Council of European National Top Level Domain Registries

Highlights of RIPE 70 Amsterdam

11 - 15 May 2015

more to come ...

Table of Contents

Root Zone KSK-Rollover: Checking the brakes of the car

ICANN, VeriSign and the NTIA have re-started the initiative to roll the root zone key signing key. Edward Lewis, Senior Technologist to the ICANN CTO presented what is an ambitious time plan at the RIPE70 and the preceding OARC meeting.

Plans for the root zone KSK rollover (every five years) which is part of ICANN's contractual obligation as IANA operator were first considered in 2013 when ICANN started an earlier public consultation on the ZSK rollover. SSAC followed up with advisory 063. The initiative was shelved, though, presumably because of ongoing work with the introduction of new Top Level Domains during that time.

With five years now over since the signing of the zone, the partners involved in the cryptographic protection of the root zone, ICANN, VeriSign and the NTIA, have agreed to go ahead with the KSK rollover. Interestingly, Lewis answered a question of this reporter on the relation of the rollover and the IANA transition efforts, that the rollover came close to being a pre-condition.

Perhaps this explains the quite ambitious time plan of the activity. The members of the new design team for this activity (Joe Abley, John Dickinson, Ondrej Sury, Yoshiro Yoneya, Geoff Huston and Paul Wouters, in addition to members from the „partners“) are expected to present a draft proposal for the key rollover in June, to allow for the public comment period to be opened around the ICANN meeting in Buenos Aires. After a four-week comment period a final report shall be prepared in August, go through ICANN, VeriSign and NTIA again and then be executed.

During the OARC meeting there was a little more in-depth discussion about concerns the design team will now look into. These are mainly large answers and truncation (due to serving two keys during rollover, and especially with IPv6; see for a good explanation SSAC063), validation failure due to lack of support to RFC5011 (the standard for automated trust anchor updates) and potential risks resulting from algorithm change.

The biggest problem according to Lewis was with validating resolvers as the group had no way to know which ones would fail to validate due to bugs or wrong implementation. Given the fragmented resolver market „we have no chance to know, and i think we should not know“, Lewis said. Sury said while the Design Team assembled knowledge on various platforms (Debian, Red Hat, etc) there was no way to foresee problems with the many distributions in the wild. A concern were large organizations like telco companies who had not touched the keys in their systems for years.

The experts expect some regional differences due to the number of validating servers, and also the provenience of Google DNS, which in turn is expected to perform just fine. Statistics about how much validation is done globally are sketchy. Stats about validation by Geoff Huston are here:
http://stats.labs.apnic.net/dnssec. They show quite some validation (bigger rate than DNSSEC signing in many parts of the world), yet could be biased as they are derived a Google-ad based survey.

Adding new High Security Modules

A second effort currently underway is the renewal of the Hardware Security Modules (HSM) for tamper-resistant storage of the KSK. Reasons for substituting HSMs are potential battery failure (battery life time is said to be 10 years) of the 2010-dated HSMs, and the fact that the warranty period ends in 2015. ICANN explained in the announcement of March, 23 that the project was „distinct from the project to replace the existing Root KSK with a new Root KSK“ and „separation of the replacement“ would „allow time for the Rollover Design Team to fully develop their approach without being influenced by operational pressures relating to HSM replacement.“ Lewis said to this reporter that two new HSMs were „added“ to the Culpepper location during the April,9 signing ceremony. During the next signing ceremony in El Segundo

on August, 13, two more will be added.

The new HSM, AEP Keyper Plus, has been selected from the same vendor. Lewis did not come up with an answer to questions if there had been a tender or considerations to choose an alternative HSM. On tamper-proofness in the light of the Snowden revelations (on question from Shane Kerr, Beijing Internet Institute), Lewis made more general remarks, noting to this reporter that DNS data was more fleeting, so no resistance for several decades would be expected as it was for highly classified data.

Mehmet Akcin, then ICANN and now Microsoft, came up with what he remembered as the rationale for choosing the AEP Keyper in 2010, which that it was seen as the only model meeting FIPS 140-2 Level 4 security certification. This security level is a requirement for the Root KSK in the IANA functions contract.

Algorithm change, or not?

The new model, according to Lewis, was also allowing to import the Root KSK from the older models without needing to regenerate the Root KSK. The question if there should be an algorithm change when exchanging the KSK is still under discussion in the design team and the DNSSEC partners, according to Lewis. The new hardware would allow a later change. One major question here is if the parties want to move away from RSA to ECDSA.

Currently in use is a RSA 2048 bits key pair. While ECDSA would allow for shorter, more efficient keys compared to RSA, Geoff Huston in an analysis of potential client behaviour (using Google-adds survey again) comes to the conclusion that the move at this point in time could result in a decline of validation, because even when clients support for ECDSA has moved up considerably, resolvers were lagging behind (see Huston's ECDAS vs RSA analysis, presented both at RIPE and OARC). VeriSign, according to Duane Wessels, is currently not recommending an algorithm exchange and recommends to keep the ZSK size at a maximum of 2048 bits (for experiments with different key sizes, see Wessel's OARC presentation). We shouldn't forget either that validation of ECDSA signatures is computationally more expensive for validating resolvers than doing so with RSA.

A question posed to Lewis during the RIPE meeting was on how the DNSSEC partners would get the word out on the ZSK roll over to create awareness, which Lewis said was on the agenda.

Address Policy: give us more v6 and give us more v4, please

For both, IPv4 and Ipv6, requests were made to allow allocation of bigger chunks of addresses. One policy proposal of the British Ministry of Defense and German retailer Kaufland asks for exemptions from current documentation policies for large organizations. Problems of the growing transfers market were discussed based on a plenary talk by Jim Cowie, Chief Scientist at Dyn.

Troups, Retailers, Countries want bigger IPv6 chunks

For IPv6 German large retailer Kaufland engaged in an unusual cooperation with the British Ministry of Defense to call for special treatment for very large organizations when it comes to address allocation (the partnership was somehow sponsored by RIPE itself as both parties asked for big IPv6 chunks). According to standing policy the RIPE NCC can assign a /32 as standard initial allocation and, on request, organizations or companies can receive up to a /29 of IPv6 resources. The authors of the new policy proposal, Alexander Brinkman (Kaufland Information Systems) and Mathew Newton (British Ministry of Defense) argued in Amsterdam that for their organizations a /29 was a „showstopper“. At the same time they were unable to come up with the kind of required documentation that would allow RIPE NCC to hand out bigger blocks to them. Especially the obligation to document the number of customers was a problem for them.

The MoD representative explained in several written comments that his organization was unable to provide that kind of documentation as it was classified information. He just pointed to the size of the British Forces (among the fifth largest globally, with a budget of 52 Billion Euro) and the need for a large block enabling hierarchical assignments to the various units (air, sea, land, space). Brinkman argued more with the number and geographical spread of locations of Kaufland (which has kept adding other retailers in recent years).

There is quite some difference between the two proposers, since Newton underlined that the vast majority of external routing announcements the British Forces would make would „only be visible to coalition partners (other nations' military networks, NATO infrastructure, etc) and will not appear in the routing tables". Brinkman on the other hand seemed to go the other way explaining the need to announce all locations directly.

The problem with announcing said 10.000 locations in the public routing table would put pressure on network operators to go to larger routers, Address Policy WG Chair Gert Doering explained. Several participating members at the meeting warned that to go easy on large requests could approve bad address planning, too. Also, if multinationals would announce their locations under global prefixes geographical filtering (to keep one's routing table manageable) would become more difficult.

The policy proposal to allow for exceptionally large allocations nevertheless received quite some support from the government camp. Switzerland, who had its request declined earlier from the RIPE NCC, was supportive and Tahar Shaa, consultant for the German Ministry of the Interior got up during the WG session in Amsterdam declaring support from the German and other governments (he did point to Switzerland and Spain, in particular).

A higher burn-rate for IPv4?

With regard to IPv4 address block allocations from the last /8 block, address-trader Elvis Velea, V4Escrow, and Radu-Adrian Feurdean, Coriolis Telecom, laid a proposal on the table to increase the size of the final allocation to RIPE members and also newcomers.

Every RIPE member receives a final /22 from the last block, according to RIPE's so-called „soft landing policy". The policy has resulted in a fast growth of RIPE, with 1500 new members from May 2014 to May 2015. The total number of members now has reached 11750.

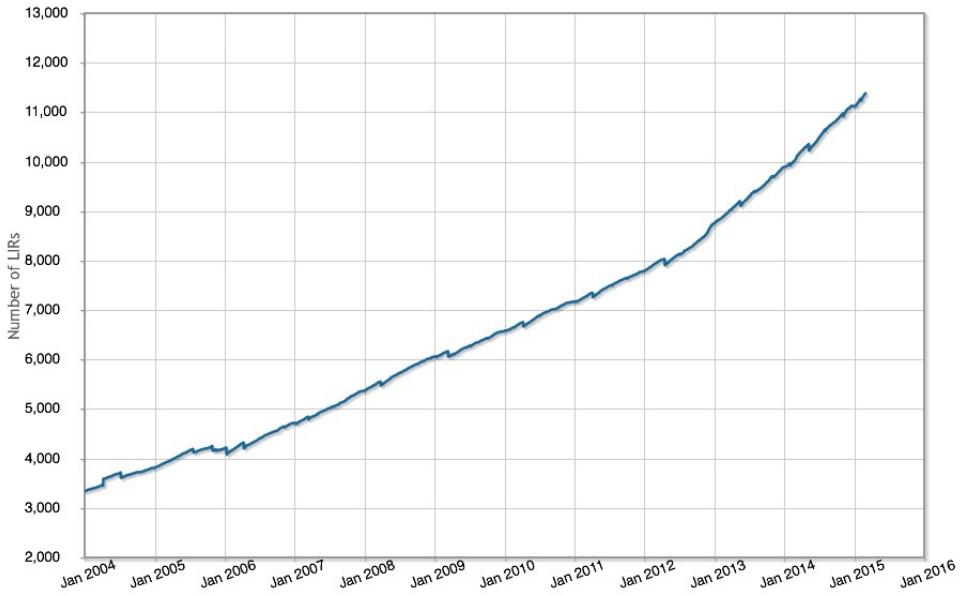

Feurdean propose to allow for more flexible allocation sizes and even hand out /21 or /20 blocks, when old and incoming members can document the need. The two proposers for one pointed to difficulties for some people due to the current limits on block sizes. They even argued that three years after running out of IPv4 RIPE in fact has more addresses than before because it had been allocated recovered resources from IANA. IANA is distributing recovered blocks to the RIRs two times a year (March, 1 and September, 1), handing out equal shares. If sticking to the current policy, RIPE would not run out of IPv4 for some time, while keeping additional v4 allocation small and potentially painful. Velea and Feurdean promoted their idea with the argument that a higher burning rate would signal to everybody that IPv4 was out and IPv6 had to be used.

While more flexibility in allocations from the last block was seen as potentially beneficial by several participants, there also were reminders that, if burned slowly, future incoming companies or organizations will still be able to receive small IPv4 allocations and configure for dual stack.

Shopping Ipv4? Buyers beware!

Despite the tiny net growth of the IPv4-pool in the RIPE meeting, IPv4 remains scarce, even more so in APNIC, ARIN and LACNIC. APNIC is down to 0.75 of a /8 (having reserved numbers for newcomers only), LACNIC to 0.28 of a /8 and for ARIN, Aaron Hughes, ARIN Board of Trustees reported, that the registry was at 0.19 of a /8 with the still available blocks (one /11, 13, 14, 16 each, plus smaller blocks) „going fast", Hughes said, he would describe ARIN as „out of IPv4". address reserves would never go down to zero really, he said, yet for now the only question was if the registry had to decline all or most address requests. „We are pretty close to the point where it is an all-transfers market", Hughes said in Amsterdam.

The pitfalls of IPv4 transfers transfers were highlighted in a plenary talk by Jim Cowie, Chief Scientist of Dyn. Between 17 October 2012 and 1 May 2015 a total of 2,252 unique address blocks (14,526,720 IPv4 addresses, equivalent to 86.6% of a /8) have been transferred in the RIPE region, according to Cowie's statistics. The transfers market had been quickly picking up since last fall (see also RIPE NCC listing service), with a peak of transfered addresses in the RIPE region in November 2014 (380 blocks).

The big problem of the transfers, according to Cowie, is that some of the sold blocks are accidentally announced twice in the routing table, which sends queries toward different networks around the world. The problem with IP-addresses was that, contrary to a used car which only the buyer could use once in possession of the keys, IP addresses could be „driven" elsewhere and the old owners still could receive „speed tickets" for them. The contradicting routing announcements could result either from accidental mistakes or from malice of address sellers.

Cowie reported on one exemplary case, a /17 range at 46.51.0.0. After being transfered from Netserv Consult SRL in Romania, to Mobile Communication Company of Iran, the Iranian Mobile Provider began to announce the prefix under AS197207, but Level 3 under AS 3356 had announced more-specific prefixes (46.51.16.0/21, 46.51.24.0/21, 46.51.32.0/21) within that range since early 2012. Users who tried to announce specific ranges within the /17 were black-holed, mobile users trying to reach content in the US from these spaces were unable, too. To get control Mobile Communications of Iran announced even more specifics in December. Some prefixes within the range as of today are announced by both Iranian and Romanian entities, resulting in a split view of what sits at the respective addresses.

His presentation also gave a very interesting overview over regional trends (see also nice RIPE labs story on these) for the RIPE region, with Romania currently being the biggest „exporter" of IPv4 space and companies in the Arabic peninsula being the big „importers", especially address hungry mobile telecommunication providers, like the Iranian Mobile and Fixed Network Communication Company, Saudi Telecom and Emirates Telecommunication Company under the top buyers. 930 out of 1856 (50%) blocks transferred since January 2014 were from Romania and 33 percent of the 4500 prefixes that today originated in Saudi Arabia had been Romanian only a few month ago.

With transfers between the five RIRs possible in the future (RIPE finalized its respective policy in April, see RIPE 644) the situation could get even more confusing, Cowie suspects, and made several recommendations to buyers:

- Research historical routing of prefix for sale (including more specifics)
- Don't forget DNS, are fully qualified domain names pointing to your prefix.
- Configure aggressive routing alarms on purchased prefixes via 3rd party service
- Establish strong technical contacts within the seller's organization
- Brokers may want to explore 'clean routing' assertions (even clawbacks?)

Intercessional on special names: Many concerned against IETF stepping on ICANN's turf

Two years ago a document on the allocation of „special names" passed the IETF RFC document stream as an individual submission written by Apple Engineer Stewart Cheshire. It described guidelines how the IETF should assign TLD-like names, with the IESG the body in charge to grant the TLD-like names expected to not being resolvable in the DNS. Now, that RFC 6761 bites back at the IETF, as it has to deal with a list of applicants for alternative/innovative functionalities (.onion, a list of p2p service related names) and names deemed as potentially creating confusion when delegated by ICANN in its gTLD application process (mail, home, corp; potentially also .onion due to its installed base).

Nothing has been decided at the long planned intercessional teleconference of the IETF DNSOP WG, which took place

during the RIPE 70 week, and was attended jointly by 16 DNS experts gathered at the RIPE meeting (for the meeting minutes are here, one nice summary from one of the P2P supporters is here). Deciding over the different proposals was in fact not on the agenda set by the WG Co Chairs Suzanne Woolf (ISC) and Tim Wicinski (Salesforce.com), as it is not fully clear how the applications for the „special names" should be processed by the IETF and IESG.

Will WG adoption of the drafts on special name applications result in an automatical approval by the IESG for the respective special names? Will adoption of the documents make an IESG approval more likely? Or will the deliberations of the WG have no bearing? No answer to these questions was given during the meeting. The process forward is not yet clear at this point.

The role of the IETF in names

Woolf noted during her introduction to the debate that basis for the process was RFC 2860, which is the Memorandum of Understanding between IETF and ICANN. The RFC pointed out, Woolf said, that ICANN had responsibility over policy, but IETF had a role in assignments of domain names for technical uses. RFC 2860 reads:

4.3. Two particular assigned spaces present policy issues in addition to the technical considerations specified by the IETF: the assignment of domain names, and the assignment of IP address blocks. These policy issues are outside the scope of this MOU.

Note that (a) assignments of domain names for technical uses (such as domain names for inverse DNS lookup), (b) assignments of specialised address blocks (such as multicast or anycast blocks), and (c) experimental assignments are not considered to be policy issues, and shall remain subject to the provisions of this Section 4. (For purposes of this MOU, the term "assignments" includes allocations.) In the event ICANN adopts a policy that prevents it from complying with the provisions of this Section 4 with respect to the assignments described in (a) - (c) above, ICANN will notify the IETF, which may then exercise its ability to cancel this MOU under Section 2 above.

Earlier TLD reservations by the IETF included

1. *.test (testing of current or new DNS related code)*
2. *.example (use in documentation or as example)*
3. *.invalid (online construction of domain sure to be invalid)*
4. *.localhost (traditionally been statically defined in host DNS implementations*

These reservations and the RFC reserving them (RFC 2606, June 1999) predates ICANN. RFC 6761 updates 2606 and sets out seven critieria to be discussed in the process of adopting a potential candidate for a „new name". RFC 6761 sets out an IETF standards process or an IESG designated experts process that would describe the new functionality and would check how the new name would have to be handled by DNS operators, registries, software developers, etc (said seven criteria to check) and if it had technical merit. Woolf spoke of „innovation" and „interoperability" for the DNS.

Current applications pending in front of the WG are (a nice list is here):

.HOME, .CORP, .MAIL (Chapin and McFadden) for preventing name collisions between names used in private networks and applications filed with the ICANN in its first regular new gTLD round. Considerable debate is going on about the question if the IETF should not stay clear from what some say would be policy-laundering from ICANN. While the authors pressed the „operational" and therefore „IETF-nature" concern, IAB Chair Andrew Sullivan finally got explicit in an email following the intercessional:

„The point that I keep trying to make is that, if that's what we think, we should _not_ be attempting to use DNSOP or the special names registry as a policy-preference enforcement body. If the issue is that you don't think ICANN will do the right thing in managing the policies of the root zone, then you need to go work on ICANN, not try to use the IETF as a second control. Doing that puts the IETF itself in jeopardy."

.ONION, a name in use by Tor routers (which claims to be in use by around 30.000 nodes). Applicants including activist Jacob Appelbaum want the name to be blocked from future gTLD rounds and serving NXDOMAIN answers. The .ONION draft was split from the P2P draft due to a deadline of the CA Browser Forum with regard to the use of certificates. Starting October, 1st 2015 only delegated TLDs will receive certificates.

PTLDs: a set of several P2P domains including "GNU", "ZKEY" (for the Gnu Name System), "I2P" (for Invisible Internet Project) and "BIT" for the dot.Bit (Namecoin timeline) project. Still listed in the P2P draft are both .ONION (Onion Routers) and .EXIT (exit Nodes) for the Tor Project, which meanwhile have their own draft. Statements were made at the intercessional that the IETF should not accept a „batch"-application. The applicants shot back after the meeting arguing that their draft made the case P2P resolution of names (the applicants, especially Hellekin O. Wolf, speak of pTLD) – contrary to the hierarchical DNS resolution, plus at least some of the names were part of one system (like .TOR and .EXIT, or .GNU and .ZKEY). Partly the names were architecturally linked, Christian Grothoff (INRIA) noted in an email to this author. Yet there is discussion by the group to split the documents if need be.

.ALT, as a potential was for the IETF to create a space for all the future non-DNS like name experiments or applications. There was very broad support for .alt from all sides and camps. Jonne Soininen, who participated in the intercessional as the IETF liaison of the ICANN Board, welcomed .alt as an experimental place. For the IETF the proposal looks like one potential way to evade competition with ICANN.

Two basic camps

In essence two camps seem to be out there (apart from the applicants who hope to receive the delegation or reservation by the IETF). One questions the wisdom of the IETF getting into names and therefore stepping into a policy-heavy area. IAB Chair Andrew Sullivan seems to lean on this side (see above). Peter Koch, DENIC, also argued that the collision domains were dealt with by ICANN and the IETF had not reason to engage in „policy-laundering". Koch even questions the reasoning for .onion because of the rather political motivation (pushing for more privacy), which has considerable support and looks for the observer like the one sure winner from the process. An argument mentioned for example by Warren Kumari (Google) with regard to deployment numbers was that it would allow to easily game the system, use bots to propagate queries to a certain, undelegated name and stop their competitors from gaining a delegation from ICANN.

The other camp includes quite some supporters for .tor, with supporting a privacy friendly space as major motivation. To prevent collision and confusion for some observers is a technical and operational issues, therefore .onion's deployment in the wild and also the use of .mail, .corp, .home in private networks (and leaking thereof) must motivate the IETF to act. The most extreme position on the IETF should delegate special names-camp might be the one represented in statements like the one of John Levine who argues „this isn't an ICANN issue, it's an IANA issue. ICANN can't sell .corp, .home, and .mail for the same reason they can't sell .arpa or .invalid: they're already spoken for". An argument rejected by Sullivan, for example. Would the IETF tend to not approve any of the special names, certainly RFC 6761 would be called into question, said .tor and .p2p supporter Hugo Maxwell Connery, Head of IT from the Technical University of Denmark.