RSSAC047 Study

By SIDN Labs and NLnet Labs
In collaboration with Verisign and ISC

Agenda

- Background
- Motivation
- Goals
- Approach
- Results
  - Code review
  - Deployment
  - Availability metric and measurements
  - Independent measurements
  - Publication delay metric and measurements
- Recommendations
- Conclusions

RSSAC047

- Two versions: Current version RSSAC047v2 from 1 February, 2022
- Presents a set of metrics for the DNS root servers as well as for the RSS
  - **availability**, response latency, ~~correctness~~, **publication latency**
- Initial implementation
  - https://github.com/icann/root-metrics
  - implements the measurements defined in RSSAC047
  - runs and collects the measurements
  - generates reports

Motivation

- Reports from the initial implementation raised some eyebrows:
  1) On multiple occasions, the RSS did not meet the availability threshold of 99.999%
  2) In May 2024, C-Root had a high publication delay but the reports did not show anything unusual for the publication delay metric of C-Root

Goal

1) Evaluate whether the timeouts reported by the initial implementation were actually caused by sites of the RSIs
2) Evaluate why the reports by the initial implementation did not pick up on the publication delay at C-Root in May 2024

Approach

- Analyse whether the measurement results were influenced by the initial implementation:
  - by studying the source code
  - by studying the reported timeouts
- Analyse whether we can independently confirm the reported timeouts with RIPE Atlas measurements
- Run the initial implementation and analyse the reported measurements
- Analyse whether the metric to calculate the publication delay is sufficient to identify the publication delays

Results: Code review

Code review

- Initial implementation open source: https://github.com/icann/root-metrics
- Are there implementation specific aspects that impact reported results?
- Aspects:
  - Usability
    - Ran our own deployment of the initial implementation
  - Code organization
    - future proof design and the setup of the code bse
  - Code quality
    - craftsmanship, consistency in coding practice and general tidiness
  - Code readability
    - the mechanisms in place to assess and guarantee correctness
  - Security and privacy
- Assess expectations enumerated in the second recommendation of section 8 of RSSAC047v2 *(for transitioning from initial to official implementation)*

Code review - conclusions

- The initial implementation is a one-on-one thorough and of solid quality literal implementation of the measurements and metric calculations from RSSAC047
- We did not find any implementation specific aspects that would impact results
- We did find improvements for RSSAC047:
  - Coordinated removal of measurement results (to prevent disk space full)
  - Separate initial deployment from initial implementation
  - No hardcoded data (currently root server IP addresses)
  - **Publication Latency metric calculation *was* wrong**
  - **RSS response latency calculation *was* wrong**
  - **Unit tests for measurements and metric calculations**
- Not all expectations from RSSAC047 section 8 recommendation 2 are met, so the initial implementation cannot transition into the official implementation

Code review - some more observations

- RSSAC047 Section 8 Recommendations
  1. A list of expectations before initial implementation becomes official implementation
     - ❌ No equal distribution of vantage points
     - ❌ No reports are publicly available (implementation is unfinished)
     - ❌ List of vantage points is not publicly available
  2. Insight learned from the implementation will inform future revisions of the report
     - Already worked for version 2
     - We have feedback for version 3

Results: Initial implementation deployment

Deployment of the initial implementation

- 20 measurement vantage points (VPs)
  - not evenly distributed over regions (see RSSAC047v2 3.2)
- 1 collector
- Performs measurements every 5 minutes
- Coverage:
  - 253 out of 1,501 sites seen in January 2024 (16,86% coverage)
  - Coverage of individual RSIs varies between 7.54% and 100%
  - 69% of all sites are only reached by one VP at once
- Traceroute measurements:
  - One traceroute measurement towards each IP per VP and measurement interval
  - Between 51.25% (IPv6) and 61.14% (IPv4) successful
  - 9 hops between VP and site (median)

Results: Availability metrics

Availability measurements: Failed tests

Month	IP Version	Transport	Availability
2023 July	4	UDP	99.997%
2023 August	4	UDP	99.992%
	6	UDP	99.992%
2023 September	4	UDP	99.997%
	6	UDP	99,996%
2023 October	4	UDP	99.993%
	6	UDP	99,987%
2023 November	4	UDP	99.986%
	6	UDP	99,984%
2023 December	4	UDP	99,996%
	6	UDP	99,996%
2024 January	6	UDP	99.996%
2024 March	6	UDP	99.844%

Months in which the RSS did not meet the 99.999% availability threshold

Availability measurements: Timeout characteristics

- 84.65% of timeouts lasted for one measurement interval (max 5 minutes)

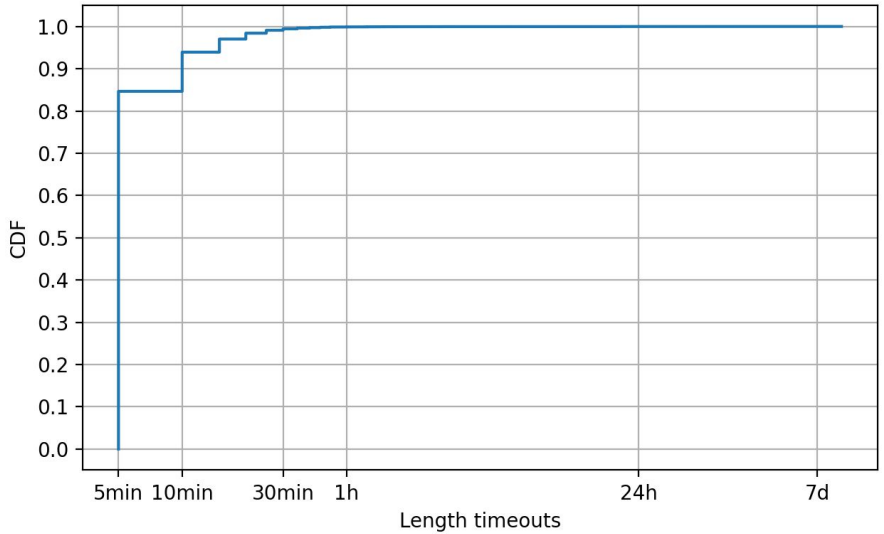

Availability measurements: per RSI

- Timeouts not evenly distributed across RSIs
- Timeouts not evenly distributed across sites of RSIs

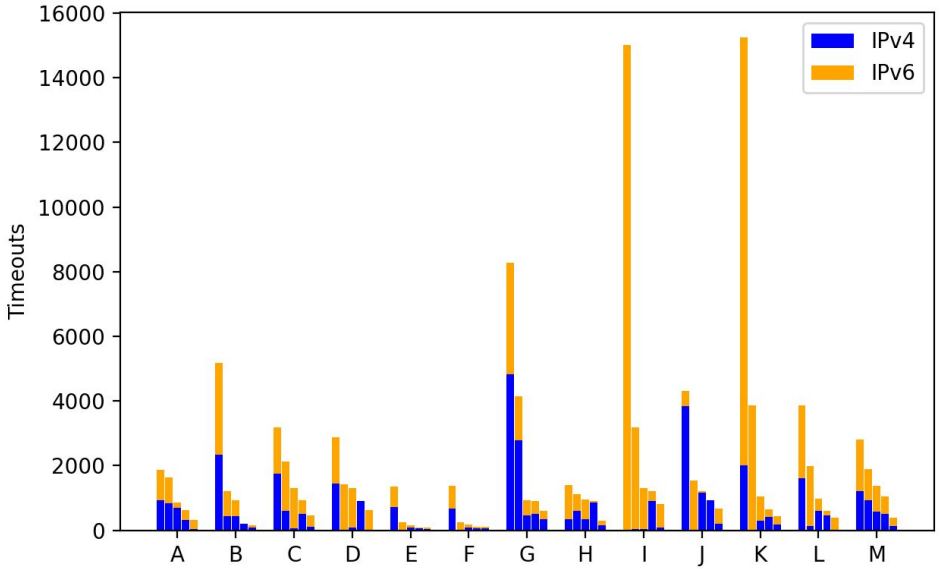

Timeouts per site of the RSI (top 5 sites only)

Availability measurements: per VP

- Timeouts not evenly distributed across VPs

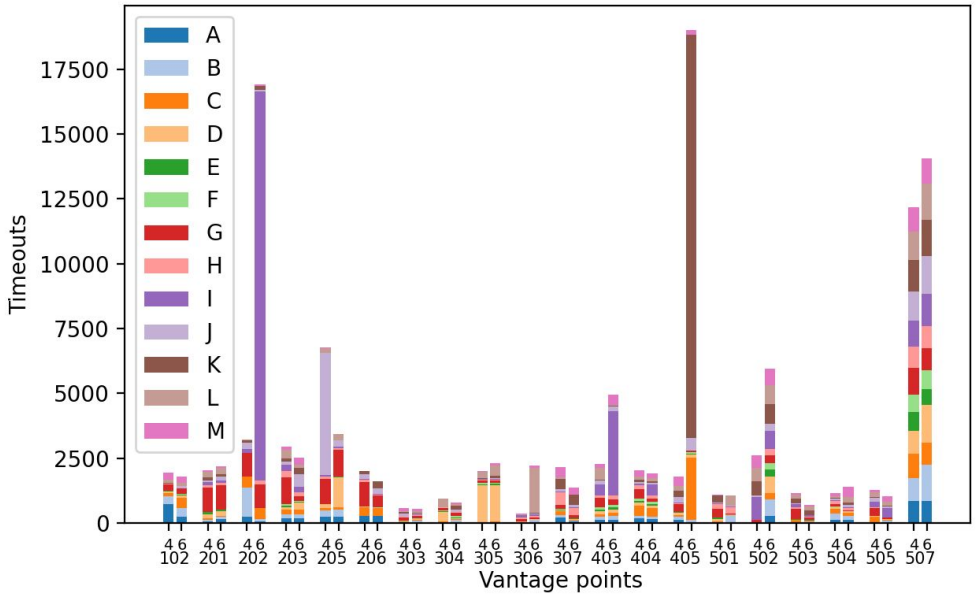

Timeouts per VP, RSI and IP version

Availability measurements: per VP

- Timeouts not evenly distributed across VPs

Some VPs see almost the same number of timeouts across all RSIs

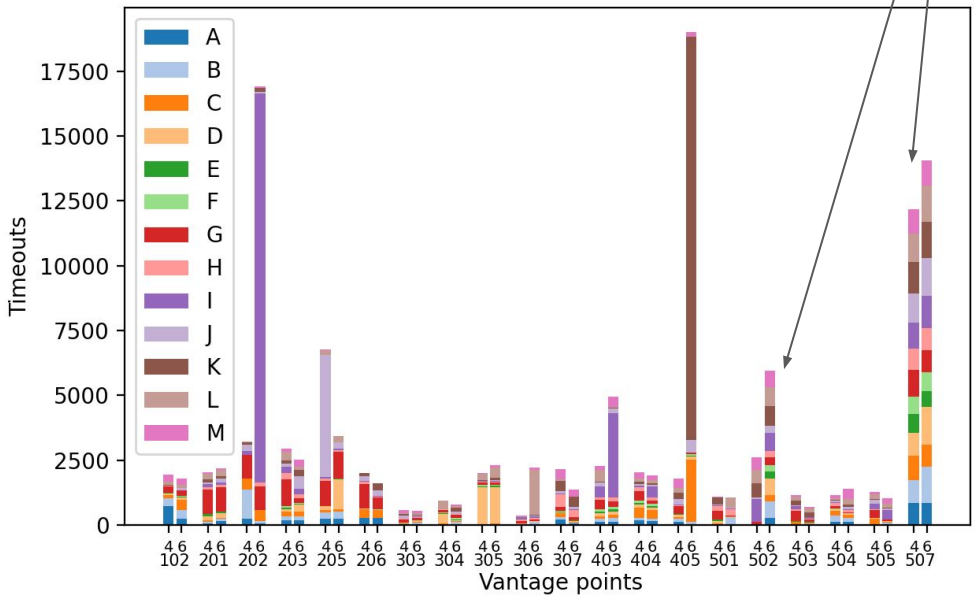

Timeouts per VP, RSI and IP version

Availability measurements: per VP

- 98.35% of all timeouts were only observed by one single vantage point
- Concurrent timeouts:
  - 26.16% same VP, same measurement interval, same IP version (thus different RSI)
  - 37.32% same VP, same measurement interval
  - 3 VPs have concurrent timeouts more than 50% of the time they observed a timeout

Availability measurements: per VP

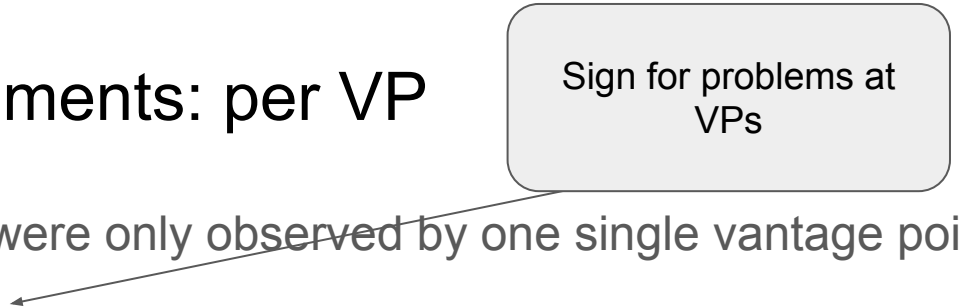

- 98.35% of all timeouts were only observed by one single vantage point
- Concurrent timeouts:
  - 26.16% same VP, same measurement interval, same IP version (thus different RSI)
  - 37.32% same VP, same measurement interval
  - 3 VPs have concurrent timeouts more than 50% of the time they observed a timeout

Availability measurements: per VP

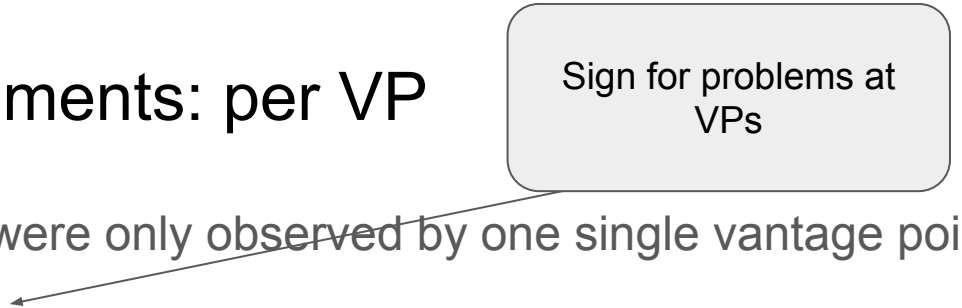

- 98.35% of all timeouts were only observed by one single vantage point
- Concurrent timeouts:
  - 26.16% same VP, same measurement interval, same IP version (thus different RSI)
  - 37.32% same VP, same measurement interval
  - 3 VPs have concurrent timeouts more than 50% of the time they observed a timeout
- Routing loops:
  - During 3.39% of the timeouts the traceroute terminated because of a loop

Availability measurements: per VP

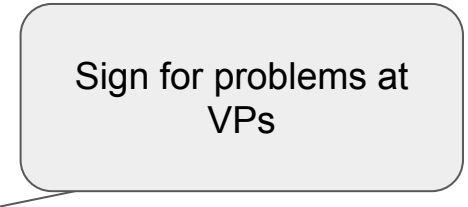

- 98.35% of all timeouts were only observed by one single vantage point
- Concurrent timeouts:
  - 26.16% same VP, same measurement interval, same IP version (thus different RSI)
  - 37.32% same VP, same measurement interval
  - 3 VPs have concurrent timeouts more than 50% of the time they observed a timeout
- Routing loops:
  - During 3.39% of the timeouts the traceroute terminated because of a loop

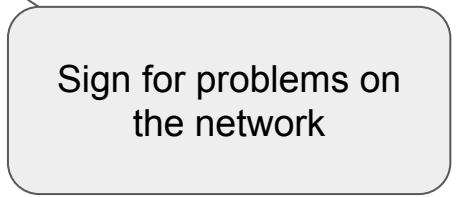

Availability measurements: Correlating with RIPE Atlas

- Motivation:
  - impossible to determine the cause of a timeout for certain after the fact
  - use RIPE Atlas measurements to understand if also others have observed the timeout
  - No: likely problem at the VP
  - Yes: likely problem at the RSI
- Between 2023-11-09 and 2024-04-03
- Filter probes to improve reliability of the measurements
- RSI site coverage: 65.76%

Availability measurements: Correlating with RIPE Atlas

- Measuring availability:
  - Number of probes that reached a site per 5 minutes (X)
  - Identify decrease in probes reaching a site by:
    - Calculating mean 12 hours before and after the timeout (μ)
    - Calculating the standard deviation (σ)
    - Significant decrease in reachability, when:
      - $X < \mu - 3\sigma$ (conservative approach)
      - $X < \mu - 2\sigma$ (liberal approach)
- Decrease reachability at Atlas correlates with timeout when:
  - decrease appears max 5 min before or 5 min after the timeout

Availability measurements: Correlating with RIPE Atlas

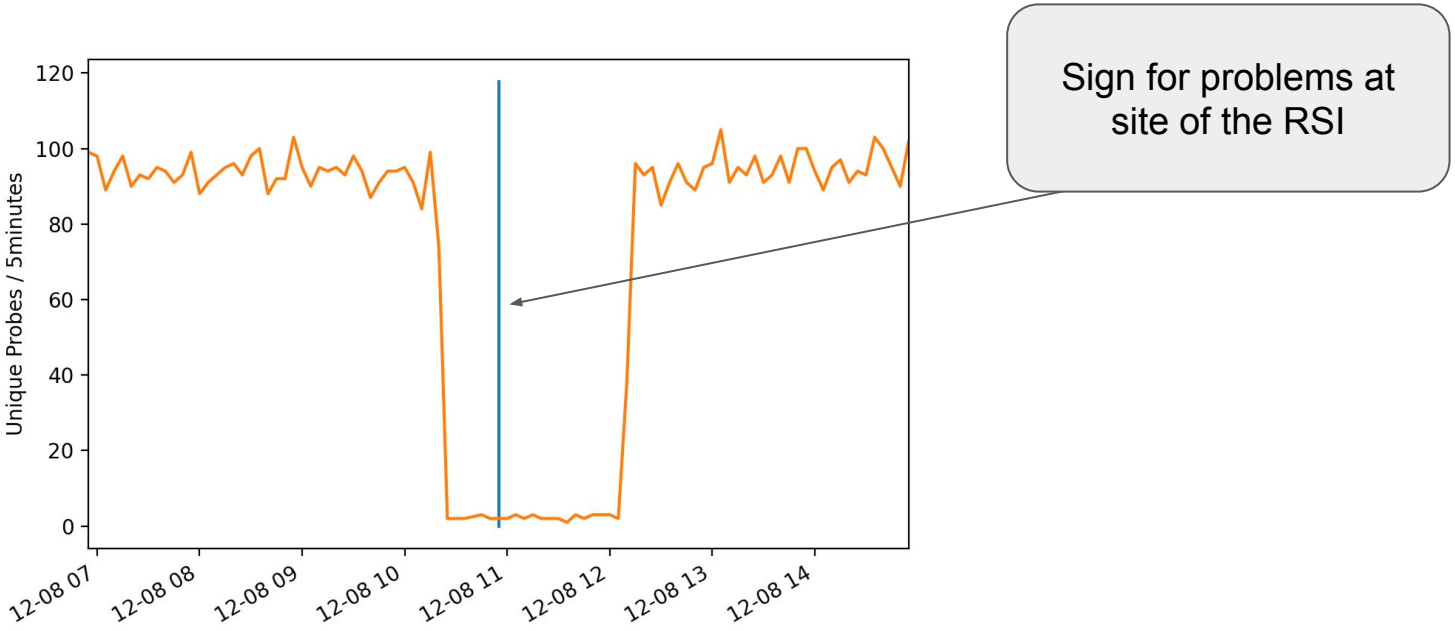

Example of a correlated timeout with RIPE Atlas

Availability measurements: Correlating with RIPE Atlas

- For 87.48% of the timeouts we could find corresponding Atlas measurements
  - For these timeouts, we could identify a decrease in reachability in
    - 2.10% of the cases (conservative threshold)
    - 11.97% of the cases (liberal threshold)

Availability measurements: Classify timeouts

- 3 classes
  - VP problem
  - Network problem
  - Site problem
- Differentiate between *likely* the root cause and *very likely* the root cause

Availability measurements: Classify timeouts

- VP problem
  - *Very likely* the root cause, when:
    - 1) The timeout **has not been observed** by another vantage point at the same time, and
    - 2) the vantage point **has observed** another timeout at the same time, and
    - 3) the timeout **does not coincide** with a traceroute measurement that resulted in a loop, and
    - 4) the timeout **could not be correlated** with a drop in reachability with low confidence.
  - *Likely* the root cause, when:
    - all the above criteria hold, expect 2)

Availability measurements: Classify timeouts

- RSI problem
  - *Very likely* the root cause, when:
    - 1) The vantage point **has not observed** another timeout at the same time, and
    - 2) The timeout **does not coincide** with a traceroute measurement that resulted in a loop, and
    - 3) The timeout **could be correlated** with a drop in reachability **with high confidence**.
  - *Likely* the root cause, when:
    - 1) The vantage point **has not observed** another timeout at the same time, and
    - 2) The timeout **does not coincide** with a traceroute measurement that resulted in a loop, and
    - 3) The timeout **could be correlated** with a drop in reachability **with low confidence**.

Availability measurements: Classify timeouts

- Network problem
  - *Likely* the root cause, when:
    - 1) There **exists** a traceroute measurement towards a site of an RSI that resulted in a routing loop at the time of the timeout, and
    - 2) there **does not exist** a traceroute measurement towards a site of an RSI that resulted in a routing loop one measurement interval before the timeout, and
    - 3) the timeout **could not be** correlated with a drop in reachability with low confidence.

Availability measurements: Classify timeouts

Problem location	Likely	Very likely
Measurement vantage point	81.3%	19.1%
Network path	4.3%	-
RSI	9.6%	1.4%
Unknown	4.8%	79.5%

Availability measurements: Classify timeouts

Problem location	Likely	Very likely
Measurement vantage point	81.3%	19.1%
Network path	4.3%	-
RSI	9.6%	1.4%
Unknown	4.8%	79.5%

- 7 vantage points with > 20% than timeouts classified as VP problems

Availability measurements: Adjusted root server metrics

- Calculating RSS availability, adjusted by classified timeouts:

		Availability		
Month	**IP version**	**Default**	**Low Confidence**	**High Confidence**
2023-12	**v4**	99.995848	100.000000	100.000000
	v6	99.995920	100.000000	100.000000
2024-01	**v4**	99.999160	100.000000	100.000000
	v6	99.995519	99.999650	99.999650
2024-02	**v4**	99.999467	100.000000	100.000000
	v6	99.999163	99.999848	99.999848
2024-03	**v4**	99.999638	100.000000	100.000000
	v6	99.844015	99.998117	99.998117

Availability measurements: Adjusted root server metrics

- Calculating RSS availability, adjusted by classified timeouts:
  - RSS meets threshold 7 out of 8 times
  - 2024-03 v6 likely caused by missing RIPE Atlas measurements

		Availability		
Month	**IP version**	**Default**	**Low Confidence**	**High Confidence**
2023-12	**v4**	99.995848	100.000000	100.000000
	v6	99.995920	100.000000	100.000000
2024-01	**v4**	99.999160	100.000000	100.000000
	v6	99.995519	99.999650	99.999650
2024-02	**v4**	99.999467	100.000000	100.000000
	v6	99.999163	99.999848	99.999848
2024-03	**v4**	99.999638	100.000000	100.000000
	v6	99.844015	99.998117	99.998117

Availability measurements: Takeaways

- Between 19.1% and 81.3% of timeouts are caused by VP or the network
- Between 1.4% and 9.6% of all timeouts are caused by the RSI

→ RSS meet the availability threshold most of the time

Results: Independent measurements

Independent measurements

- Goals:
  - Test initial implementation on a larger scale
  - Test possible extensions
- Setup:
  - 17 vantage points
  - 1 month test period
- 2 extensions:
  - Non-DNS related measurements
  - Additional ICMP based monitoring of VPs

Independent measurement results

- Deployment on multiple vantage points easy
- Availability on par with or lower than reported by initial implementations
- 3 vantage points reported 48.2% of the timeouts
- 54.5% of the timeouts on IPv4 correlated with failed traceroute measurements to public DNS services of Google and Cloudflare
- 41.0% of the timeouts on IPv4 and 9.4% of the timeouts on IPv6 correlated with ICMP based timeouts

Independent measurement conclusions

- At least 20.85% of the timeouts observed by our deployment are not caused by the servers of the RSS
- Additional measurements can give some additional insights

Results: Publication delay metrics

Publication Delay: Measurements by the initial implementation

- In May 2024, 74 out of 85 zones published by C-Root were on time (not more than 65 minutes late)
- 3 zones late for more than 65 minutes
- 8 zones not observed at all

Publication Delay: Measurements by the initial implementation

- In May 2024, 74 out of 85 zones published by C-Root were on time (not more than 65 minutes late)
- 3 zones late for more than 65 minutes
- 8 zones not observed at all

→ Initial implementation **has picked up** on the delayed zone updates

Publication Delay: Calculating the RSI publication delay

- In a nutshell:
  - For each zone and RSI, measure the time between zone publication and the time the zone has been fully deployed
  - Calculate the median across all values
- The initial implementation calculated the metric in accordance with RSSAC047v2*

  → The metric is not sufficient to pick up even large delays in zone publication

* The initial implementation contained a bug, but this did not affect the results in this case

Publication Delay: Flaws in the RSI metric

- Does not take zones into account that are never published
- Median is not sensitive enough to have picked up delays at C-Root
- Suggestions:
  - Take missing zones into account, e.g. by measuring time between missing zone was first published and the time the next zone was seen first
  - Use the mean instead of the median

Publication Delay: Modified RSI metric

RSI	Measurements	Median publication delay (s)	Adjusted publication delay (s)
A	85	0.0	21.2
B	85	0.0	17.6
C	85	0.0	11,523.5
D	85	0.0	17.7
E	85	0.0	28.2
F	85	0.0	24.7
G	85	0.0	14.1
H	85	0.0	14.1
I	85	0.0	17.6
J	85	0.0	3.5
K	85	0.0	17.6
L	85	0.0	35.3
M	85	0.0	17.6

Publication Delay: Modified RSI metric

→ RSS publication delay meets threshold (35 minutes) in both cases

RSI	Measurements	Median publication delay (s)	Adjusted publication delay (s)
A	85	0.0	21.2
B	85	0.0	17.6
C	85	0.0	11,523.5
D	85	0.0	17.7
E	85	0.0	28.2
F	85	0.0	24.7
G	85	0.0	14.1
H	85	0.0	14.1
I	85	0.0	17.6
J	85	0.0	3.5
K	85	0.0	17.6
L	85	0.0	35.3
M	85	0.0	17.6

Recommendations and conclusion

Recommendations

- Additional monitoring of VPs
  - external: monitoring, e.g. from the collector
  - internal: DNS independent connectivity tests
- More VPs
  - timeout observed from multiple VPs generate stronger signal
  - Open question: How many VPs is enough?

Conclusions

Put bluntly:

The initial implementation reported something that didn’t happen and didn’t report something that did happen.

- Availability of the RSS very likely higher than reported
- Metric to calculate the publication delay not sensitive enough
- Measuring the availability of a highly distributed system like the RSS externally is challenging
- More VPs can increase the reliability of the measurements

Questions? Comments?

Contact:

Moritz Müller (SIDN)	moritz.muller@sidn.nl
Marco Davids (SIDN)	marco.davids@sidn.nl
Willem Toorop (NLnet Labs)	willem@nlnetlabs.nl