

# Anycast Monitoring with passive TCP RTT

## (abstract)

---

**Giovane C. M. Moura** (1)   John Heidemann (2)   Jeroen Bulten (3)  
Wes Hardaker (2)   Joao Ceron (1)   Cristian Hesselman (1,4)   Maarten  
Wullink (1)   Marc Groeneweg (3)   Marco Davids (1)

**DINR2020 Worskhop**

2020-07-22

Virtual Conference

1: SIDN Labs

2: USC/ISI

3: SIDN

4: University of Twente

- Authoritative DNS operators strive to reduce latency
- They often deploy IP anycast on dozens of global sites
  - so they can be closer to clients
- Still, operators *don't* really know what latency most of their clients experience

# Measuring your client's latency nowadays

## Active methods

- Ripe Atlas [4, 5]
- ThousandEyes [8]
- Verfploeter [1]

They have pros and cons

- Either few VPs, or require some extra effort

## Passive methods

- DNS TCP RTT - this presentation
- Measure latency for *real clients*
- Works with IPv6, and it's *easy*
- It's being used in production at SIDN (.nl ccTLD)
  - Tech report available [3]:  
<https://www.isi.edu/~ejohnh/PAPERS/Moura20a.pdf>

# Measuring your client's latency nowadays

## Active methods

- Ripe Atlas [4, 5]
- ThousandEyes [8]
- Verfploeter [1]

They have pros and cons

- Either few VPs, or  
require some extra effort

## Passive methods

- **DNS TCP RTT** - this presentation
- Measure latency for *real clients*
- Works with IPv6, and it's *easy*
- It's being used in production at SIDN (.nl ccTLD)
  - Tech report available [3]:  
<https://www.isi.edu/%7ejohnh/PAPERS/Moura20a.pdf>

# DNS over TCP? Really?


- Passive TCP latency has been used since 1996 [2]
- Used in CDNs [6]
- But nobody thought of using it for DNS
- Why?
  - DNS over TCP is often overlooked
  - It shouldn't

## DNS over TCP? Really?

See tech report for details [3]. But in short:

- 2.6–6% of .nl traffic is over TCP
- 20+% of resolvers *eventually* send TCP queries
  - but they are responsible for 30–57% of **ALL** queries
- 44% of ASes covered
- and it's *free* – passive.

## Architecture used at SIDN



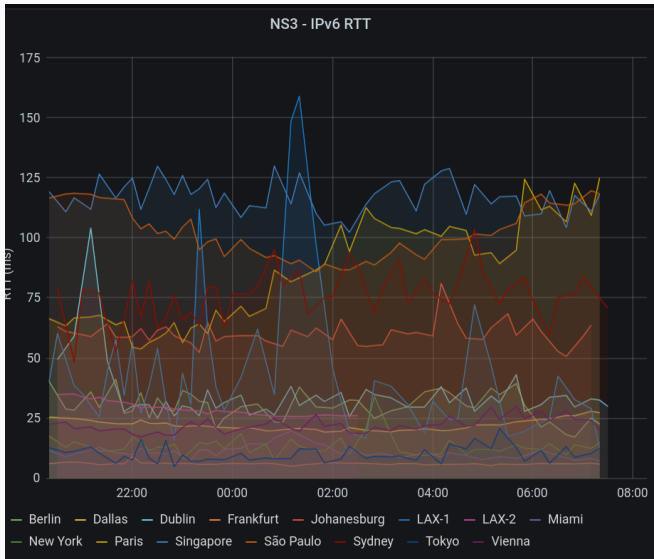
**Figure 1:** DNS TCP RTT Monitoring Architecture

## What can we do with that?

1. Real time monitoring
2. Evaluate your most important client ASes
3. Evaluate individual anycast sites
4. Detect near-real time BGP misconfiguration



# Real time monitoring



# Prioritization by Client AS




Figure 2: NS1 IPv4




Figure 3: NS1 IPv4

# Fixing latency for Google





Figure 4: Google: AS15169 - IPv4 - RTT (ms)

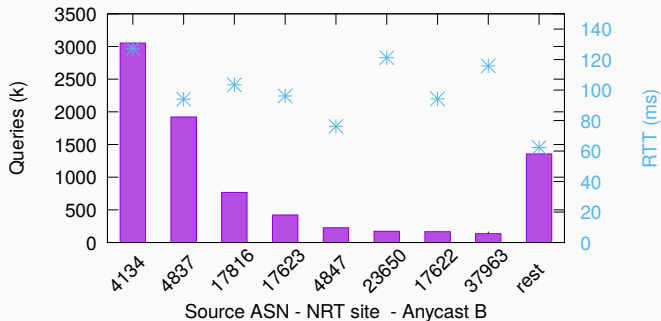
# Anycast Site analysis



**Figure 5: NS3 IPv4**

- See NRT (Tokyo, Narita)? High inter-quartile values . But why?

# Anycast Site analysis



**Figure 6: NS3 IPv4**

- NRT clients: top 10 from China
- And China has poor Intl' peering [9]

# Detecting Near real-time BGP misconfiguration





Figure 7: Anycast B SYD site: Latency for IPv4.



## Summary

- TCP RTT is old but *gold*
- We use it on `.nl`
- Provides insight into real clients
- Which can be used to fix real problems
- ENTRADA does the math for you – open-source

- [1] DE VRIES, W. B., DE O. SCHMIDT, R., HARAKER, W., HEIDEMANN, J., DE BOER, P.-T., AND PRAS, A.

**Verfloeter: Broad and load-aware anycast mapping.**

In *Proceedings of the ACM Internet Measurement Conference* (London, UK, 2017).

- [2] HOE, J. C.

**Improving the start-up behavior of a congestion control scheme for tcp.**

In *Proceedings of the ACM SIGCOMM Conference* (Stanford, CA, Aug. 1996), ACM, pp. 270–280.



- [3] MOURA, G. C. M., HEIDEMANN, J., HARDAKER, W., BULTEN, J., CERON, J., AND HESSELMAN, C.

**Old but gold: Prospecting TCP to engineer DNS anycast (extended).**

Tech. Rep. ISI-TR-740, USC/Information Sciences Institute, June 2020.

- [4] RIPE NCC STAFF.

**RIPE Atlas: A Global Internet Measurement Network.**

*Internet Protocol Journal (IPJ)* 18, 3 (Sep 2015), 2–26.

[5] RIPE NETWORK COORDINATION CENTRE.

**RIPE Atlas.**

<https://atlas.ripe.net>, 2015.

[6] SCHLINKER, B., CUNHA, I., CHIU, Y.-C., SUNDARESAN, S.,  
AND KATZ-BASSETT, E.

**Internet Performance from Facebook's Edge.**

In *Proceedings of the Internet Measurement Conference* (New York, NY, USA, 2019), IMC '19, ACM, pp. 179–194.

[7] SIDN LABS.

**ENTRADA - DNS Big Data Analytics, Jan. 2020.**

<https://entrada.sidnlabs.nl/>.

[8] THOUSANDEYES.

**Digital Experience Monitoring.**

<https://www.thousandeyes.com/>, June 2020.

[9] ZHU, P., MAN, K., WANG, Z., QIAN, Z., ENSAFI, R.,  
HALDERMAN, J. A., AND DUAN, H.

**Characterizing Transnational Internet Performance and the  
Great Bottleneck of China.**

In *Proceedings of the ACM SIGMETRICS conference* (Boston, MA, USA, June 2020), ACM, p. (to appear).