OpenINTEL

an infrastructure for long-term, large-scale and high-performance active DNS measurements

UNIVERSITY OF TWENTE.

DACS
Design and Analysis of Communication Systems

SIDN labs

SURF NET

Why measure DNS?

- (Almost) **every networked service relies on DNS**
- DNS translates human readable names into machine readable information
  - e.g. IP addresses, but also: mail hosts, certificate information, …
- Measuring **what is in the DNS over time** provides information about the **evolution of the Internet**
- (we started this because we were interested in the rise of DDoS Protection Services)

Goals and Challenges

- Send a **comprehensive set of DNS queries for every name** in a TLD, **once per day**

- But can we do this at **scale**? How does this **impact** the global DNS?

  .com + .net + .org ≈ over 150 million names
  (about 50% of the global DNS namespace)

- How do we store and analyse this data efficiently?

Data collection stages

- We distinguish **three stages** for data collection:

  - **Stage 1:**
    Collection of zone files for TLDs to scan, compute daily deltas

  - **Stage 2:**
    Main measurement, perform queries for each names, collect meta data, store results

  - **Stage 3:**
    Prepare data for analysis

High-level architecture

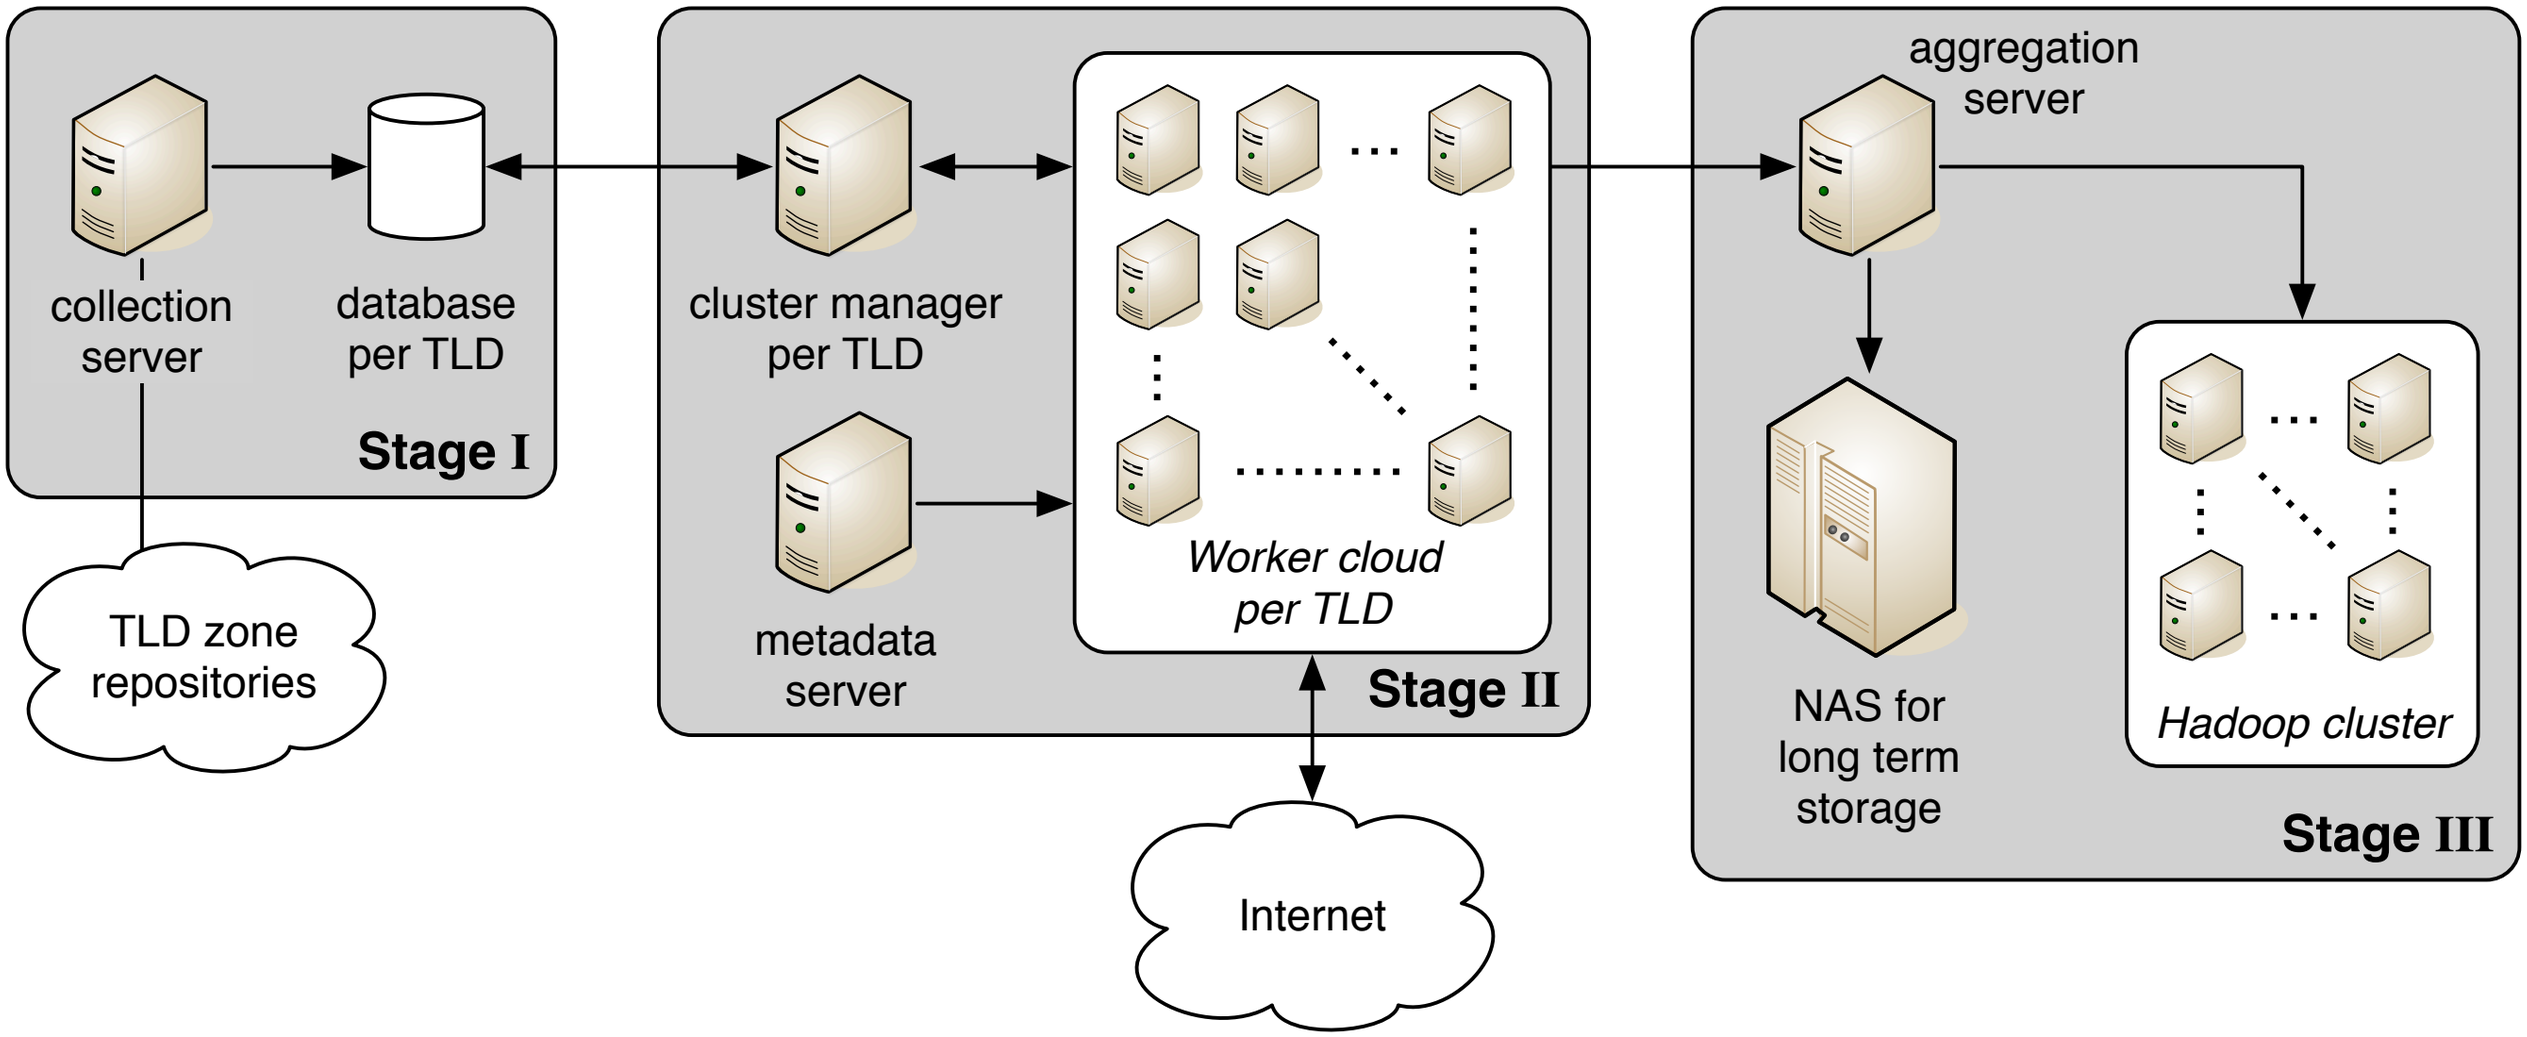

What do we query and store?

- We ask for:
  - SOA
  - A, AAAA
    - (apex, ‘www’ and ‘mail’)
  - NS
  - MX
  - TXT
  - SPF
  - DS
  - DNSKEY
  - NSEC(3)
- We store:
  - All records in the *answer* section
  - CNAME expansions
  - DNSSEC signatures (RRSIG)
  - Metadata (Geo IP, AS)

Impact on the global DNS

- Our measurement is clearly visible in SURFnet's traffic flows

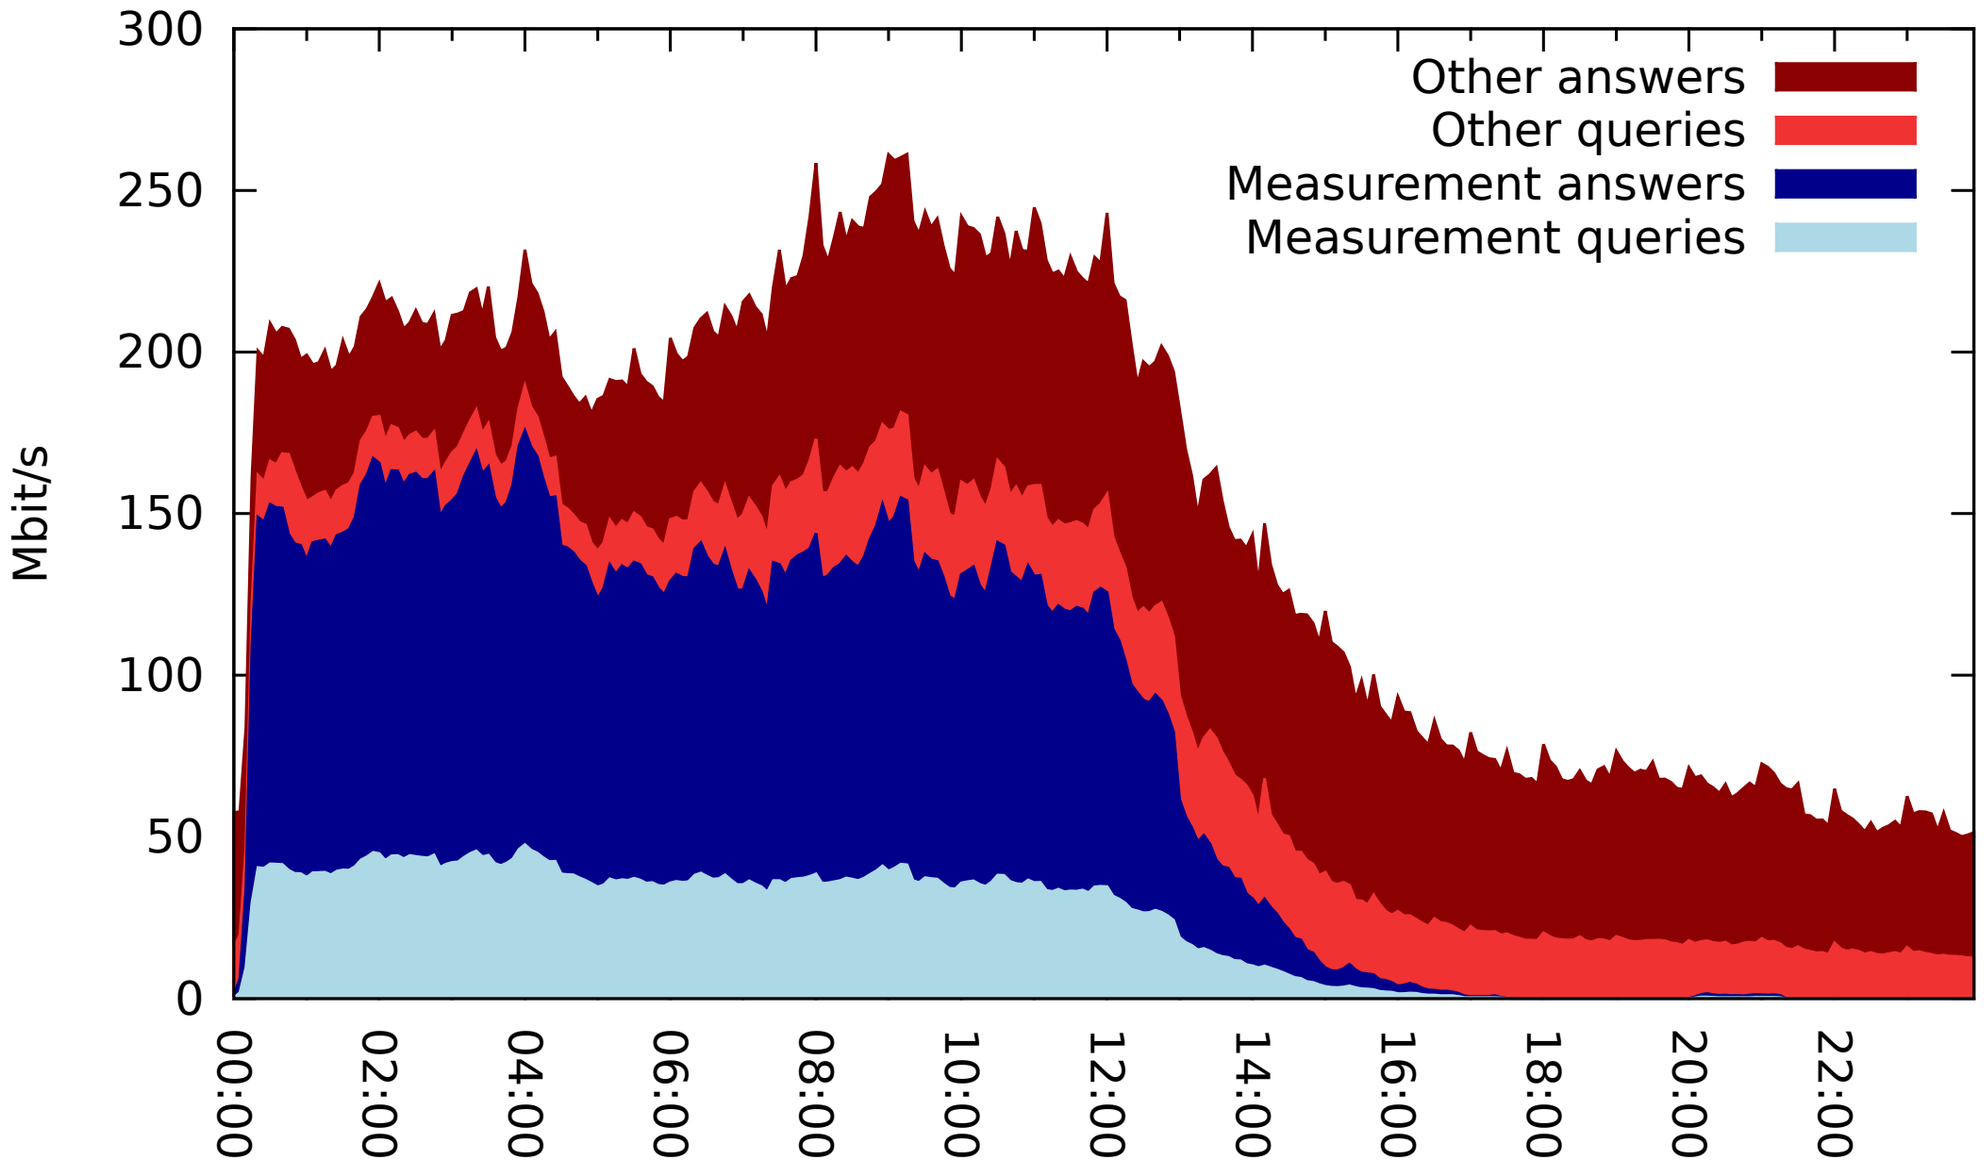

Impact on the global DNS

- Deeper analysis shows very few top talkers (less than 35 receive more than 100 packets/sec.)

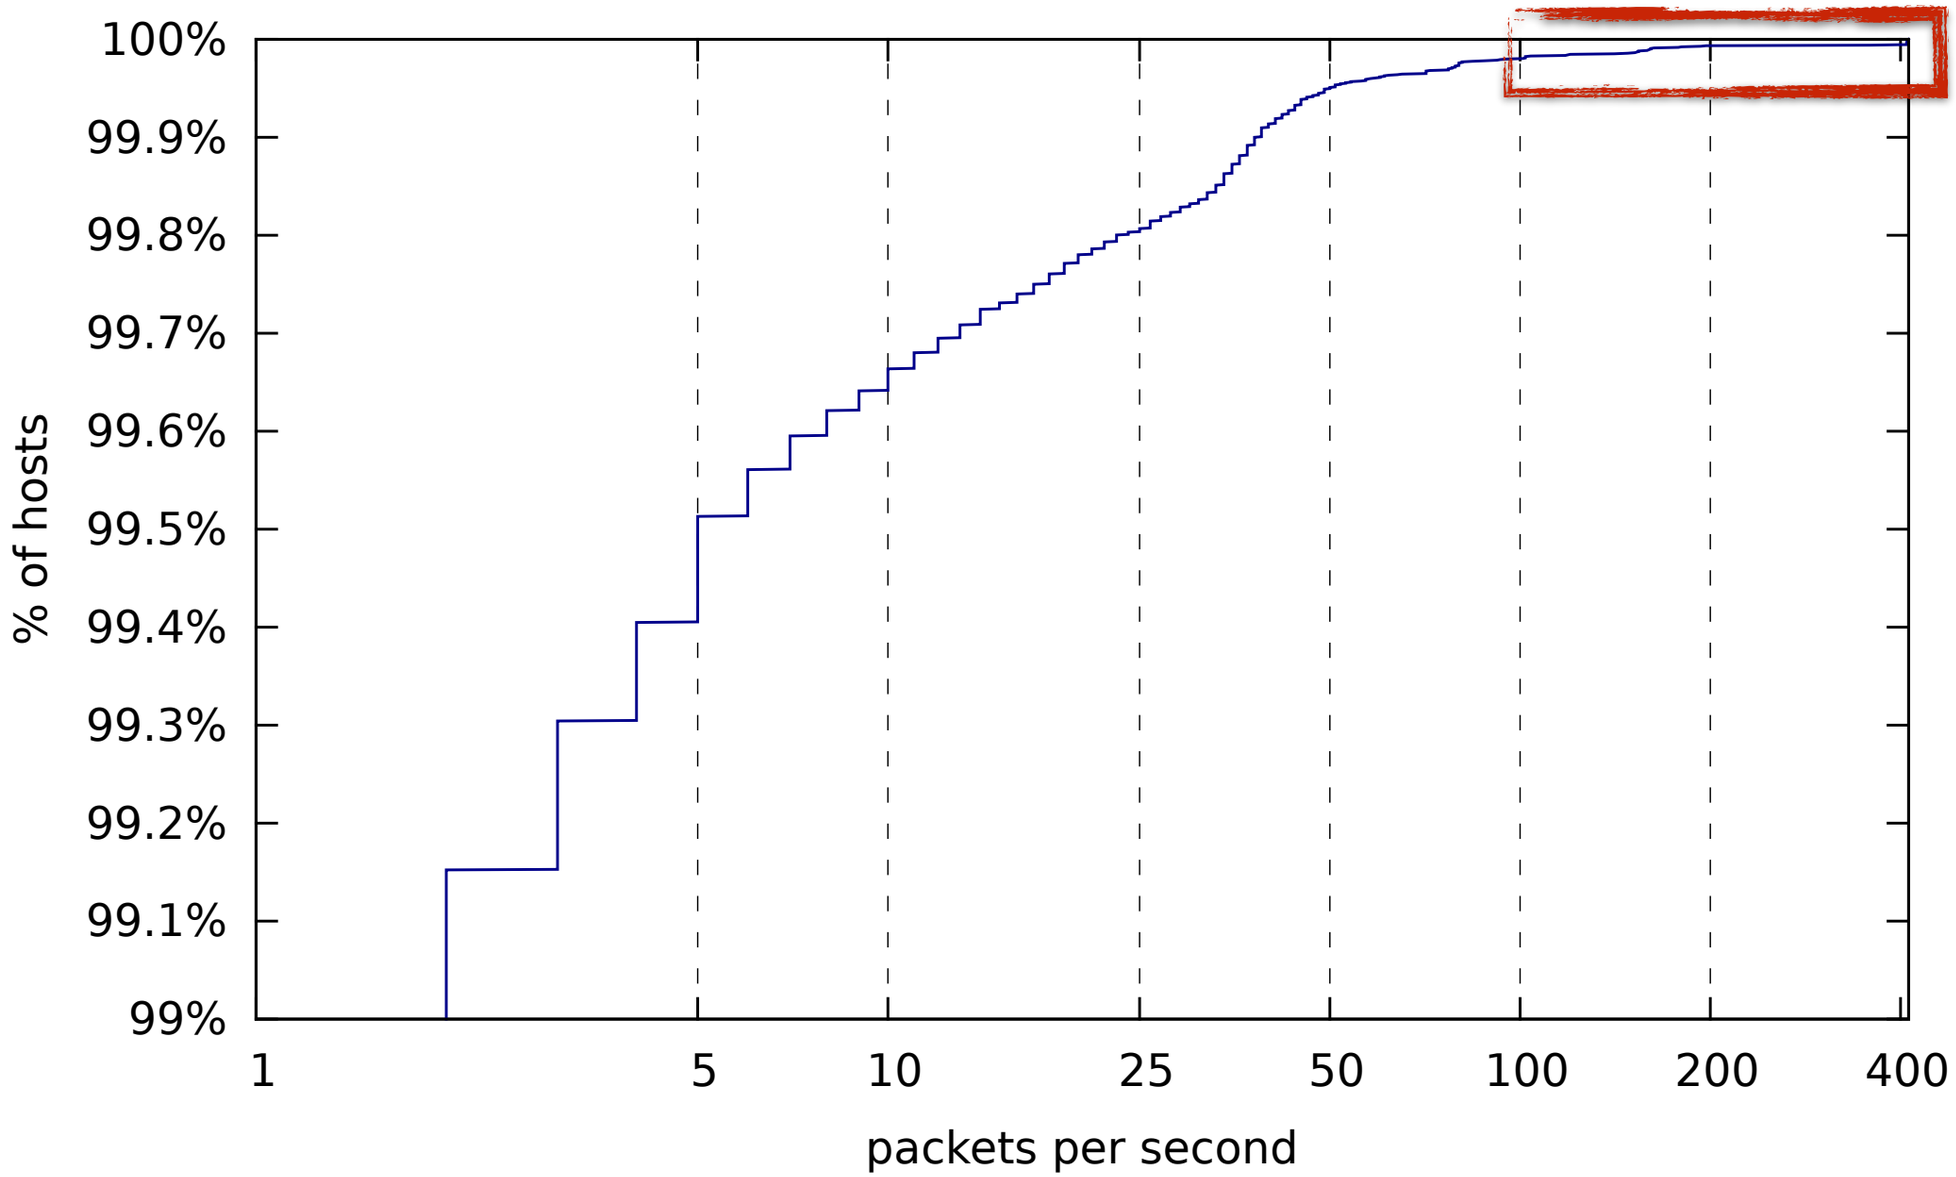

Big data? Yes!

- Calling your research “big data” is all the rage

- So would our work qualify as big data?

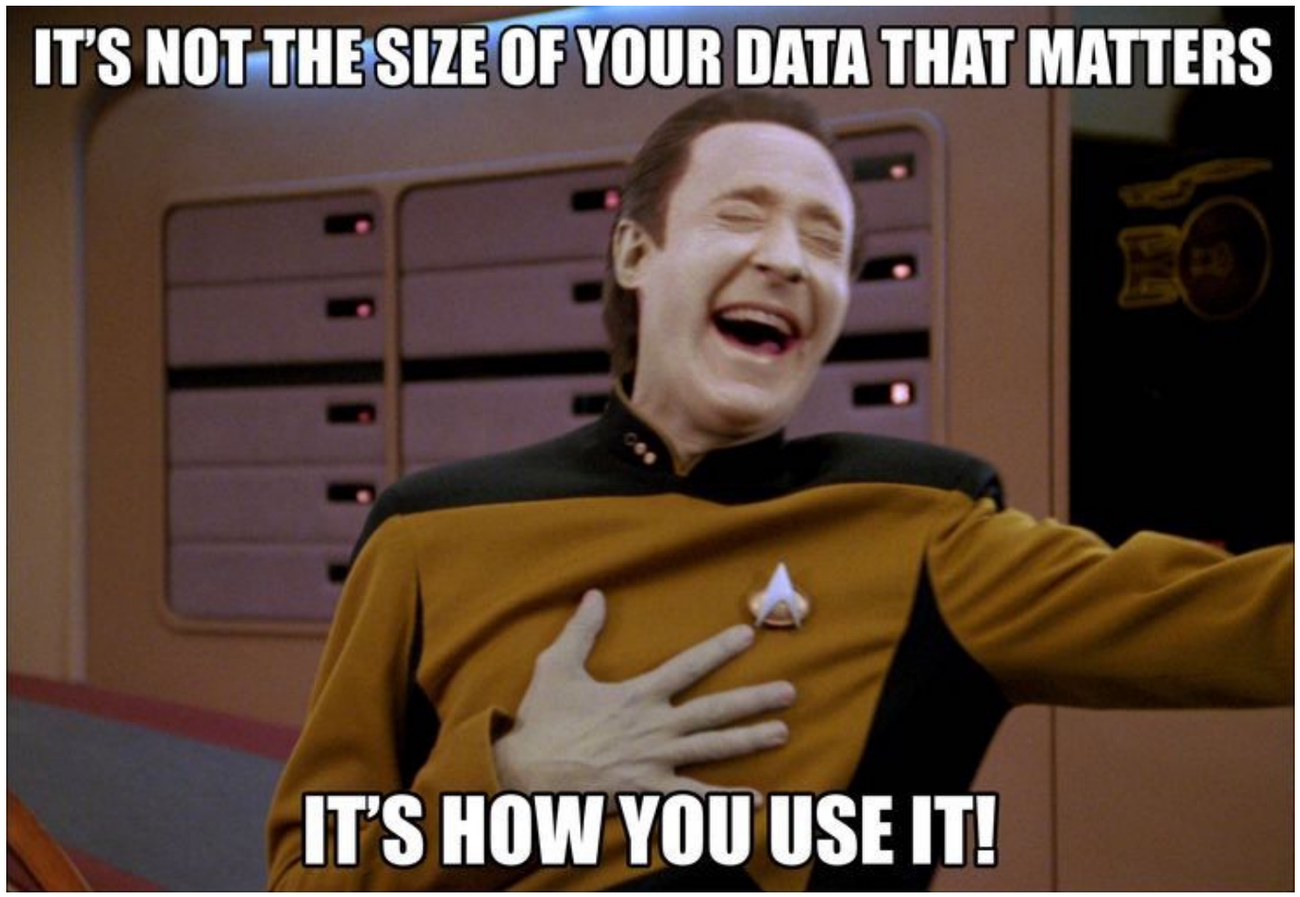

- The **human genome** is about $\mathbf{3 \cdot 10^9}$ base pairs

- We collect around $\mathbf{1.8 \cdot 10^9}$ DNS records **per day**

- Since February 2015, through December 31st we collected $\mathbf{511 \cdot 10^9}$ (511 **billion**) results

Some numbers

- Workers: 1 CPU core, 2GB RAM, 5 GB disk

TLD	#domains	workers	measure time
.org	10.9M	10	7h19m
.net	15.6M	10	14h29m
.com	123.1M	80	17h10m
.nl	5.6M	3	3h09m

- Data collected daily:

TLD	#domains	(failed)	#results	Avro	Parquet	uncompressed
.org	10.9M	(1.2%)	125M	2.6GB	3.2GB	18.5GB
.net	15.6M	(0.9%)	166M	3.5GB	4.3GB	24.4GB
.com	124.0M	(0.6%)	1419M	30.0GB	36.8GB	213.4GB
.nl	5.6M	(0.5%)	112M	8.5GB	11.8GB	27.8GB
total	**156.1M**	**(0.6%)**	**1.8B**	**43.3GB**	**54.7GB**	**284.1GB**

- We collect almost **16TB** of **compressed** data **per year**

Big data? Use the right tools

- With 3 partners invested in a Hadoop cluster (SURFnet, SIDN, UTwente)

- Use latest & greatest tools for analysis, **Impala, Spark, Flume,** …

- Working on making datasets accessible to other network researchers

Query performance

- Example query:
  top 10 countries A records
  geo-locate to in the .com TLD

- Storage format matters a lot!

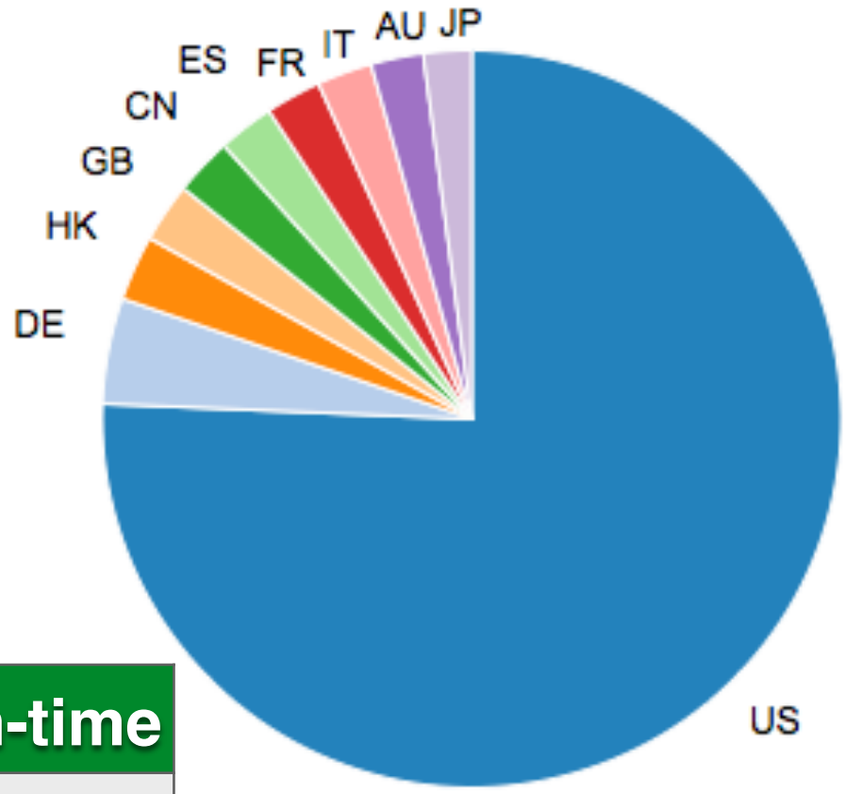

Storage format	Compression	Relative size	Query run-time
Avro (row oriented)	`none`	100%	25.1s
	`deflate`	17%	15.5s
	`snappy`	**23%**	**9.3s**
Parquet (columnar)	`none`	44%	17.5s
	`gzip`	10%	5.7s
	`snappy`	**17%**	**4.3s**

Sweet spot!

An example: cloud e-mail

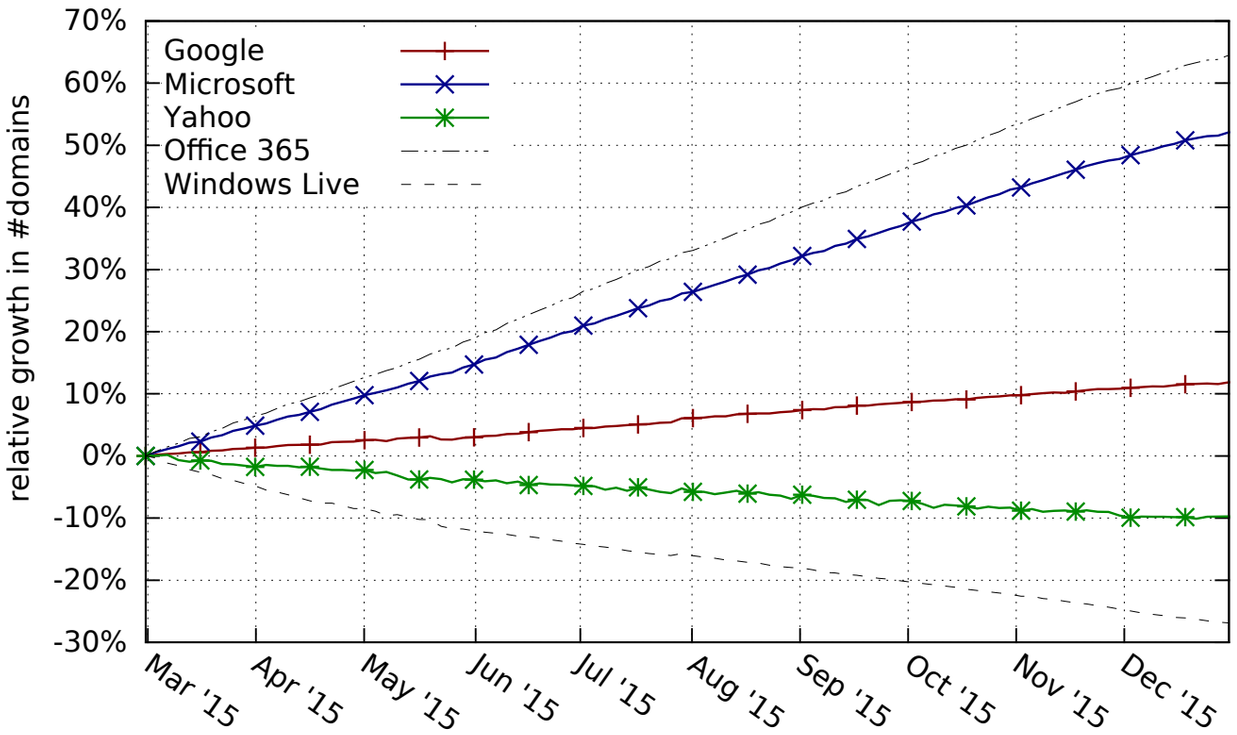

- Google largest (4.57M)
- But Microsoft grows much faster!
- Yahoo in decline

- SPF protects against e-mail forgery
- Microsoft users show (near) ubiquitous SPF use
- Google users at only one third

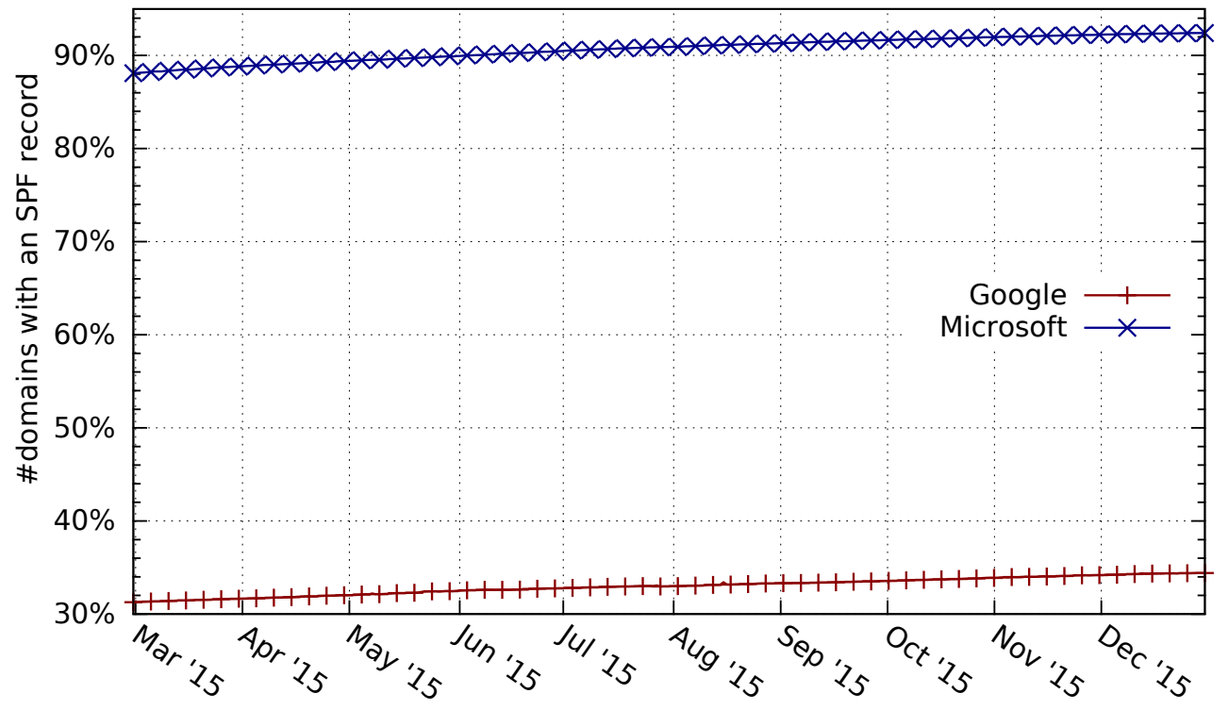

Data access

- Working on ways to make this resource accessible to the measurement research community
- Problem: contracts for zone file access (com/net/org/nl/…) are (very) restrictive
- Current thinking:
  - Publishing aggregate data sets is OK
  - “Toy” cluster with open data (e.g. Alexa 1M) to allow others to write queries & scripts, then execute “on behalf”
  - Anonymisation of data?

Thank you for your attention!

Questions?
(come see us for a live demo)

nl.linkedin.com/in/rolandvanrijswijk
nl.linkedin.com/in/mattijsj

@reseauxsansfil

r.m.vanrijswijk@utwente.nl
m.jonker@utwente.nl

UNIVERSITY OF TWENTE.

SURF NET