete with each other (see also Huston's measurement above with the two BBR streams intermittently getting a big chunk of capacity). But while it is still unclear how BBR will work once it is unleashed more broadly, Huston argued it was another shift of control from network operators to the app world. "What we are going to see is encrypted BBR", said Huston, pointing to BBR being implemented in Quic ("it is a package"). "It will run the way it wants to, not the way the network operator wants it to run. This could well be a real moment for the network operators, and they might well be saying while burying their head in their hands, oh, my God, you can't do this to us, it's so bad. I can't do selective packet drop or rated red, all of those instruments don't work, it ploughs on relentlessly."

Currently, BBR is used by Google for YouTube and Google.com. Google released it into Linux distribution. The question is if BBR2 (under discussion at the IETF as well) will be less aggressive. According to Huston's diagnosis, BBR "will never play fairly with everyone else: it's aiming at a different control point."

While network operators at the RIPE meeting said to this reporter they currently were not seeing effects from BBR, Roland Bless, one of the academics who measured the effects, warned that using BBR by Google was premature. "The effects have not been fully reviewed, especially not by independent researchers." While some users of certain Google services might benefit from it, users of other services would be discriminated.

ITU ideas for IPv6 addressing

Two Study Groups at the International Telecommunication Union are considering special addressing plans for IPv6 related to IoT. The RIPE IPv6 WG had invited the author of a possible standard for a subnet addressing schema for IoT to bring the proposal currently under discussion at the ITU Study Group 20 to the RIPE IPv6 WG.

Sebastian Ziegler, from Mandat International, a research consultant, presented the concept via Skype. When asked by Tahar Schaa (representing the German Minister of the Interior), Ziegler did not clearly disclose who was funding the IoT addressing proposal.

As Ziegler underlined, the draft standard text doesn't touch on the global prefix, nor on the local (or host) ID. Instead, it proposes a scheme to number the IPv6

subnet ID for IoT. Ziegler picked the prevention of a new digital divide – developing countries not using IPv6 – as the main motive for the proposal. According to the proposal, the subnet addressing scheme shall give guidance to interested ITU members on how to design subnet addressing. During the RIPE76 session, Ziegler underlined that it was optional.

The proposed subnet scheme includes:

- a hexadecimal digit (A), equivalent to 4 bits for buildings and locations
- a hexadecimal digit (B), equivalent to 4 bits, for different categories of subnets (demilitarized zone category for public servers, internal servers, regular local area network, Internet of things, an "other" category)
- hexadecimal digits for (C) and (D), equivalent to 4 bits each for specific subnets.

When IPv4 addresses require mapping with IPv6, the first and the fourth hexadecimal digits are set to 0 (see graph from the proposal below).

	Dual IPv6 - IPv4						Pure IPv6				
	IPv6				IPv4		IPv6				
Allocation	A	B	C	D	octet	Nb	A	B	C	D	Nb
DMZ	0	0	0-f	0	0 - 15	32	0-f	0	0-f	0-f	16 x 8'192
	0	1	0-f	0	16 - 31		0-f	1	0-f	0-f	
Internal Servers	0	2	0-f	0	32 - 47	32	0-f	2	0-f	0-f	16 x 8'192
	0	3	0-f	0	48 - 63		0-f	3	0-f	0-f	
Regular LAN	0	4	0-f	0	64 - 79	64	0-f	4	0-f	0-f	16 x 16'384
	0	5	0-f	0	80 - 95		0-f	5	0-f	0-f	
	0	6	0-f	0	96 - 111		0-f	6	0-f	0-f	
	0	7	0-f	0	112 - 127		0-f	7	0-f	0-f	
IoT & Building Automation	0	8	0-f	0	128 - 143	64	0-f	8	0-f	0-f	16 x 16'384
	0	9	0-f	0	144 - 159		0-f	9	0-f	0-f	
	0	a	0-f	0	160 - 175		0-f	a	0-f	0-f	
	0	b	0-f	0	176 - 191		0-f	b	0-f	0-f	
Others	0	c	0-f	0	192 - 207	64	0-f	c	0-f	0-f	16 x 16'384
	0	d	0-f	0	208 - 223		0-f	d	0-f	0-f	
	0	e	0-f	0	224 - 239		0-f	e	0-f	0-f	
	0	f	0-f	0	240 - 255		0-f	f	0-f	0-f	

Figure. 3 Subnet ID addressing plan

Participating RIPE members all pushed back against the proposal. Schaa said the German Government had made quite an effort to design its own subnet structure for the different states and federal bodies sharing the /23 address block requested from the RIPE NCC (instead of the /26 allocated earlier). Should the ITU proposal become a standard – and according to ITU procedures, would become mandatory in some way – the German addressing plan might become void.

Other RIPE members questioned assumptions made in the proposal. Jordi Palet, one of the IPv6 consultants for a number of governments (Spain, Latin American countries), said the proposal did not make sense. Jan Zorz, Slovenian IPv6 expert, warned that the proposal had too many mistakes and could not proceed in its current form.

Benedikt Stockebrand warned that the addressing schema might result in creating large patches of unused space and a run-out of IPv6 20 years earlier than necessary.

Global prefix based on telephone numbers

There is an additional proposal on the agenda of the ITU study group 20. This proposal, not yet discussed at the RIPE, wants to give the global prefix of IPv6 the shape of a telephone number, according to the e164 standard. The proposal will be presented at the ITU Study Group 20 meeting in July. The proposal is linked to an IETF draft proposal by the main author, Andreas Foglar, which has not been assigned an official IETF number as of now. Innoroute explained to this author that the advantage of the concept was ease of addressing and cut-through routing. Reduction of latency, he said, was key for IoT and machine-to-machine communication in industrial automated networks.

Foglar's company was part of an EU project (Charisma) for which he developed a router/forwarder that uses the e164 address scheme. For the moment, the experimental use was indeed squatting on unallocated address space, Foglar acknowledged. The respective prefixes, which are now also used for a wider test bed, currently were still free (+49, +41, etc).

Certainly, re-using the e164 numbering scheme which the ITU is controlling looks interesting to the UN organisation. Members of the RIPE NCC, on the other hand, are concerned about another edition of a fight over IP address allocation with the ITU.

There have over they years been clashes between the RIRs and the ITU over ideas that the ITU could become a sixth IP address registry. One motivation cited by the ITU was support for developing countries, similar to the now discussed study group 2 proposal.

How to select a RIPE Chair, and other accountability questions

Marseille was the largest RIPE meeting ever with 737 attendees (22% newcomers). If the organisation continues to grow at that speed, it will soon become more difficult to find venues, RIPE Chair Peter Holen

said during the closing plenary. During this plenary, the RIPE community discussed core aspects for a future procedure to select its own Chair. Holen, hand-picked as successor by the late Rob Blokzijl, the first RIPE Chair, said that he certainly did not think he should chose his successor.

There was a broad consensus during the plenary session about using a nominating committee. The nominating committee shall be selected in some way by the Working Group Chairs and the RIPE Program Committee. Members of the nominating committee could be community members, but not necessarily include members from the two bodies.

The task of the nominating committee will be to review candidates and prepare a shortlist for comment. Final selection supposedly shall be made after weighing the community comments. Participants in Marseille also broadly supported a 4- to 5-year term of office for the Chair. At the same time, the community present seemed to prefer a longer total term, so there was a lot of support for two or even three terms. However, there was also broad support for selecting a RIPE Vice-Chair in the future, according to the same procedures.

Chair Hans Peter Holen said in a conversation with this reporter that he thinks a similar, perhaps somewhat more lightweight procedure should be considered for the selection of the working group chairs. Holen acknowledged in a BoF on the accountability of the RIPE community that the lack of a single process for elections of WG Chairs resulted in allowing all WGs to decide for themselves.

Election of WG Chairs is very informal in the RIPE. From the outside, there seems to be a reluctance in the community to run against long-standing Chairs when their term is up. At times, people also seem to avoid criticizing Chairs because they are not prepared to step up themselves. An example for a WG rather patient with this is the Abuse WG, which only rarely sees its second Co-Chair (the WG usually is run by Brian Heanet alone). On the other hand, it was the Abuse WG that experienced the sole procedure to remove a Chair some years ago. For Chair removal, there is also no formal procedure at the moment.

Formalizing or not formalizing – the discussion about accountability of RIPE processes

In general, at least part of the RIPE membership clearly favours a minimum of formal process. Yet as a result from the IANA transition, there was a concern that RIPE procedures could become the topic of an extension of the "accountability" discussion. RIPE and the RIRs had been described as exemplary in terms of accountability, Athina Fragkouli, Head of Legal at RIPE, said. But given the informality of some of the procedures, there was a feeling that structures and rules should be reviewed for both RIPE NCC and also the RIPE community. "We think we are accountable", Fragkouli said. Still, a review and possibly documentation was seen as a proactive step to be prepared in case questions from the outside (for example, governments) would come up. 18 months ago, an accountability task force started to document RIPE processes and bodies.

During the BoF meeting on accountability, the Accountability Task Force also collected feedback to what community members view as core values of the RIPE, how define consensus and how trust could be secured for RIPE in the future.

A draft document is under development by the accountability group and it includes one of the motivations for the exercise of additional formalization of RIPE processes and procedures:

"Respected and trusted community members fill critical positions within RIPE. They follow unwritten processes in a way that everyone feels comfortable with and that preserves the legitimacy of the community. There is a concern that newer community members may come to fill these roles without having an understanding of the spirit behind the RIPE community's procedures. This could erode the community's trust in its procedures over time."

"Traditionally, RIPE participants from comparable backgrounds would gather in an informal manner to make technical decisions with shared goals in mind. Over time, people from different backgrounds and with other interests have also come to participate within the community. It is a principle of RIPE to accommodate everyone with an interest in Internet infrastructure.

However, the values, aims and approach of these newcomers may be at odds with the traditions of RIPE. If RIPE's structures are not sufficiently established/ fortified and accountable enough to accommodate these changing demographics, there could be outcomes that weaken the predictability, trust and legitimacy of the community's processes in the future. Documenting not only processes, but also core values, could help the community to protect against this, by having some kind of an agreed-upon affirmation to refer back to."

During the accountability BoF session, interestingly, answers varied as to what the core values are. Answers included "advocacy for the internet", "making the internet more accessible" to "cooperation of operators and stewardship for the numbering system". Some participants warned against mission creep: RIPE already was the most "expansionist" of the RIRs, said Randy Bush IIG. Daniel Karrenberg, one of the founding members of RIPE, called the open internet advocacy part of the "motherhood and apple pie" stuff that was not originally part of the mandate. The question of how far RIPE and RIPE NCC engages in political advocacy is answered different by different members in the community – see for example a call to the RIPE community by Alexander Isavnin in the Cooperation WG to not let governments like the Russian government "abuse" the internet (see Cooperation WG).

The working document of the accountability task force lists the processes and bodies which need to be covered in a future accountability framework:

RIPE processes:

- WG Chair Selection
- RIPE Chair Selection
- Policy Development
- Emerging "N/A" Issues
- Creation of new TF/WG
- Closing down of TF/WG
- BoFs
- Agreement/adoption of ad hoc community statements
- Removal of chairs
- NRO NC elections
- PC elections
- Plenary programme selection

RIPE structures (bodies and people):

- Working Groups
- Task Forces
- Bird of Feather meetings
- RIPE Chair
- Working Group Chairs
- Program Committee
- Task Force and Task Force Chairs
- NRO NC
- RIPE NCC (secretariat, IP registry operator)

IPv4 dust

In about two years' time, there will be only "dust" left of IPv4 for the RIPE NCC to scratch together and give out to potential new members in need. The RIPE NCC asked the community how to proceed with IPv4 allocation.

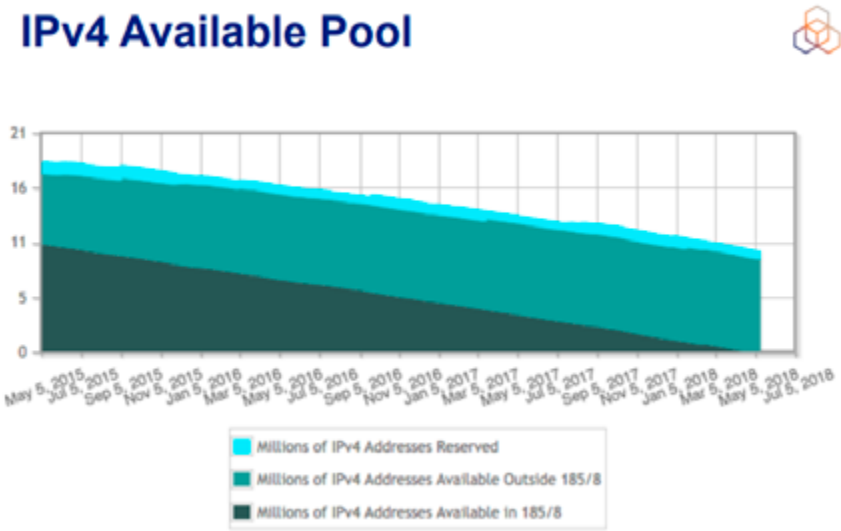

The regular IPv4 addresses have been exhausted. The special pool 185/8 is close to exhaustion after around 7,700 /22 have been handed out to new members. The extra block has pushed numbers of members to over 19,000 by now. Andrea Cima from the RIPE NCC asked members to consider various options for allocating IPv4 addresses in the future.

The RIPE NCC has already started to combine smaller blocks to build /22 allocations and is asking if it should go ahead with this. Another option to stretch the time for allocations would be to reduce the temporary assignment pool, which is currently used for conferences or similar things from a /13 to a /14.

After the final depletion of the RIPE resources, how should recovered space and "dust" be allocated: should a waiting list be governed by a new policy? Another pool quickly drying out is the one for new IXPs and here the question is if the pool has to be stocked up. The address WG has to come up with answers.

Nearly a non-topic for RIPE: GDPR

Several presentations were dedicated to the much-debated implementation of the GDPR during the RIPE week. However, the gist of the presentations came down to a "no need for big changes" mantra for the RIPE NCC.

As the RIPE had dug down into its database policies when implementing the 2006 Data Protection Directive as implemented in Dutch law, important issues were already addressed, the RIPE lawyers said, in particular a clear definition of the purpose for data collection and publication. Also, some limitations have been introduced back then following recommendations of a data protection task force.

The RIPE NCC lawyers nevertheless have addressed a few issues they identified as potentially problematic.

Historical handle and personal information in the data base – so far, handle and personal information of resource holders were also available retrospectively. This will be made impossible by filtering out personal data from historical records.

In addition, queries for the names of individuals, which so far resulted in full objects related to these individuals, will not be possible in the future.

And finally, RIPE request for consent from new members is reconsidered. Currently, those applicants have to provide personal information for the initial creation of RIPE Database objects. Now, the RIPE is investigating potential modifications that would allow to demonstrate that the relevant individual has consented to this processing before a person object is created for them.

A good update on RIPE NCC's steps can be found here.

Naturally, there are still open questions. Consent is an issue, as it cannot be forced. One member also asked for the removal of private street addresses from public query results. These were not necessary to solve problems with the network.

Working Groups, BoFs and Plenary Bits

IoT WG

The brand new RIPE Internet of Things Working Group had its first official meeting in Marseille. It is still a little unclear what the RIPE community's role will be in the IoT environment, which is characterized by various old and newly established industry-led or government-initiated organizations (like the EU initiated AIOTI), who all want to play a role in this sphere.

For the RIPE NCC, Marco Hogewoning explained that the aim was to keep track of developments for the members and also informing the discussion from the RIPE point of view, including answering questions from the media. RIPE NCC is representing the RIPE membership in bodies like AIOTI (Alliance for Internet of Things Innovation) and follows developments elsewhere, for example the ITU.

In one of the interesting presentations of the first IoT WG meeting, Jelte Jansen of SIDN, presented a list of tasks to be taken on for IoT, including:

- better practices for manufacturers
- better (free) standard software libraries?
- international policy, regulation, and certification?
- generate market demand for secure products?
- quarantine bad actors at ISP level?
- educate users
- empower use

As one practical step, SIDN is working on a project that empowers users by offering software to control what individual devices connected in the smart home network can and cannot do. The SPIN project (Security and Privacy in the In-house Network) shall allow users to see what data is sent out by the various devices and the applications bundled with them, and limit the traffic to what the user thinks necessary. Different profiles can be set for the different devices and in the future, an incident reporting to the ISP or domain registry could be realized. According to SIDN, SPIN will allow users to take back (some) control over smart things, and will at the same time protect domain registries and ISPs against DDoS attacks.

On how to better secure IoT devices, Hugo Vincent, leader of the security research group at ARM and one of the experts behind ARM's IoT work, pointed out that the traditional patching, currently the most important tool for operators in security, met with some challenges in IoT. One was that no human users were at hand to patch – and make sure the patches were legit – another was that patches might put heavy strains on the battery life of small devices. Work is under way at the IETF with the SUIT (Software updates for the Internet of Things) work, as Matthias Wählisch, Professor at the FU Berlin and founder of Riot, a small operating system optimized for IoT, said in Marseille. The IETF has for many years worked on small devices in constrained environments (see the earlier work of the Thing2Thing Research Group at the Internet Research Task Force) and produced a whole suite of standards, including Coap and so on.

Vincent predicted that Moore's law would provide for better capabilities to encrypt and isolate mechanisms in smart things (and thereby harden them). A trend to "platformization" was also visible. ARM itself offers IoT as a platform services to customers that do not want to build out a well-maintained platform for smart gadgets/devices of their own. Manufacturers and operators would grow to understand that security is indispensable to turn regulator action away. Vincent was optimistic that the expected large number of applications and users for IoT can be realized. He said his company expected a trillion connected devices by 2035, and productivity improvements across all industries amounting to ~3% global GDP. A white paper with predictions from ARM can be found here.

Domains and IP numbers for addressing in the IoT?

Outreach work for the RIPE and the CENTR community towards IoT manufacturers and providers might be needed in order to convince them to use standard internet technology for IoT,

Sandoche Balakrichena from Afnic recommended. Legacy IoT providers could need help of the domain and IP address communities to move away from walled garden systems and use DNS for naming instead. Balakrichena pointed out that there were already other hierarchical naming systems like the Electronic Product Codes (EPC, with its own root .gs1) or the Digital Handle System (DOI) (practically

based on DNS). Additional new IoT TLDs, like .gs1 and in the future possible .lora (from the LoRA alliance) might be used in the IoT and put a strain on the DNS. Considerations also had to be given to the gateways between the non-Internet IoT systems and the Internet (only with interconnection an internet of things can come to life) and the necessary processing at the gateways (re-configuration, decryption and re-encryption).

DNS WG

Wither DNS Privacy?

DNS Privacy once more topped the agenda with the second of two DNS sessions and several plenary talks addressing various aspects, from implementation status to potential unintended consequences of the DNS privacy work.

A quick look at how operators viewed the various DNS privacy specifications was presented by Vicky Risk (ISC). In a survey she found that 68 percent of the 170 respondents said that privacy concerns of their users were important and 50 percent said they had already plans to implement Qname minimization. Between 10 to 15 percent (depending on the set of respondents) also said that they had already deployed it. While answers were difficult to weigh, said Risk, and some GPPR effect might be there, numbers for DNS privacy didn't look too bad.

Moving away from UDP to TCP for DNS transport could be one first step to ease the move to DNS over TLS, said research student Baptist Jonglez. He compared quality and latency of UDP and TCP connections and found that while there were drops in the processing of queries, with larger numbers of clients (he had 6 million TCP clients participating), Unbound still could handle around 50k queries per second per CPU core.

One proof of concept for DNS over TLS was presented by the RIPE NCC which turned on DNS over TLS (Qname minimization and aggressive caching) for the RIPE76 meeting, giving proof that implementation was possible with some slight tweaks.

The first necessary tweak was to change from the usually used BIND (no support for Qname minimization and TLS as of the time of preparation for the meeting) to Knot. All three tested DNS softwares (BIND, Unbound and Knot) also did not allow to only send the TLS secured answers only to some IPs (using DNS 64).

During implementation with Knot, some funny behaviour was experienced, according to Colin Petri from RIPE NCC, namely the stop of Qname minimization for queries for structured name spaces. For co.uk, after referral and the reception of the authoritative server for co.uk minimization was shut down completely, said Petri.

Deployment is advancing slowly, according to Sara Dickinson, Sinodun. She looked at both DNS over

Mode		Load Balancer	Recursive				
Software		dnsdist[d]	Unbound	BIND	Knot Res	CoreDNS[f]	Tenta[f]
General	QNAME minimisation	n/a	✓		✓		
TCP/TLS Features	TCP fast open[b]	✓	✓	✓	✓		
	Process Pipelined queries	✓	✓	✓	✓		
	Provide OOOR	(h)		✓	✓		
	EDNS0 Keepalive[c]			✓			
TLS Features	TLS encryption (Port 853)	✓	✓	(e)	✓	✓	✓
	Provide TLS auth credentials	✓	✓	(e)	✓	✓	✓
	EDNS0 Padding (basic)			✓	✓		
	TLS DNSSEC Chain Extension						

https://dnsprivacy.org/wiki/display/DP/DNS+Privacy+Implementation+Status#DNSPrivacyImplementationStatus-Servers

TLS and the newer DNS over HTTPS, which seem to be in a race now for the DNS privacy standard in the network. Currently, there are 19 servers running DNS over TLS (DOT), and in addition the 9.9.9.9 server of Quad9 (9.9.9.9) and Cloudflares 1.1.1.1 offer a big resolver alternative. While Google seems to be more set on the DNS over HTTPS solution – which seems to be the natural solution for the browser operators – Cloudflare interestingly offers both solutions. Dickinson also reported that one browser is preparing to send all queries from the browser over DNS over TLS. On the client side, there is also a considerable movement, with Stubby and Android starting to support both variants DNS over TLS and DNS over HTTP, while the Knot and Unbound supporting DOT on the one hand and Firefox DOH on the other side.

While welcoming the added DNS privacy, at the same time Dickinson called on the community to consider the change the DNS is undergoing with these developments. With the outsourcing of DNS resolution to the browser (or their chosen DNS resolvers) DNS was moving away from the current DNS providers infrastructure. Detecting and resolving failure of DNS resolution will become much more difficult, for example. A big question also was how DNSSEC was implemented. While she expected that more and more browsers would be shipped directly with DNS over HTTPS, one question was if DNS resolution in the end would be concentrated in some servers only, which would make nice targets for attacks. Therefore, Dickinson called for a broader discussion on the evolution.

Please update to benefit from aggressive caching!

A plea for operators of authoritative DNS servers to sign and for those operating recursive resolvers to validate was made by Petr Spacek, cz.nic. Using aggressive caching (Aggressive Use of DNSSEC-Validated Cache) helps privacy, but is also a valid tool for recursive and authoritative against subdomain attacks. This was illustrated in a test run with three domains under .cz, as presented by Spacek.

The tests performed showed that expectations that the use of NSEC/NSEC3 resource records to allow DNSSEC-validating resolvers to generate negative answers within a range and positive answers from wildcards in fact helps to decrease latency and resource utilization and increase performance. While CPU utilization of the attacked system was still on a record, random sub domain attacks to the authoritative site could be nearly eliminated, with NSEC being more efficient than NSEC3. Knot Resolver 2.4 will include NSEC3 support for aggressive cache.

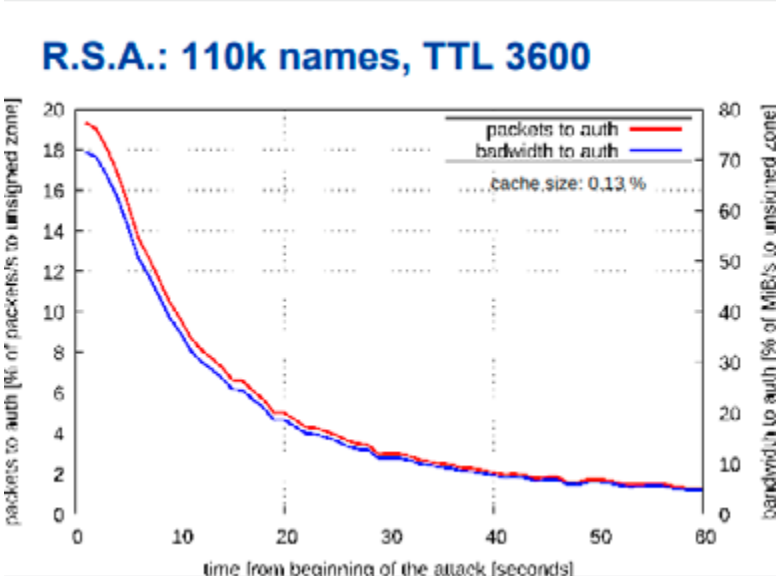

A measured roll of the crypto-algorithm for DNSSEC signing

An SIDN research project measured delays and failure rates during the successful DNSSEC algorithm roll-over of the .se zone in December. .se has been the first DNSSEC-signed zone back in 2005, five years before the Root Zone was signed. Only Sha1 had been available then, but is no longer considered secure enough today. With .se being a zone with a high rate of signed domains, (50 percent of the 1,4 million domains are DNSSEC-signed) and also the highest percentage of validating resolvers, in 2017 the .se registry IIS decided to change from RSA-Sha1 to RSA-Sha256. Newer Elliptic Curve Crypto (ECC) algorithms were decided against by IIS due to a lack of current propagation in validating resolvers. In order to allow for the automatic algorithm roll-over provided by OpenDNSSEC (ODS), IIS first changed from ODS 2.1.3, which then allows to roll in five steps:

1. Update KASP database
2. Generate new keys
3. OpenDNSSEC will start using the new keys, publishing KSK, ZSK and double signing all records
4. Publish new DS in parent and remove old DS from parent
5. Drop old KSK, ZSK and RRSIG from zone

At the RIPE meeting, Moritz Müller from SIDN Labs described the measurement and monitoring project (which will be published in a paper later on). Using 10,000 RIPE Atlas probes, SIDN Labs measured a publication delay for the keys using the new algorithm of 10 minutes. Propagation delay measured was 48 hours. Following a TTL of the DS and root 24 hours was expected, but a very small number of resolvers (roughly 1 percent or less) took another day to update. Müller said that if registries rolling the crypto algorithm wanted to make sure every resolver picked up the new DS, they should consider 48 hours and ten minutes. While IIS took even more time, that was conservative and presumably not necessary.

The paper on the measurement which will allow other registries to use the set-up will be published here. More information on the role of IIS is here and here. RIPE NCC's earlier algorithm roll is here.

Truncation, fragmentation and ATR measurements

APNIC researcher Geoff Huston considered the net benefit of a new draft in the IETF pipeline that intends to cut short the time for a fall-back from UDP to TCP after fragmentation. ATR (Additional Truncated Responses) hopes to speed up DNS resolution by adding a truncated response just after a fragmented response. Instead of waiting and retrying, the mechanism triggers an immediate switch to TCP, which can handle the longer responses.

In a large measurement campaign for IPv4, Huston saw failure rates of 40 percent of resolvers for large packets in UDP and 21 percent failure to do TCP at all. IPv6 failure rates reached a staggering 50 percent for UDP and 45 percent for TCP fall-back.

Mitigation by ATR was visible, but not too high. In IPv4 the ATR loss rate was 29 percent (so a little over 10 percent of resolvers that were incapable of receiving a fragmented UDP response were able to switch over the TCP). For IPv6 ATR failure rate was 45 percent (only making the situation better in 5 percent of the cases. Counting users instead of resolvers, the drop was a little bigger, according to Huston. Yet, when weighing the benefits of ATR, it had to be considered that it also allowed for better DDoS attacks due to added traffic.

Protocol	Visible Resolvers	Fail Large UDP	Fail TCP	Fail ATR
IPv4	113,087	40%	21%	29%
IPv6	20,878	50%	45%	45%

Table 1 – Failure rate of visible resolvers.

Huston concluded from the measurements that the internet was "pretty broken" in the first place. Adding additional "tricks" like ATR while mitigating the problems was also more straws on the

Abuse WG

The Abuse WG had another discussion on an addition to the Abuse Contact policy, which in the future will oblige the RIPE NCC to validate the much-discussed abuse contact email once a year. Initiated and presented by representatives of a French mobile operator and Europol, the addition in essence adds one sentence to the existing policy:

"The RIPE NCC will validate the 'abuse-mailbox:' attribute at least annually. Where the attribute is deemed incorrect, it will follow up in compliance with relevant RIPE Policies and RIPE NCC procedures."

Thomas Schmidt, Head of Policy at the RIPE NCC, reported that about 10 to 25 percent of the existing abuse contact addresses are wrong, which represents about 70,000 individual addresses.

During the debate in Marseille, both proponents and sceptics of the policy reiterated well-known arguments from the policy debate that originally introduced the abuse-C obligation. Proponents want to push for correction and also, where needed, closure of LIRs in an effort to have a "clean" data registry. Opponents see the potential for closures following false positives as an issue and consider the policy a method to provide false security. A considerable number of fat finger-problems (wrong email addresses listed in the field) could be erased, yet those up to malicious behaviour will not necessarily be caught by annual validations of an abuse-C contact.

For the time being, legacy holders will not be included in the validation efforts and these efforts, according to Schmidt, will follow an automatized procedure in order not to overburden the RIPE NCC. The latter resulted in more questions asked about what type of problems will be caught. Meanwhile, the policy has been adopted and implementation by the RIPE NCC is under way.

Should RIPE position itself against state interventions harmful to internet infrastructure?

Should the RIPE also discuss a policy that sanctions governmental abuses of the internet? Alexander Isavnin, representing a loose group cooperating under the title of the Internet Protection Society (IPS), hinted at such a discussion during the Abuse WG session. A proposal from a number of African members calling for sanctions against African governments for broad blockings on internet traffic (during or before elections, for example) had resulted in fierce debates at the Afrinic meetings last year.

Instead of the RIPE membership saying it was only "technical", Isavnin asked it to take a stand against the growing interference of state actors into content delivery and routing.

The IPS was not incorporated, which allowed it to campaign for internet rights. Isavnin described in his presentation the spiral of Russian legislation toward additional control over content and also routing. Legislation still under discussion on critical infrastructure intends to establish a national TLD and a national Internet Exchange (currently owned by RosTelecom). The legislation also included an enforced routing registry, which was declared to be a government controlled safe version of the RIPE data base, and Bgp blackholing as a blocking method.

Isavnin also described a number of cases before the Russian courts, namely the Telegram case in which the Telegram should be obliged to hand over encryption keys and the jailing of Tor node operator Begatov. Begatov was freed after President Putin declared (in relation to potential Russian interference in the US elections) that IP addresses were not a sufficient source of evidence.

Isavnin's idea about the RIPE positioning itself against such state interventions, let alone a formal policy to sanction state internet abuse, will quite definitely be met with staunch opposition at the RIPE, which had already criticized the Afrinic draft proposal. It is not seen by the majority as the task of the operator community, despite the fact that Isavnin pointed out that some interventions were harmful to the infrastructure and the business of technical operators.

Projects the IPS is currently pursuing are the Internet Freedom Index, the Internet Connectivity Index, a "Repressions mapping" and campaigns to create awareness for caveats of Russian Internet regulations.

Cooperation WG

Chris Buckridge from RIPE NCC underlined the noticeable trend of the EU legislature to broader regulation of internet communication. The aim is to tackle specific issues such as illegal content and cyber security on the one hand and address the more general problem of the cross-border nature of the net on the other hand.

E-evidence, IG and regulation

RIPE NCC is working with a consultant company to keep up with the developments and expects rather big changes for its own operations from the e-Evidence directive currently in the legislative

process in Brussels/Strasbourg. The core change for RIPE NCC will be that instead of answering only to Dutch Court orders (to give access to data or evidence), it will have to answer to Court orders from all EU member states.

Other legislative developments the RIPE NCC is following closely is the implementation of the NIS Directive (which some EU members have decided will extend in scope to the domain name operators, ccTLDs, and in RIPE's case its DNS root server service) and also the developments in regulating measures against illegal content. On the latter, there are recommendations from the Commission, and a future regulation is under debate after a consultation.

Using multi-stakeholder norm-setting and multi-stakeholder enforcement of these norms (for example at ICANN) should be considered for securing the public core of the internet, professor Joanna Kulesza (University of Lodz) recommended at the Cooperation meeting. There was a need to counter a trend to relay on "shooting back": better to "hack back" laws on the rise, she said, pointing to the proposed US AC DC (Active Cyber Defense Certainty) Act as one example. The Act would allow companies under attack to hack back after being granted so by the FBI.

Kulesza said that while international law-making (Cybersecurity Framework Law) might take another 10 to 20 years, more easily available options were lending from the existing standards and norm-setting bodies (to which she added ICANN and also RIPE) and look for soft law to be agreed upon by countries and stakeholders. A new body for IG discussions was not in the cards, though, so enhancing cooperation between the various bodies was necessary.

One body promoting this is the Global Committee on Security in Cyberspace, which has called for protection of the core of the net end of last year, while also defining what is part of this core.

Meanwhile, the UN itself has reacted by setting up a new multi-stakeholder Commission that was tasked by UN Secretary General Antonio Guterres to establish a broader debate about norms for the net (there will be consultations and at least two events after September 2018) and in the end come up with recommendations. That body is chaired by Melinda Gates and Jack Ma.

BoF on IP/ASNs for Governments

Representatives and experts acting as consultants for a number of governments got together in a BoF to discuss special IPv6 allocation needs for public administrations. The organizer of the BoF was Iljitsch van Beijnum from Logius, an agency of the Dutch Ministry of the Interior and Kingdom Relations. According to van Beijnum, the purpose was to consider potential changes in address allocation and ASN assignment policies for the public entities which often comprised many independent sub-organizations with their own internet connections.

Experiences with the allocation policies of the RIPE were shared by Tahar Schaa, consultant for the German Ministry of the Interior and Jordi Palet, Consulintel, who helped the government of Spain and other governments with their allocations and addressing plans.

In Germany, public address planners realized after an initial allocation round that the allocation of a /26 was too small to satisfy the various ministries and federal states. According to Schaa, the larger federal states needed /32 subnets to in turn start to build their subnets for city governments, police and other state institutions.

In the end, the German government, being a LIR since 2009, applied for a larger-size block and just the week after the RIPE meeting, was allocated a /23. While some renumbering might be possible, the original block had been chosen so that neighbouring space was available for the larger block of continuous address space. While the German example illustrates that once started IPv6 address need might be considerable, it also shows that IPv6 is advancing slowly, with Schaa reporting that some lobbying had been necessary.

Van Beijnum said the problem he called for the BoF was the difficulty to get public AS numbers for private use, needed by various government entities in The Netherlands. For address space, the Dutch address planners had used a trick to receive space beyond their original /28 from which they, too, cut subnets for various public entities including 280 cities in the Netherlands. Instead of giving back the /29 to receive a new /28, the Dutch Ministry of the Interior became

a LIR and applied for its own /29 before closing down that LIR again to consolidate the resources.

On the question of how to receive the AS numbers for private networks of public authorities, a change of RFC 1930, which dates back to 1996, was considered. Certainly, given the fact that the number of globally routed AS numbers should not go overboard, there was a need to find a middle ground between the needs of public authorities and the number of ASN allocated.

The next RIPE meeting will take place in Amsterdam on 15-19 October 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