SSR2 Subgroup

DNS SSR

Date: Mon., 14 August 2017

Time: 20:00 UTC

Attendees: Geoff Huston, Eric Osterweil, Kim Davies, Yvette Guigneaux, Bernard Turcotte, Karen Mulberry

GH: First set of questions are around TLD label management

KD: Two-character TLDs that are not ccTLDs would have to be a policy decision which would need to be approved by the Board.

GH: Under the existing policy framework – would you allow a single character Unicode which would be a multi-character ASCII.

KD: At the top-level, a single character in either Unicode or ASCII would not be allowed by current Policy. We understand some of the community has concerns about single character TLDs.

GH: Underlying this are the RFCs including the LDH rules – are these overarching constraints or can the policy go beyond that?

KD: This is an area that has lacked full exploration. As policy is developed Staff are typically involved, including technical specialists, providing an opportunity to flag potential issues that conflict with technical standards or have other implementation issues. Should the community approve policy that conflicts with an RFC, and if the Board approved, it would be sent to staff for implementation and the staff could flag implementation issues which could go back to the Board for further review. It is unlikely we would get that far down that path before the issue is identified.

GH: Unicode and the IDNA rules. Is this a grey area? EU in Greek?

KD: IDNA specs are straightforward. The only unusual issue at the moment is that the IAB has asked IANA to not publish IDNA tables beyond version 7.0.0. Are you looking at something specific? When new domain strings come to IANA their eligibility as legal IDNA strings is usually pretty much decided elsewhere, we are there to implement.

GH: This comes about my concern about emoji characters in 2nd and lower level domains but the Unicode is displayed differently on different systems. What criteria are you applying?

KD: Today the only way to get IDNs in the root is one of two ways – either via ccTLDs and it has to be a meaningful representation of a country name. In gTLDs there are evaluation process. By the time IANA gets them they have been vetted by one of these two processes. Of course, it

has to be IDNA conformant to go into our systems and therefore emojis are not allowed. ICANN is working on implementing variants and root LGR efforts – what we expect is that eventually the root LGR will be an additional constraint on IDN registrations in the root zone.

GH: Next three sets of questions which point to the fact that there is some coordination of ccTLDs with ISO3166 and gTLDs. Any input from staff or is this simply applying policy.

KD: We have a relationship with the ISO 3166 Maintenance Agency. This is currently a topic of discussion in the ccNSO. We use the ISO 3166-1 standard to identify countries and the two letter code for an ASCII ccTLD. For IDN we just use it to determine what is a country and then there is another process for selecting the string. For "exceptionally reserved" codes we have 3 buckets: those delegated prior to a 2000 Board resolution that set the rules for their use, codes that meet the 2000 Board resolution, and codes that do not and those are being phased out. ISO approached ICANN to participate in the administration of the standard about 10 years ago. We contract a non-staff member to be our delegate – we have a small coordination group internally which liases with our delegate and we have a company policy that our delegate always abstains from votes on adding or removing codes. Does that cover it?

GH: Does this cover all possible ASCII 2 letter codes or just ccTLDs?

KD: Given there is, on average, at least an annual change to the standard makes it impossible to rely upon a static view of the standard – therefore as a practical implementation all 2 letter alpha codes are considered as being covered. There are some codes marked as "user defined" that could be used without risk of collision with future countries but that has not been explored by the community to my knowledge.

GH: Anyone else have other questions on this topic. (none). Onto changes to the DNS, the records in the root. When you get a change request how do you know its authentic?

KD: We do not need the submitting party to have authority to approve the change. We have two mechanisms to verify it should proceed: Firstly, we require positive confirmations from the admin and tech contacts which is automated via our system. Second for NS and DS records we cross match with the child zone and make certain it matches, which proves it is supported by the party that has custody of the TLD zone. Those are the primary mechanisms but all RZ changes are manually verified by both PTI and Verisign and it needs to pass a smell test. For example, replacing a full NS set - this would be evaluated as a potential change of control of the TLD. Also, worth noting that we are developing a new system for authorization which is more flexible which could allow private parties which are not listed in the WHOIS to approve change requests. You could also have any number authorizing a change. Under development and getting feedback from users and hope to document the proposed approach by the end of the year.

GH: There has been exploration of different crypto algorithms. Do you check the DS and NS records work? What happens if they fail?

KD: NS records - we check all of them and none can be lame, that glue records match, not open recursive, whole set of names fits in 512 bytes without truncating. Some of these are absolute requirements, others we have latitude and will make a call after talking with the operator. We

also check serial numbers of NS's but may want to refine the approach in the future for less false positives. DS records vs keys. We do a validation of the RRSIG of the child and we expect this to become mandatory in the future. Re new algorithm types, we need to be able to support them. The current ones we support were from 2010 but were limited because of the NTIA contract. We are now looking to implement changes with Verisign and hope to have announcements later this year. We do not expect to support 100% of algorithms - we are adding new algorithms based on maturity and implementability because the RZ is critical infrastructure.

GH:Helpful. You make the DS or NS change but the user systems fail – what is your system for detecting this and what do you do?

KD: We do not do regular monitoring of these things for data quality and health. The reason for this is the last time we considered this a decade ago there was some serious sensitivity about such things and so IANA should only be reactive. This sensitivity may no longer exist but we have not revisited it recently. Also, if we publish reports publicly that would be seen as naming and shaming even though there are third parties who do this reporting.

GH: Do you have plans for this?

KD: Pending availability of staff resources we would like to ask the community to review our technical checks – but this specific issue on proactively monitoring between change requests could also be considered, including what we should do if we see issues. We have ambitions to consult with the community on tech. checks but nothing scheduled.

GH: When I was analyzing the behaviour of the various RZ letters – do they publish the same thing? Does anyone look?

KD: Not doing this, I would love to have a way to confirm this. We are confirming that Verisign is publishing the correct thing. For the Root servers that is probably more of an RSSAC thing and we have no way of knowing all the nodes out there.

GH: RZ KSK roll. Not asking about the process – what if you cannot get to both keys or if you think they are compromised – do you have an emergency procedure?

KD: PTI and Verisign are discussing this for KSK and ZSK rolls. We have produced some documentation – so that work is still ongoing but is now the responsibility of the new security director. If both KSK facilities are obliterated, we would need to regenerate everything from scratch and we have no formal procedure for reconstructing the KMFs but is also fairly far fetched but maybe that is future work to plan for that scenario.

GH: Should we take this up with your new director of security?

KD: Could be an area in your recommendations for us to consider this.

GH: Top of the hour, this has been extremely valuable and will be the basis for my report. I have no plan for a meeting next week until I hear back from the co-chairs. I hope to have a draft of this report in about three weeks which will include draft recommendations and we should have a final report by the time we have a next face to face meeting.