Comments on the Root Zone TTL document and How they are addressed

7 July 2015

Overall feedback

Factual Errors in the Report

- Both Shinta and Daniel did not find any factual errors with the document.
    - Editors' Response: No changes needed.
- Bill Manning raised the issue in section 5.3 about Mark Koster's testimony that J root did not exist as a root server operator until 1996.
    - Editors' Response: acknowledge this error, removed J-root from section 5.3

Methodology of the Report

- Daniel M: Regarding the methodology, my understanding is that there is an impressive survey of the TTL values used that are related to the one specified in the root zone. In addition to this section I think it would be beneficial to have:
    - a description of the impact of the TTL on ICANN operations on the root zone. Such operations may be for example addition/removal of a new gTLD, key roll over, emergency key roll-over - of course, the topic should remain the TTLs.
        - Editors' Response: Some of these are covered in section 4. In general we note, the primary focus of the advisory is to consider the impact of changing TTLs to root server operators and recursive name servers, *not to ICANN*. [DONE]
    - a "theoretical" section that details what RFCs recommends for the various TTL settings. From this section I believe we should be able to be able to have a theoretical model, and describe theoretically how TTLs impact the ICANN operations on the root zone as well as the expected impacts on the traffic. I would also see that section 6.4 can fall into such a section.
        - Editors' Response: RFC 1033 has a section on TTLs, which we reference toward the end. Based on my familiarity with the RFCs I do not expect

that we'll find more any specific advice to DNS operators in general, nor for the root zone in particular.

- Shinta: I could not find specific problem statement throughout the document. Thus, the motivation of this study is not clear enough. It would be better if we can set up the problem with case examples or specific patterns.
    - Editors' Response: There is no urgent nor specific problem to be solved in this case. It is an "exploratory" project. The RSSAC statement of work lists the reasons the work party was formed. Section 5 of the document itemizes some reasons to consider changing TTLs.

Findings of the Report

- Daniel: I agree with the findings. However, saying that the value should not be changed because most client ignore it does not seems satisfactory. First I would say that the boundary seems to be the "max-cache-ttl" value, and having TTL values below it would probably have a major impact. In fact for most of the traffic, max-cache-ttl overwrites the TTL and becomes the de-facto TTL. I think we should elaborate a bit more about has the right TTL value (the root zone or the max-cache-ttl), and see if there are any recommendations to do for the max-cache-ttl value. I also understand it may be out of scope of the root zone TTL, but change of this value probably impact more the root zone traffic than the current TTL.
    - Editors' Response: Disagree with the comment. Different implementations have different default max-cache-ttl values, *and operators can change the defaults as well.* The root zone does not have a single TTL value, it has different values for different types of records. I don't believe our document should be giving advice to implementors or operators of name server software.
- Shinta: no objections.
    - Editors' Response: no actions needed.

Recommendations of the Report

- Daniel: I agree with the recommendations. However, I think we should also comment on the max-cache-ttl. As TTL is related to caching, I also think there should be some additional items on how these caches may be flushed by an authoritative party -- I am really thinking of the red button Joe Abley talked about in case of an emergency key-rollover, but it looks this could be useful in other situations. Maybe we should also

see some debugging indications when a badly cached value is stored -- again I am mostly thinking of outdated signatures RRsets cached.

    - Editors' Response: max-cache-ttl -- see above. cache flushing, while interesting, is out of scope.

- Shinta: no objections. If the issue is logically problematic, it should be fixed. But the other issues which has not been seen as the current operational problem do not need to be actively addressed. The recommendations match this way.

    - Editors' Response: No actions needed.

- Bill: I am in favor of the two recommendations.

    - Editors' Response: No actions needed.

Mitigation Options

- Shinta: The parameters which is not causing the operational problem directly should not be changed easily. If we shorten the Expire period, there may be negative operational impact in the case of communication trouble between root-zone distributor and root-servers. The parameter of the DNSSEC validity period is the place we can be changed without other impact.

    - Editors' Response: This is good feedback, thank you. Added the following to the document. "Neither the work party nor the RSSAC caucus have reached consensus recommendation on which mitigation option to take. Thus, this is an area to be further studied by the RSSAC. A few caucus members expressed the opinion that changing the SOA Expire or NS TTL values could have a negative operational impact and it would therefore be better to change the signature validity period."

Editorial Nits

1) introduction: The document uses the term "Internet Community": I am wondering whether this community is identified enough to named as such or if Internet community would be more appropriated.

Editors' Response: Change "Internet Community" to "DNS operations community"

2) section 3: Types of data in the DNS Root zone:

I think that some clarifications should be made in section 3 regarding the type of data. I found it quite confusing to read that most authoritative data are 1d TTL as most NS have 2d TTLs. This comes from the fact that TLD NS are not part of the authoritative data. I looked at dratf-ietf-dnsop-dns-terminology and checked that whether TLD NS are authoritative data or not is somehow confusing. So I would suggest to explicitly state it. The text I would propose could be something like:

"In this section we considered three types of DNS data.
- authoritative data : the data the root zone has authority for. In our case, we did not consider the TLD NS nor the TLD Glue RRsets (A, AAAA).
- delegation: The definition provided by the terminology section does not fit here as it is not a process.
- glue: (cf terminology section)
"

Editors' Response: I think table 2 quite clearly states which records are authoritative, which are delegations, and which are glue. and the terminology section already describes authoritiative, delegation, and glue.

Another alternative could also consists in clarifying the usage in the terminology section.

Editors' Response: Addressed above.

At last is there any need to have a specific terminology for these authoritative RRsets and have it mentioned in the dns-terminology draft ?

Editors' Response: IMO the terminology draft already covers this quite well. In any case, providing text for the terminology draft it out of scope for this work party.

3) section 3 "glue TTLs match their associated NS TTLs" I think we should specify whether this is something observed or if this follows some operational guidance.

Editors' Response: It is something observed. I don't really know if it is general operational practice. I'm pretty sure there is no advice or requirements that they should match.

4) section 3.1 "which means it has a negative caching TTL of 1 day". I think this may impact the publication of gTLD. as well as emergency key-roll over.

Editors' Response: Yes, this is explored in Section 4 as one possible reason to consider change TTLs. No changes are needed.

5) section 6.1.1

The graph below shows each TLD's authoritative NS TTL (y-axis) and its query count in the 2014 DITL data (x-axis). While there are certainly a large number of less-popular TLDs with large RTTs, by looking at the lower right section of the graph we can see that the more-popular TLDs tend to have larger RTTs (Round Trip Times). In other words, there are no popular TLDs with small RTTs. Overall, 90% of queries in 2014 DITL data were to TLDs having NS TTLs greater than or equal to 1 day.

I do not see the relation with RTT. Isn't misspelling of TTL?

Editors' Response: replace RTT with TTL. Thanks for catching this.

6) section 6.2 "The authoritative server returns the TXT record without a 10-day TTL." Shouldn't be "with" instead of without?

Editors' Response: s/without/with/.

7) section 6.3 "The study team did find 20% of the IPs made only one request for delegated TLDs during collection . ". Shouldn't it be "data collection period" instead?
idem for "Nearly 20% of the IPs made only request for delegated TLDs during collection".

Editors' Response: s/collection/the data collection period/.

8) Section 6.3
"IPs sending queries (red line) and providing measurements (blue line)." or the legend Queries/Measures in figure 4 does not make obvious the difference between the two lines. Maybe that would be better to clarify that explicitly.

Editors' Response: Added this to the "Study Results" paragraph before the graph:

An IP address must send at least 2 same-TLD queries in order to qualify as a measurement. Addresses which did not send at least 2 same-TLD queries are counted in the queries line, but not in the measurements line.

9) section 6.3 "Figure 6 and Table 8 below shows" Shouldn't it be "show"

Editors' Response: s/shows/show/.

10) section 6.4.1/ section 7 Findings 4: lots of missing white space after ".".

Editors' Response: corrected.

11) section 6.4.1 Is there any reason the relation between the different Time is not expressed.
SOA_Expire + NS_TTL <= ZSK_validity

Editors' Response: I'm not sure I understand the question. I think Figure 7 is clear.

12) section 6.4.2 "Just before day 10, a root server instance experiences a problem which prevents the it from receiving zone" Shouldn't we remove "the"?

Editor response: remove "the".

13) Findings 2:
"This likely means that root zone TTLs could be reduced to 1 day without any significant impact on traffic levels to root name servers"

Is that intentional to repeat the sentence?

Editors' Response: Removed the second occurrence of this sentence.

14) Findings 6:

"In Section 6.4, the study team identifies two situations whereby a root name server would return data and signatures that will be cached beyond the signature validity end time. Specifically, in certain situations: "

Do we have any kind of formal proof that there is no other corner cases.

Editors' Response: No we have no formal proofs

15) section 9:
Don't we need DNSSEC mitigation mechanisms to flush cache for example.?

Editors' Response: Cache flushing is out of scope

16. I didn't see the factual errors, but came up with some questions and editorial comments. These should be addressed.

- The expressions in the terminology part varies.
- reference to the RFCs exists or not

- way to refer the RFCs
- link to the URL exists or not

- Some abbreviation and words are used without enough descriptions.
ex) AA, TLD, Root Glue, TLD Glue, NS RRset TTL (whether glue of root(.) itself or TLD delegation?)

Editors' Response: edited the terminology section to make it more have similar reference. Please identify specific instances if not already identified above.

17. In the description of the SOA, nothing is mentioned for the REFRESH and RETRY. These can be described in combination with EXPIRE, as the parameters for the secondary name servers to determine the zone transfer timing.

Editors' Response: They have no effect on TTLs and caching behavior, but I am not opposed to adding them. The following text are added:

> "The SOA Refresh field specifies how often a secondary authoritative name server should poll the primary name server for zone updates. Note that secondary name servers may also learn of updates via NOTIFY messages before the next polling interval.
>
> The SOA Retry field specifies how long a secondary authoritative name server should wait to retry a failed zone update."

18. In the last paragraph of 6.1.1, the word RTT is used to explain the Figure 2. I could not understand this expression because the Figure 2 does not contain any information about RTT. The figure might be wrong.

Editors' Response: replaced "RTT" with "TTL", agreed and already addressed.

19. Whooo. Something I don't think we remembered to mention in the doc --perhaps we should toss in a mention?

http://www.mail-archive.com/bind-users@lists.isc.org/msg21023.html

Basically upwards referrals triggering fetches (and cache exhaustion, which we *might* have mentioned). Related to the topic, perhaps worth mentioning just for completeness?

Editors response: I'm having a hard time believing that an upward referral is or should trigger a fetch. It seems like standard cache poisoning protection that the recursive should ignore the out-of-baliwick upward referral.

20. What this document highlights is a highly unlikely situation: First, a root-server does not receive updates for the root zone. Second, this goes unnoticed for 7 days. I've not ever seen that situation… but lets go with it anyway. Better be safe:

The server would serve stale data for 7 days (expire) and 30 minutes (refresh), before it returns SERVFAIL. If this situation happens, a non-validating dnssec-aware resolver that resolves and caches a root NS RRSet plus RRSIG seconds before expire+refresh, will return a stale RRSIG to its validating stub resolvers 3 days (minus 30 minutes).

As similar problem arises for the KSK. What I'm missing from the document is a survey and discussion on "priming" queries. Most, if not all resolvers issue priming queries at some frequency. Priming queries are necessary to avoid the root-hints to become stale. These priming queries are ". NS" requests, exactly those that are discussed in the document as being problematic.

If this is a non-issue for NS records and we only have to consider the KSK case the set of options are weighted differently, as before, one option (reducing the NS TTL to lower than 3 days) would solve both issues, but if the NS issue is non existing……

Anyway, I'd ask RSSAC for input on the effect of priming queries.

Editors' response: In section 6.4.2 added text to highlight that this problem are very specific and uncommon. Highlights the instances it must go wrong for this problem to occur. Also added 6.4.3 section on the proof of concept simulation to show that this scenario could actually happen.

21. What need would a validating stub resolver have for the signature over the root NS records? I'd be interested in a use case as well (to show how things could go wrong).

Editors response: See section 6.4.3 to see how it would actually go wrong.

22. The Verio/NTT server is not for public use. The list is in Table 6 inside section 6.2.

Aside: I did try it and it is indeed open. But there's a blog post (in Japanese, using Google translate) that says it isn't for non-customers of the ISP. I can't speak to the authenticity nor authority of the blog, but, so maybe I'm wrong. Maybe it social engineered me.

<http://aimless.jp/blog/archives/2089> is the blog post URL.

Given at least the intent expressed, I'd remove it from the open recursive service list in the document. I also asked the person lodging the complaint to contact Verio/NTT and ask them to close it up if that is the intent.

Editors' response: Removed the Verio / NTT.