

RSSAC047v2 Work Party (WP) Call

Wednesday, 13 October 2021 at 16:00 UTC

Attendees

Abdulkarim Oloyede
Anand Raje
Anupam Agrawal
Barbara Schleckser
Brad Verd
Dan Mahoney
Dessalegn Yehuala
Duane Wessels
Karl Reuss
Kazunori Fujiwara
Ken Renard
Kevin Wright
Mark Kusters
Mauricio Vergara
Paul Hoffman
Ray Bellis
Russ Mundy
Ryan Stephenson
Shinta Sato
Warren Kumari

Staff: Steve Sheng, Ozan Sahin, Andrew McConachie

Agenda:

1. Welcome & Roll Call (Ozan Sahin)
2. Review [Decisions / Action Items](#) from Last Meeting (Steve Sheng)
NOTE: All action items were completed.
3. Content Discussion
 - a. Review [list of issues](#) to be discussed for RSSAC047v2 (Paul Hoffman)
 - i. Warren did an implementation and he identified a number of issues with the document. **DECISION: Warren will add his list of issues to the running list for WP discussion.**
 - ii. Warren raised if there are other purposes that would be served by the RSSAC047 data, and that given that we are spending efforts on the measurement system already. It seems to be a waste of effort to focus narrowly. The additional uses would be (1) help individual

root server operators to debug issues; (2) data be made available for research and monitor the root server system. The Working group had a discussion about it. It is concluded that:

1. **Decision: (1) The RSSAC047v2 work party is to focus on correcting technical inaccuracies. Additional purposes are out of the scope for the work party for version 2. (2) Warren's point about other uses could be dealt with by another work party.**

iii. Issues related to Elapsed time and timeouts

1. Paul mentioned that the current implementation uses "dig" which does not include connection setup time. He raised to discuss how important the work party thinks this issue is.
 - a. Warren said TCP time is important. It is important to measure from set up the tear down.
 - b. Warren raised the issue on how to handle a UDP query for the NS set for .com, but getting back a truncated response, at which time a TCP query will be issued. So the question is: **How to correct account for the latency of flipping over from UDP to TCP?**
 - c. Duane said that we do not have to worry about latency switching over to TCP requests. The primary purpose of latency measurement is to establish the round trip time, not to measure truncation.
 - d. Paul said 5.2 and 6.2 measures response latency of getting the first response, and not for chaining. He suggested that in section 4, probably add somewhere: "Response for TC bit set will be ignored, except for the correctness metrics."
 - e. Warren argues as an extreme if a root server set the MTU so low as compared to others, there will be noticeably different performance, but the current metrics cannot detect such as performance change.
 - f. Paul mentioned a more real world possibility would be a root server operator makes an operational choice to enable TCP by default, and thus all the responses will have the TC bit on. He said, this is an important question to discuss, but not on version 2.
 - g. Duane said the purpose of this document is not to measure the end user experience. The WP had a lot

of discussions sending queries through actual recursive operators to measure the full user experience, but decided against that in the end.

- b. Procedures for handling outstanding issues that are out of scope for v2 but potentially for future revisions.
 - i. The WP had a discussion on this topic. **Decision: (1) A new section on the current issues list will be created listing these issues not to be addressed in version 2, but might be useful to tackle in version 3, (2) The WP have a discussion of these parking lot items after working through v2, to scope out version 3 a bit, and give it back to RSSAC for their consideration.**
4. AOB
- a. Our next call is Thursday, 28 October 2021 16:00 UTC at ICANN72.