SIDN

Your world. Our domain.

Serverless DNS Analytics using ENTRADA 2

Maarten Wullink | ICANN66
Montreal, 4 Nov 2019

SIDN LABS

Analyzing DNS data

1. Capture DNS data using Tcpdump
2. 1 file per 5-10 minutes
3. Analyze using Wireshark/Tshark

Done?

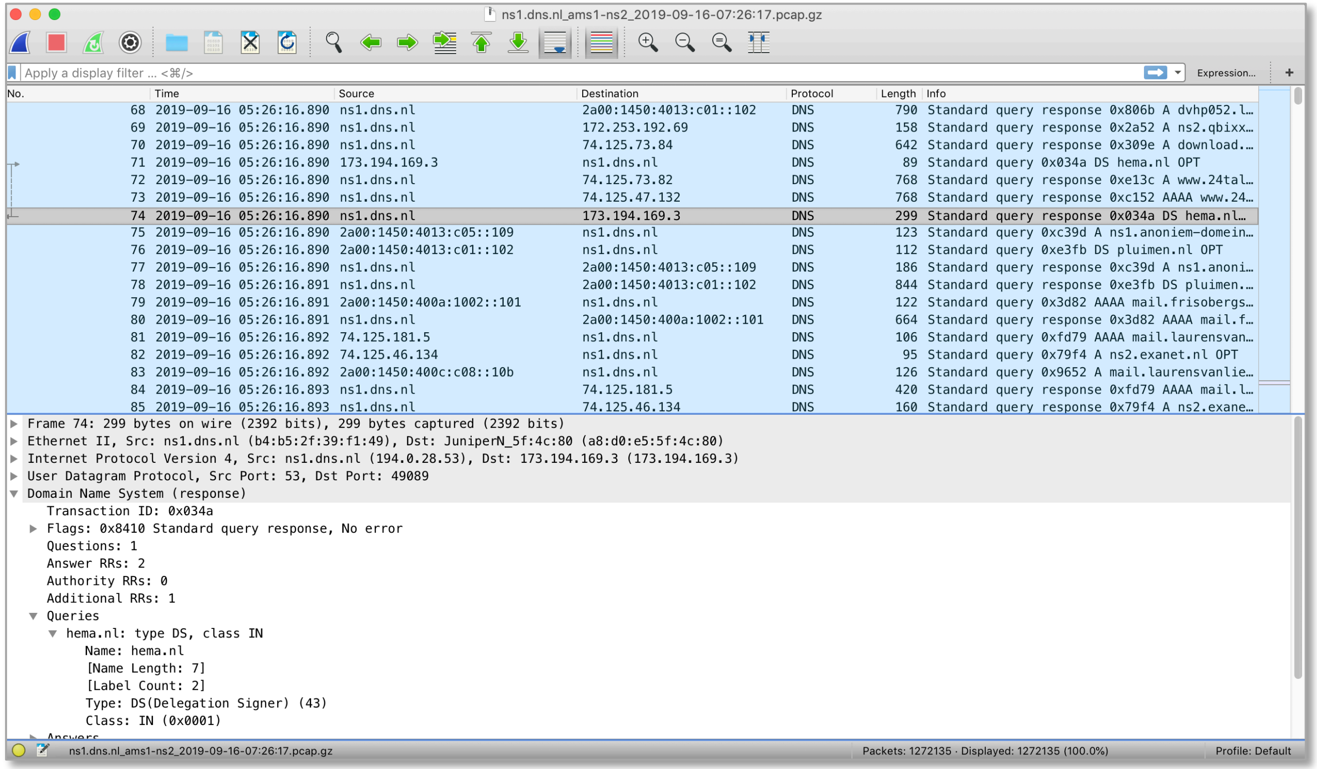
ns1.dns.nl_ams1-ns2_2019-09-16-07:26:17.pcap.gz
Apply a display filter ... <⌘/>
Expression...
+

```
No.   Time                     Source                     Destination                Protocol  Length  Info
   68 2019-09-16 05:26:16.890  ns1.dns.nl                 2a00:1450:4013:c01::102    DNS          790  Standard query response 0x806b A dvhp052.l…
   69 2019-09-16 05:26:16.890  ns1.dns.nl                 172.253.192.69             DNS          158  Standard query response 0x2a52 A ns2.qbixx…
   70 2019-09-16 05:26:16.890  ns1.dns.nl                 74.125.73.84               DNS          642  Standard query response 0x309e A download.…
   71 2019-09-16 05:26:16.890  173.194.169.3              ns1.dns.nl                 DNS           89  Standard query 0x034a DS hema.nl OPT
   72 2019-09-16 05:26:16.890  ns1.dns.nl                 74.125.73.82               DNS          768  Standard query response 0xe13c A www.24tal…
   73 2019-09-16 05:26:16.890  ns1.dns.nl                 74.125.47.132              DNS          768  Standard query response 0xc152 AAAA www.24…
   74 2019-09-16 05:26:16.890  ns1.dns.nl                 173.194.169.3              DNS          299  Standard query response 0x034a DS hema.nl…
   75 2019-09-16 05:26:16.890  2a00:1450:4013:c05::109    ns1.dns.nl                 DNS          123  Standard query 0xc39d A ns1.anoniem-domein…
   76 2019-09-16 05:26:16.890  2a00:1450:4013:c01::102    ns1.dns.nl                 DNS          112  Standard query 0xe3fb DS pluimen.nl OPT
   77 2019-09-16 05:26:16.890  ns1.dns.nl                 2a00:1450:4013:c05::109    DNS          186  Standard query response 0xc39d A ns1.anoni…
   78 2019-09-16 05:26:16.891  ns1.dns.nl                 2a00:1450:4013:c01::102    DNS          844  Standard query response 0xe3fb DS pluimen.…
   79 2019-09-16 05:26:16.891  2a00:1450:400a:1002::101   ns1.dns.nl                 DNS          122  Standard query 0x3d82 AAAA mail.frisobergs…
   80 2019-09-16 05:26:16.891  ns1.dns.nl                 2a00:1450:400a:1002::101   DNS          664  Standard query response 0x3d82 AAAA mail.f…
   81 2019-09-16 05:26:16.892  74.125.181.5               ns1.dns.nl                 DNS          106  Standard query 0xfd79 AAAA mail.laurensvan…
   82 2019-09-16 05:26:16.892  74.125.46.134              ns1.dns.nl                 DNS           95  Standard query 0x79f4 A ns2.exanet.nl OPT
   83 2019-09-16 05:26:16.892  2a00:1450:400c:c08::10b    ns1.dns.nl                 DNS          126  Standard query 0x9652 A mail.laurensvanlie…
   84 2019-09-16 05:26:16.893  ns1.dns.nl                 74.125.181.5               DNS          420  Standard query response 0xfd79 AAAA mail.l…
   85 2019-09-16 05:26:16.893  ns1.dns.nl                 74.125.46.134              DNS          160  Standard query response 0x79f4 A ns2.exane…
```

```
▶ Frame 74: 299 bytes on wire (2392 bits), 299 bytes captured (2392 bits)
▶ Ethernet II, Src: ns1.dns.nl (b4:b5:2f:39:f1:49), Dst: JuniperN_5f:4c:80 (a8:d0:e5:5f:4c:80)
▶ Internet Protocol Version 4, Src: ns1.dns.nl (194.0.28.53), Dst: 173.194.169.3 (173.194.169.3)
▶ User Datagram Protocol, Src Port: 53, Dst Port: 49089
▼ Domain Name System (response)
    Transaction ID: 0x034a
  ▶ Flags: 0x8410 Standard query response, No error
    Questions: 1
    Answer RRs: 2
    Authority RRs: 0
    Additional RRs: 1
  ▼ Queries
    ▼ hema.nl: type DS, class IN
        Name: hema.nl
        [Name Length: 7]
        [Label Count: 2]
        Type: DS(Delegation Signer) (43)
        Class: IN (0x0001)
  ▶ Answers
```

```
ns1.dns.nl_ams1-ns2_2019-09-16-07:26:17.pcap.gz        Packets: 1272135 · Displayed: 1272135 (100.0%)        Profile: Default
```

Wireshark example

ENTRADA

ENhanced **T**op-Level Domain **R**esilience through **A**dvanced **D**ata **A**nalysis

- **What we want:**
  - Good performance
  - High availability
  - Semi real-time data warehouse
  - **SQL support**

ENTRADA

Features

- Convert DNS data from PCAP to Parquet format
- Data Enrichment & optimization
- Automated workflow

Apache Parquet

- Columnar storage format
  - Developed by Twitter & Cloudera

- Why Parquet?
  - Smaller files and more efficient to read (less disk IO)

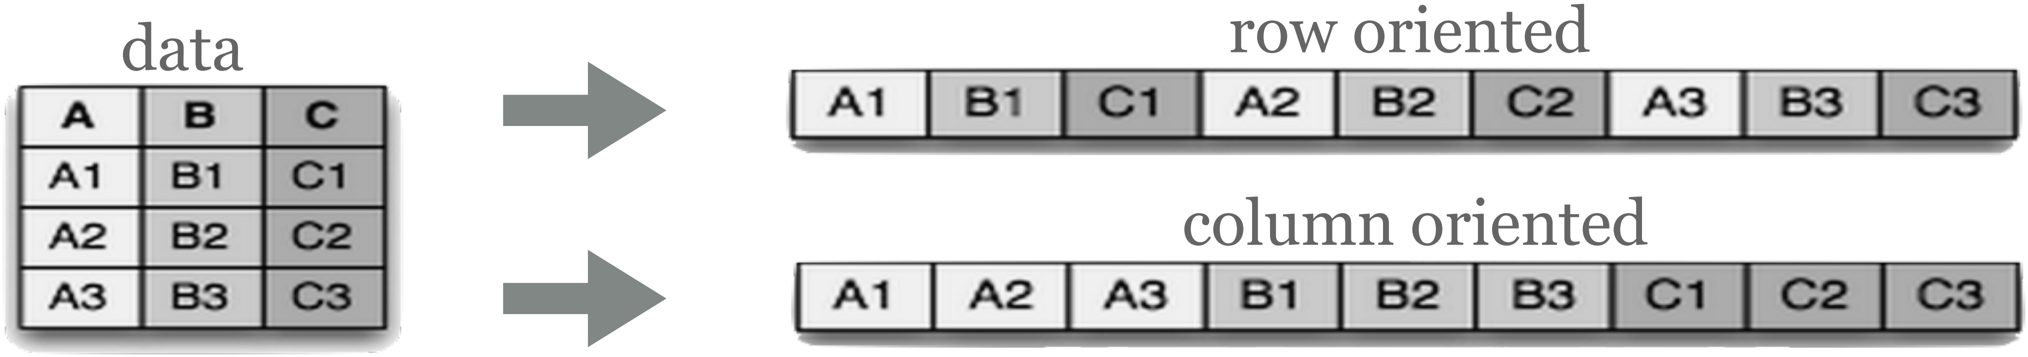

data

A	B	C
A1	B1	C1
A2	B2	C2
A3	B3	C3

row oriented

A1	B1	C1	A2	B2	C2	A3	B3	C3

column oriented

A1	A2	A3	B1	B2	B3	C1	C2	C3

Workflow

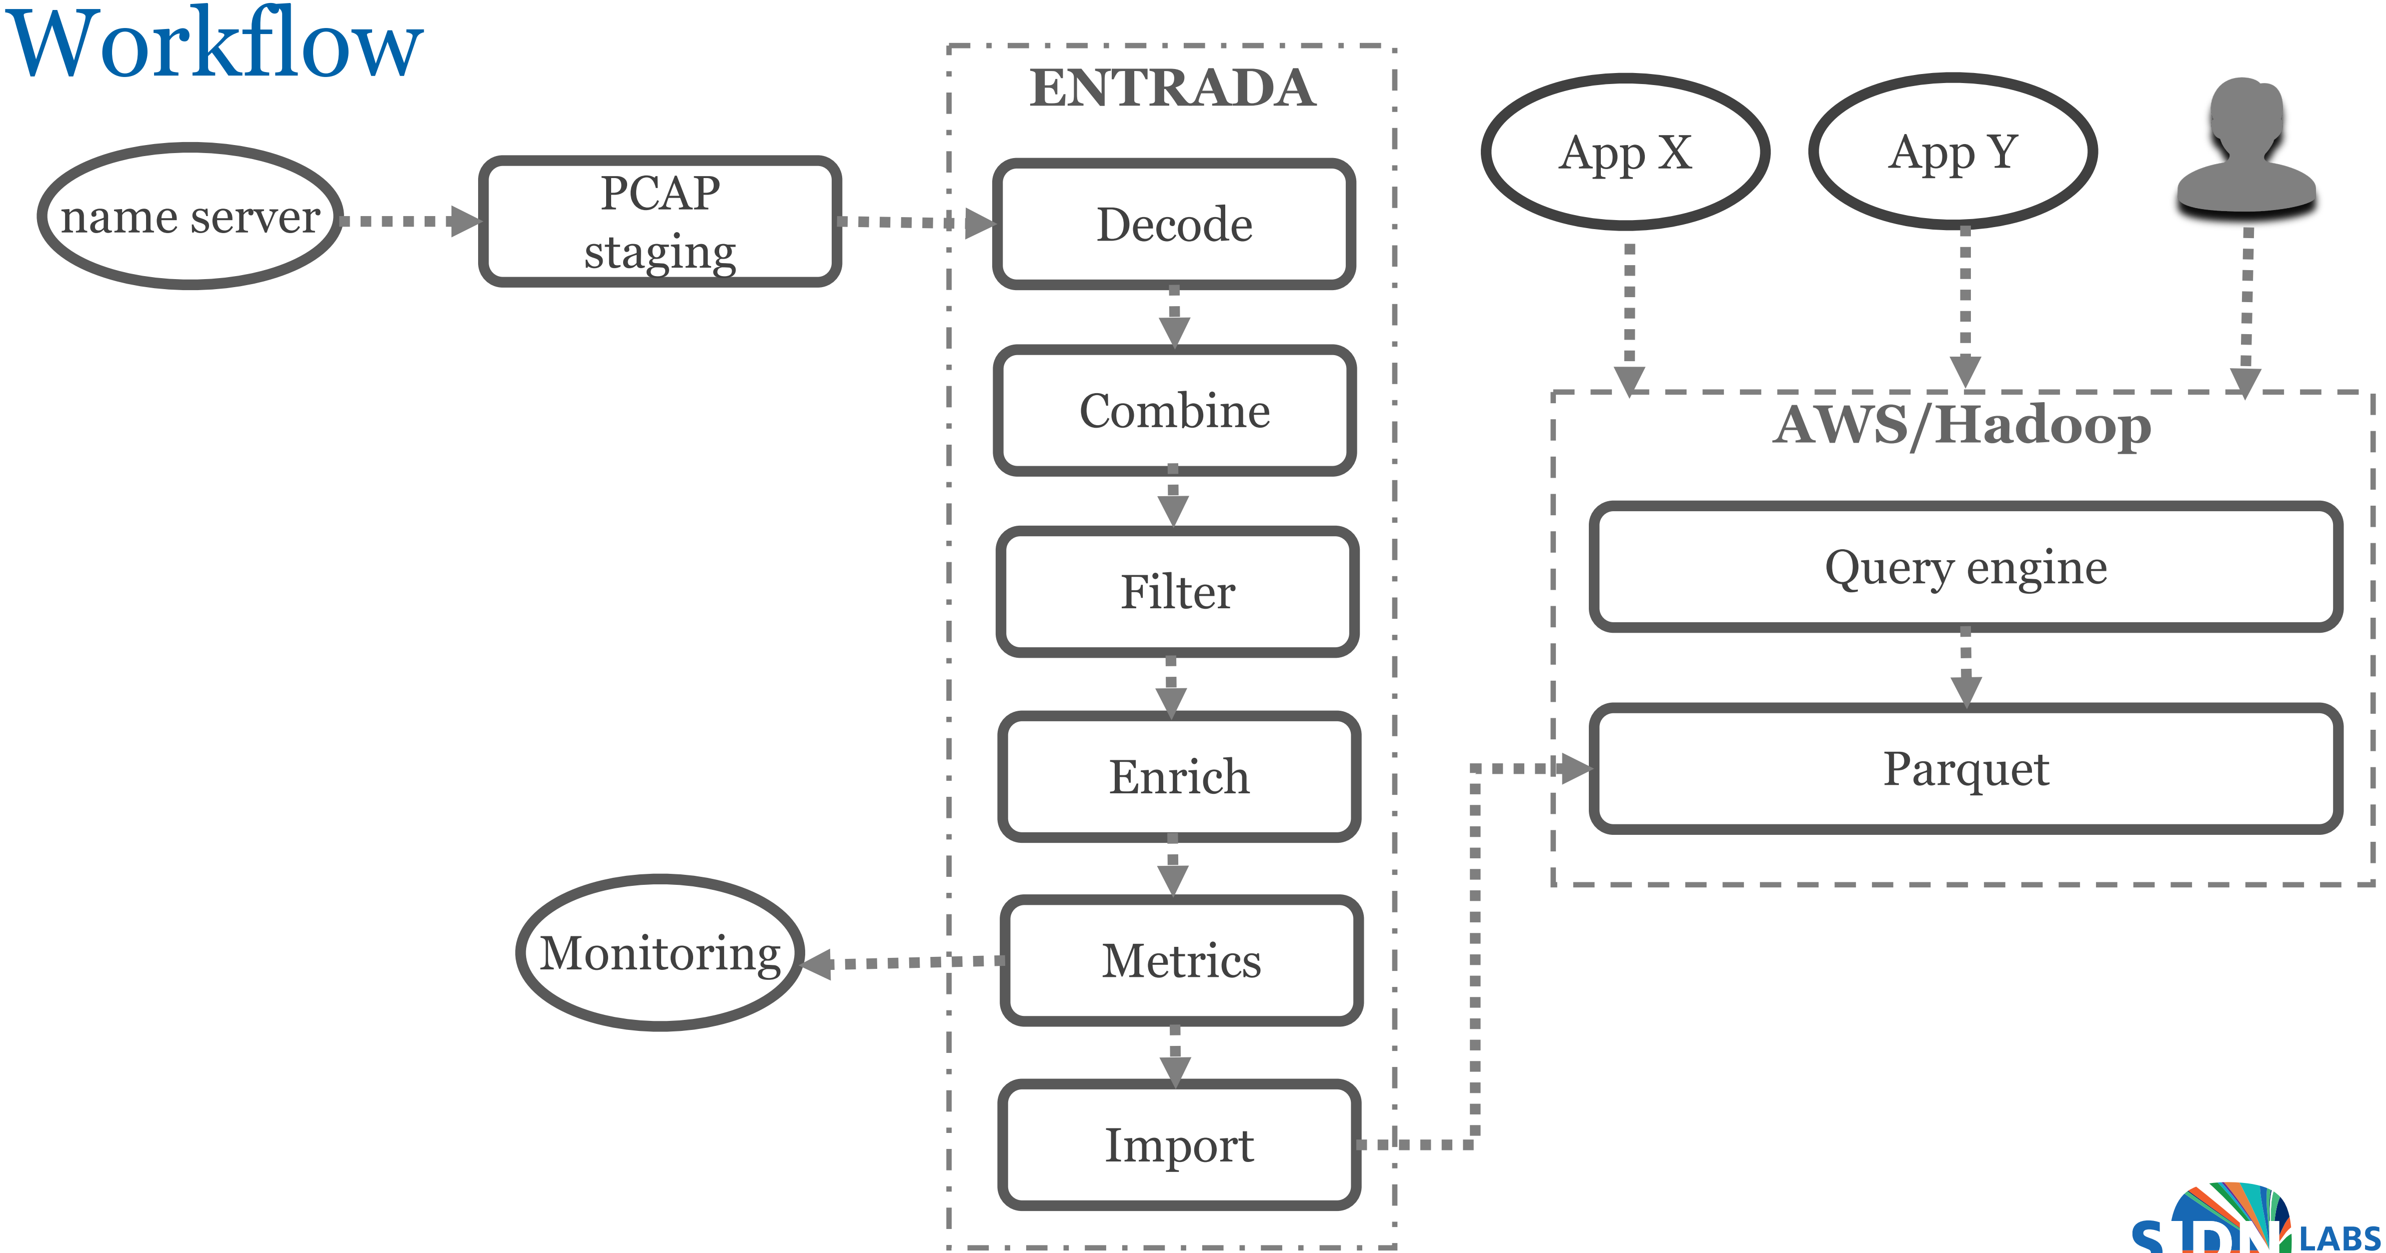

DNS data available for analysis in ~10-15 minutes

Challenges

Previous ENTRADA versions required Hadoop

- Requires more effort to install and maintain
- Requires Hadoop knowledge
- Need to have hardware of virtual cluster

ENTRADA 2

New features:

- Serverless DNS analytics
- Support for multiple SQL query engines
- Quality of service monitor, round-trip time (RTT) analysis
- Easy deployment using Docker

Serverless Computing

“Cloud-computing execution model in which the cloud provider runs the server, and dynamically manages the allocation of machine resources” [1]

[1] https://en.wikipedia.org/wiki/Serverless_computing

Serverless DNS analytics

- No need to deploy any servers
- No hardware/network maintenance cost
- Only pay for amount of data analyzed
- Focus on analyzing the data.

ENTRADA will:

- Create database schema
- Convert, upload and optimize data

Serverless DNS analytics

Support for Amazon Web Services (AWS)

- S3 for data storage
- Athena SQL-query engine to analyze the data on S3
- Pricing;
  - S3: $0.0125 per GB ~ $12,8 per TB per month[1]
  - Athena: $5 per TB of scanned data

[1] S3 pricing depends on selected AWS region and storage class

Amazon Athena

Provides API and a web-based query-interface

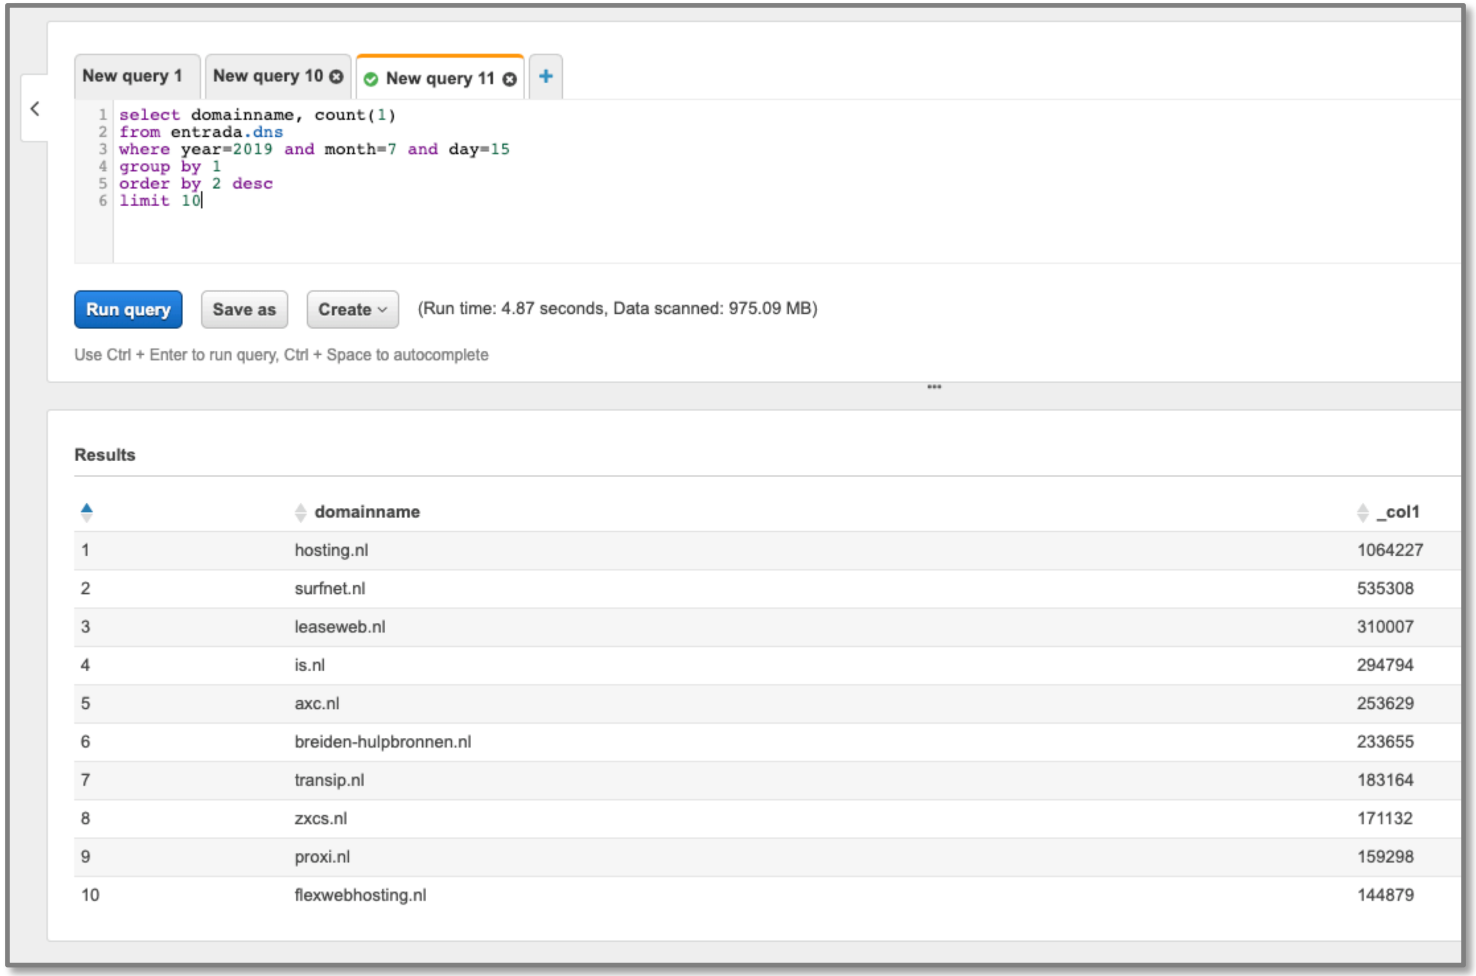

```
select domainname, count(1)
from entrada.dns
where year=2019 and month=7 and day=15
group by 1
order by 2 desc
limit 10
```

(Run time: 4.87 seconds, Data scanned: 975.09 MB)

Use Ctrl + Enter to run query, Ctrl + Space to autocomplete

Results

	domainname	_col1
1	hosting.nl	1064227
2	surfnet.nl	535308
3	leaseweb.nl	310007
4	is.nl	294794
5	axc.nl	253629
6	breiden-hulpbronnen.nl	233655
7	transip.nl	183164
8	zxcs.nl	171132
9	proxi.nl	159298
10	flexwebhosting.nl	144879

Quality of service monitoring

Use passive DNS-data from real world DNS-clients (not probes) to determine Round Trip Time (RTT) between a resolver and the authoritative name server.

High RTT can be caused by:

- Inefficient routing
- Congestion
- Router/switch issues
- …

Quality of service monitoring (QoS)

Analyze the TCP-handshake

For an average day, TCP use:

- 4-5% of queries
- 22-26% of resolvers

diff(SYN ACK + ACK) = RTT

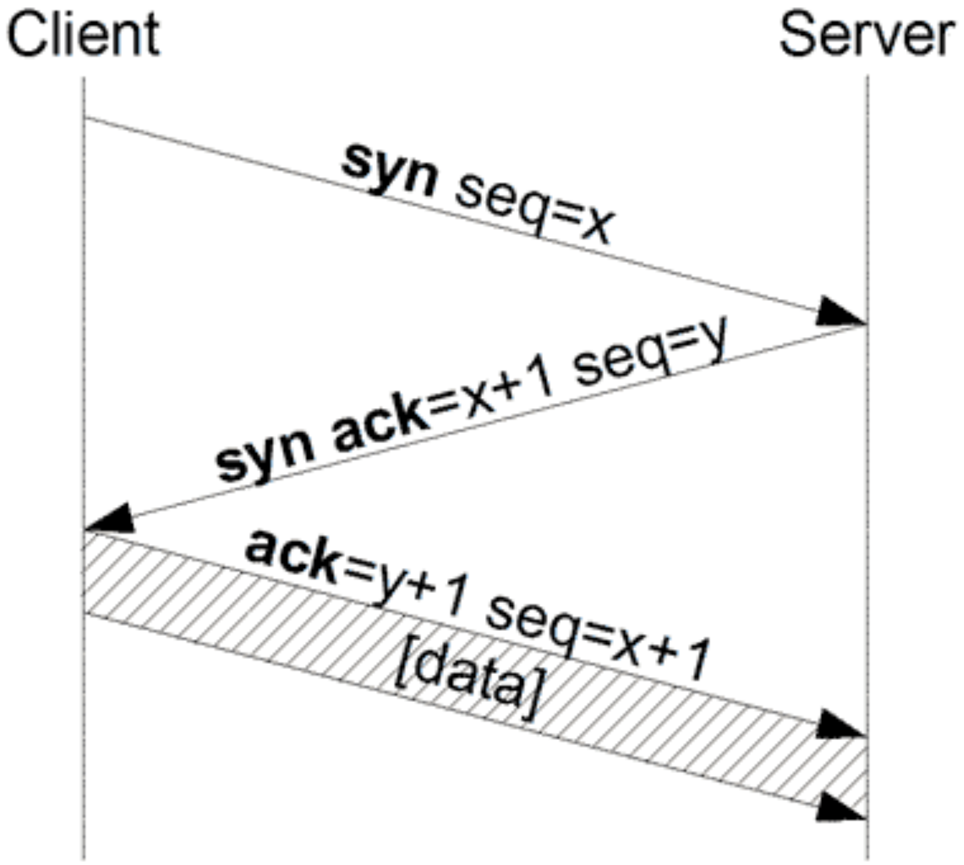

TCP-queries

Percentage of

TCP queries in nl-zone

(October 2019)

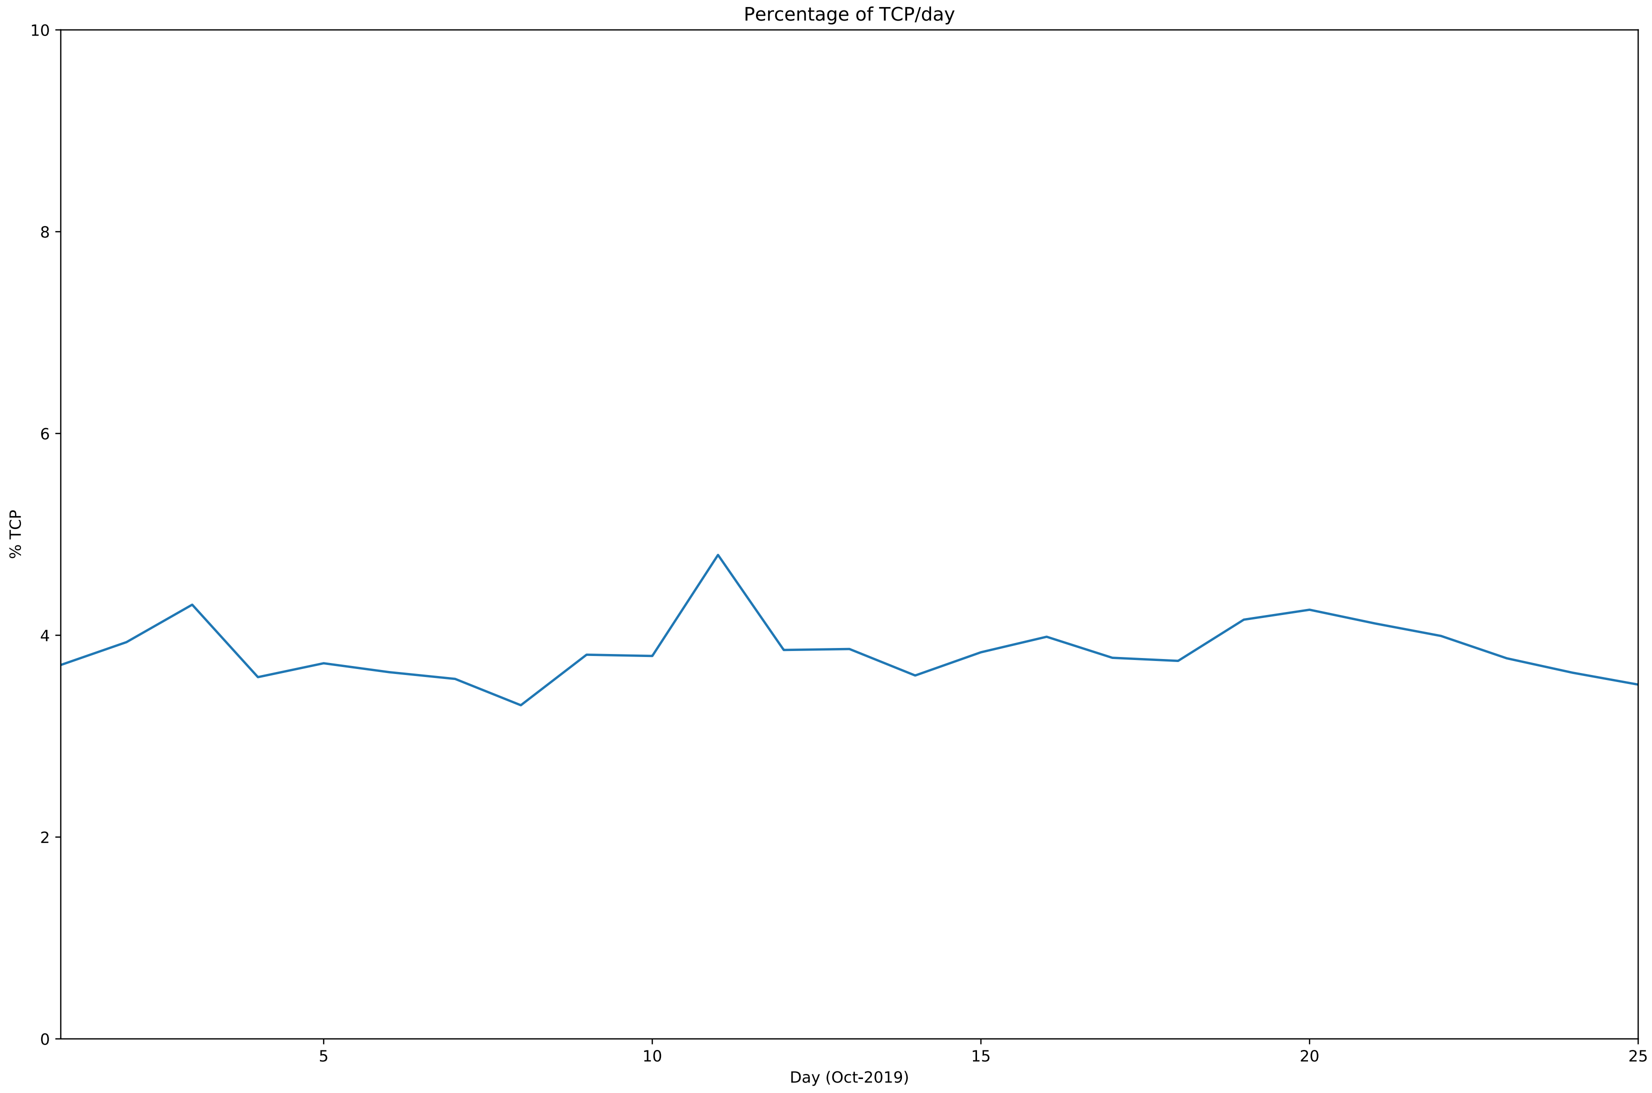

TCP-Resolvers

Percentage of

TCP resolvers in nl-zone

(October 2019)

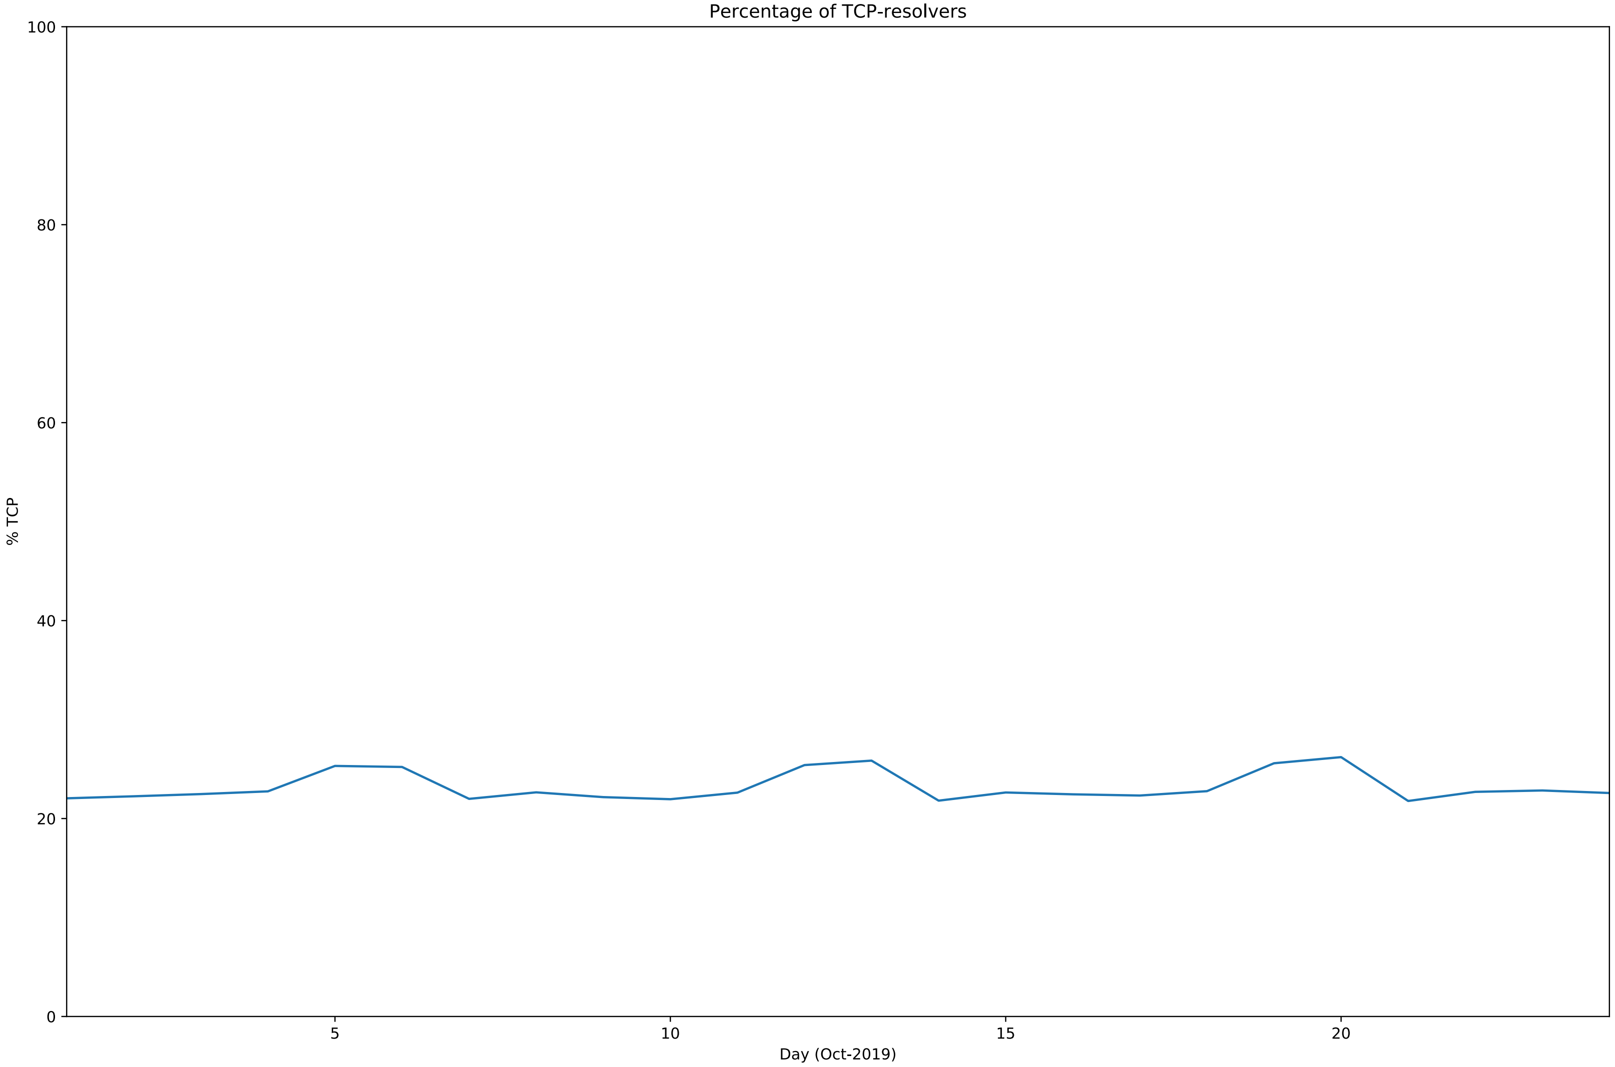

High latency ASNs

Example SQL-query to get the top 100 ASNs with a high average RTT

```
select asn, round(avg(tcp_hs_rtt)) avg_rtt, count(1) samples
from dns
where year=2019 and month=10 and day=24 and tcp_hs_rtt is not null
group by 1
having count(1) > 500
order by 2 desc
limit 100
```

High latency ASNs

asn	Avg RTT	Samples	Operator	Country
45090	1147	3130	Shenzhen Tencent Computer Systems Company Limited	CN
34205	436	716	PJSC Rostelecom	RU
24361	357	2326	CERNET2 IX at Southeast University	CN
4538	298	6256	China Education and Research Network Center	CN
56044	298	4721	China Mobile communications corporation	CN
132525	287	1733	HeiLongJiang Mobile Communication Company Limited	CN
1221	282	4478	Telstra Corporation Limited	AU
56042	280	738	China Mobile communications corporation	CN
24444	274	6361	Shandong Mobile Communication Company Limited	CN

AS45090

Relatively normal RTT until oct 20th.

Then spike in TCP queries and higher RTT.

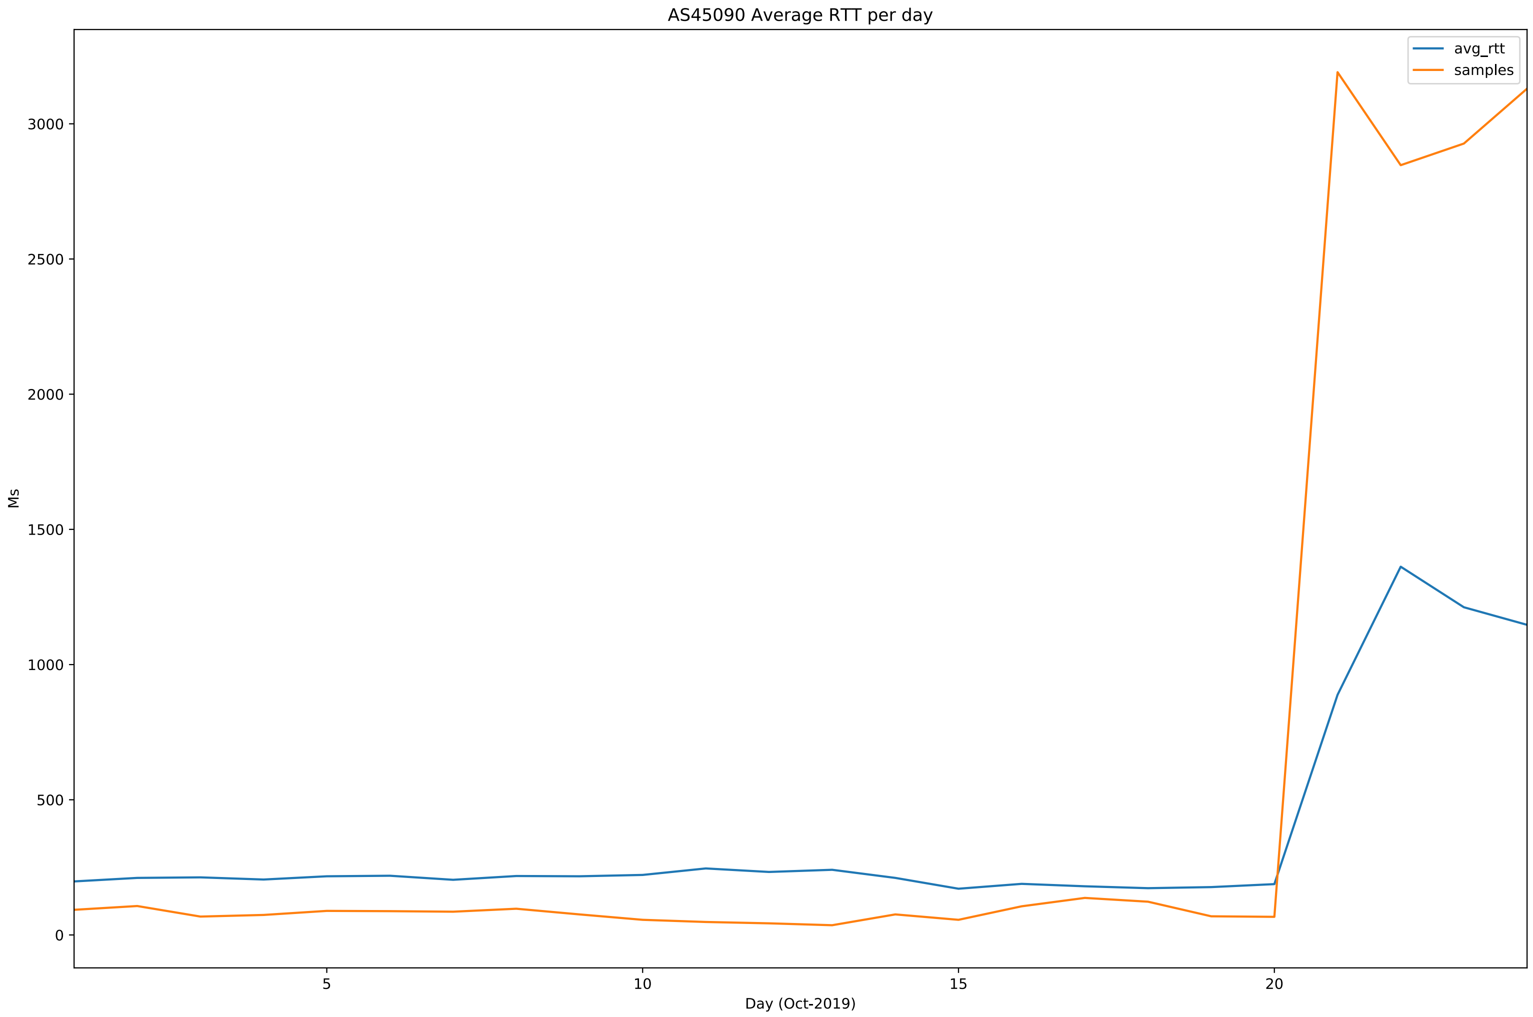

QoS monitoring

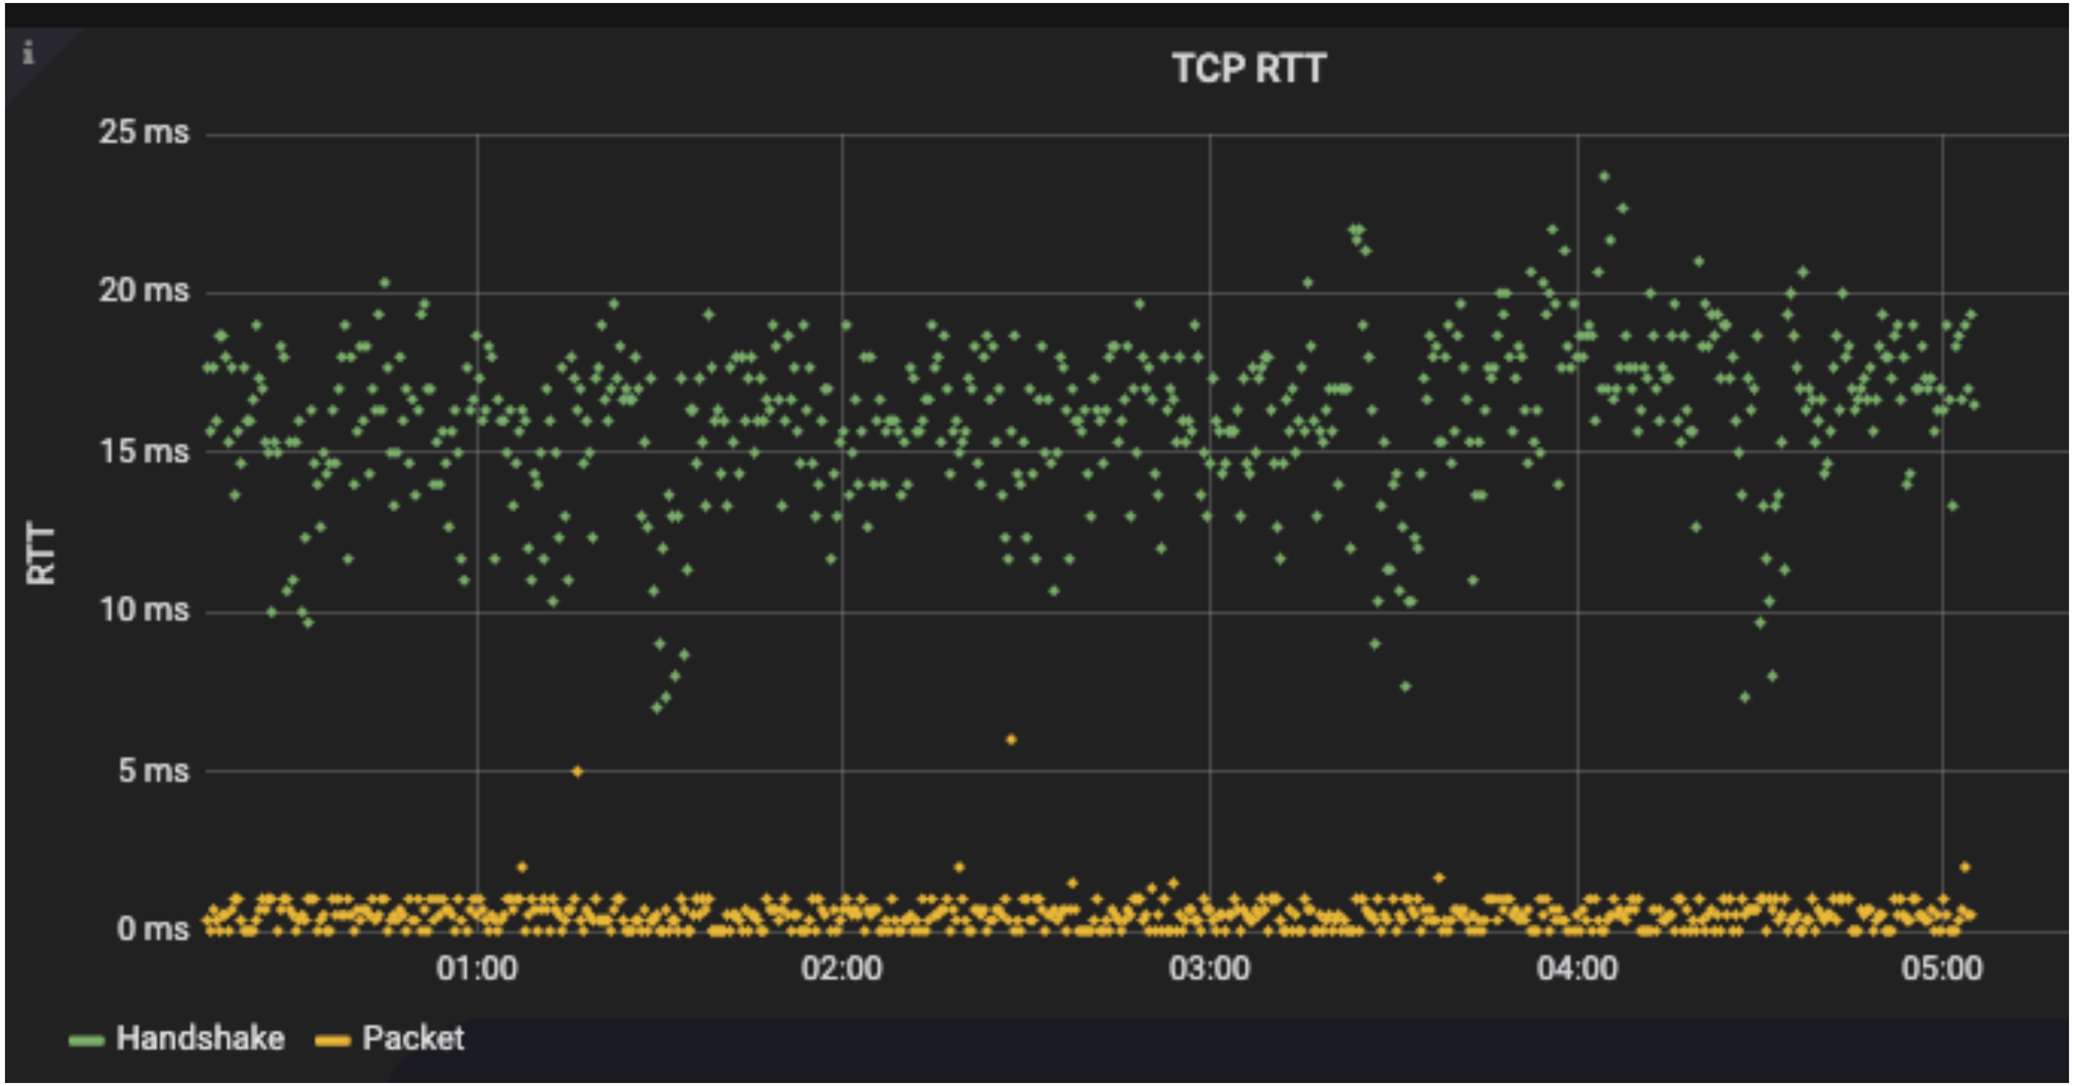

Deployment

- Deploy using Docker
- Use Hadoop, inhouse or in the cloud

OR

- Use AWS S3 + Athena

https://entrada.sidnlabs.nl

Deployment

“A Docker container image is a lightweight, standalone, executable package of software that includes everything needed to run an application”[1]

[1] https://www.docker.com/resources/what-container

Deployment

- Lightweight container
- Pull image from public repository[1]
- Easy configuration using docker-compose

[1] https://hub.docker.com/r/sidnlabs/entrada

ENTRADA@SIDN Labs

- Operational for 4 year
- 2 anycast .nl name servers (28 sites)
- 1,2 trillion (1,2 x $10^{12)}$ records (DNS query and response pair)
- 65 TB data (3x replication = 195TB)

.nl SIDN.nl

@SIDN

in SIDN

Q&A

www.sidnlabs.nl | stats.sidnlabs.nl