Big Data London November 2019

adform

Real-time Business Intelligence with Big Data

Peter Milne @helipilot50 helipilot50

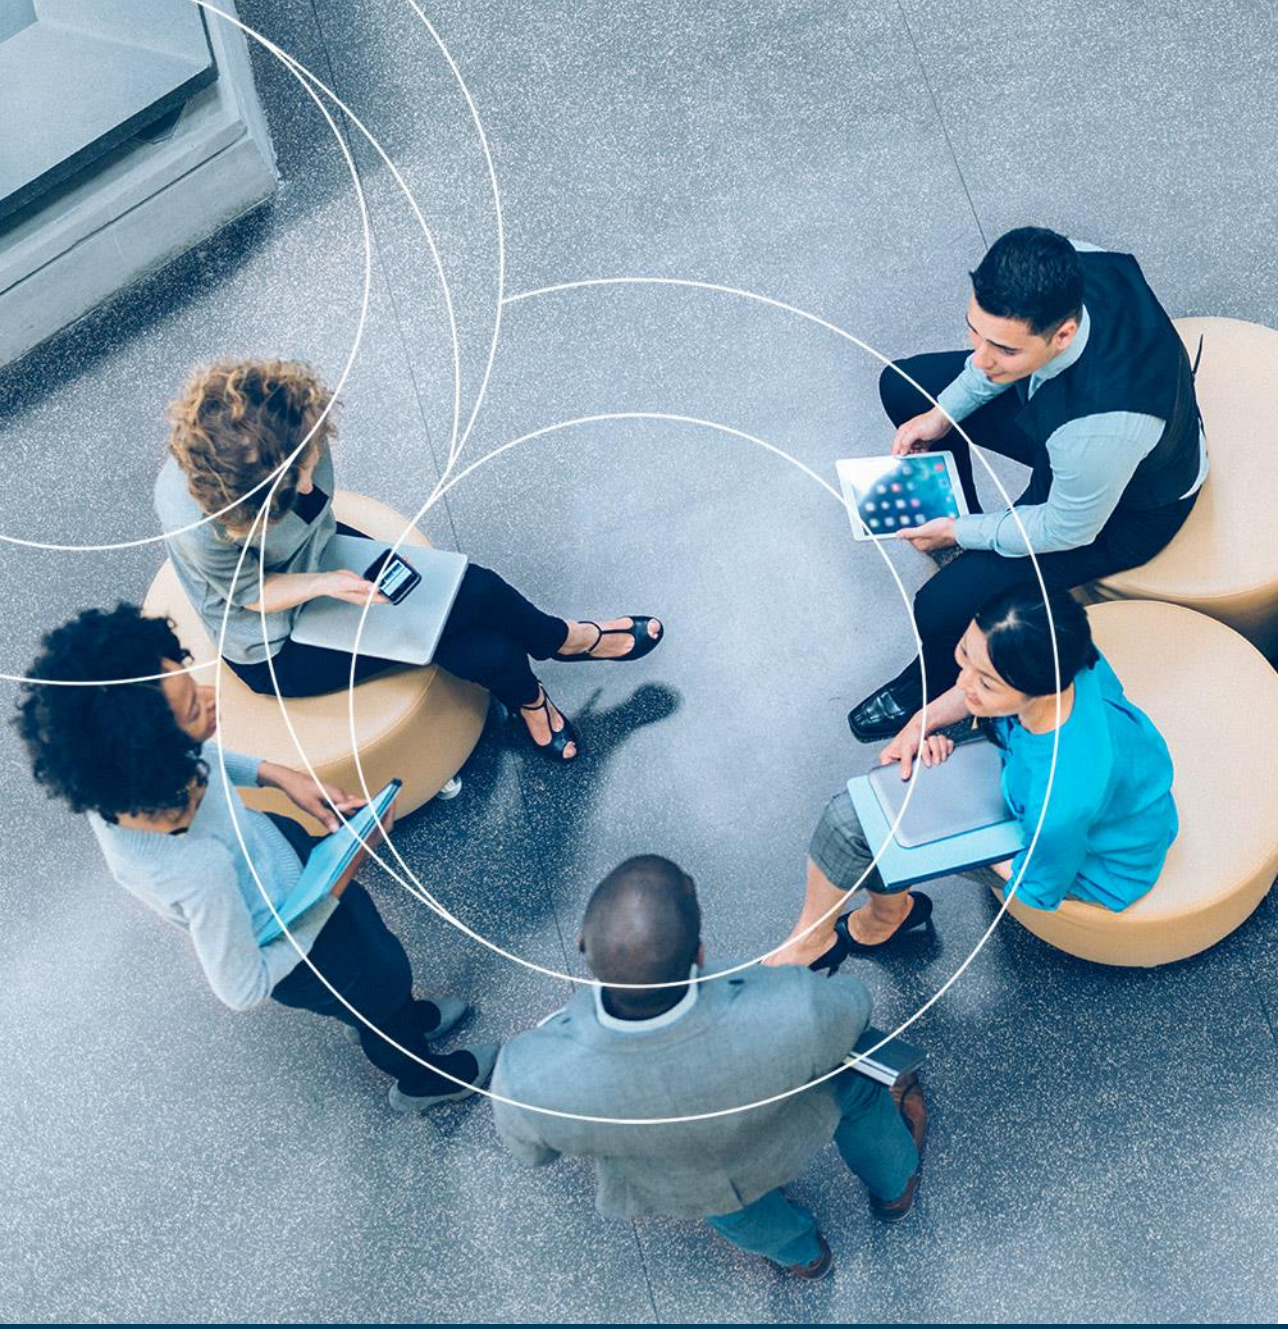

Agenda

- How big - How fast
- Real-time Bidding overview
- Real-time “ish” Campaign Reporting
- Uncle Pete’s Advice

Ad Tech
How big – How fast

How big – How fast

How big

1 byte equals 1 second, **how big is** a:

Bytes	Time
1 Byte	1 second
1 Mega Byte	12 days
1 Giga Byte	34 years
1 Terra Byte	34,700 years
1 Peta Byte	34 million years

Adform has 270-370 million "years" of data

How fast

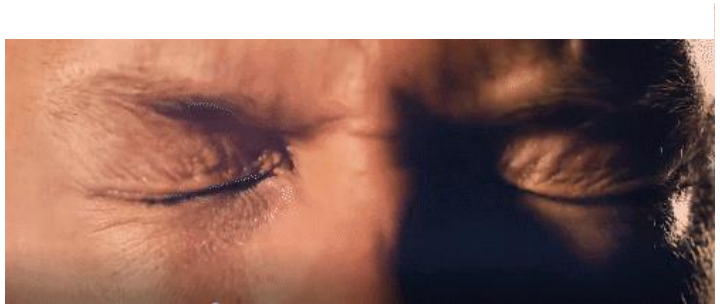

Speed of an **eye blink** is 200-400 ms.

In 1 eye-blink, Adform process 900,000 – 1,700,000 bid requests

Global Infrastructure

9
Datacenters

1800+
Physical Servers

7-11 PB
HDFS Storage

2+ Mil
Kafka Messages per Second

620+ Gb/s
Traffic via core

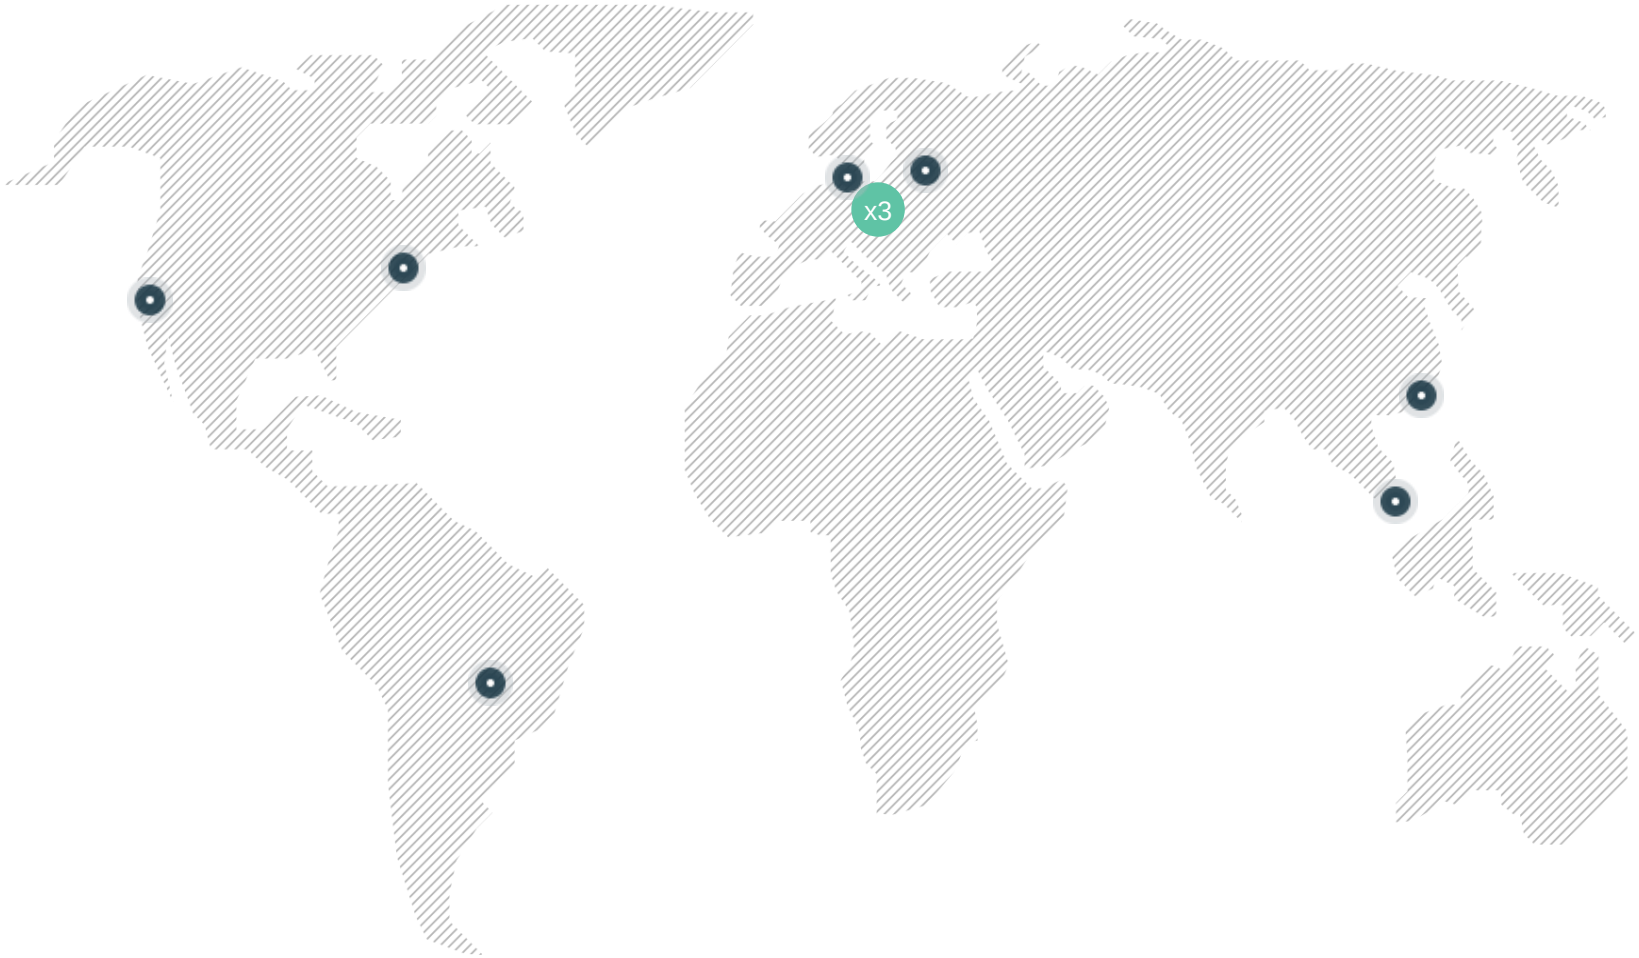

11+ Mil
Aerospike Ops/S

3200+
Openstack VMs

1250+
Hyper-V VMs

130+ Gb/s
Incoming Network Traffic

3,6+ Mil
Nginx Ops/S

Fast Data

Row store – Great for primary key queries

- Total data volume **~22 TeraBytes**
- Average query latency - **0.9 millisecond**
- Biggest set – **9 billion rows**
- **~11 million** operations per **second**


```
SELECT device, geo, segments FROM cookie_store
WHERE cookieId = 826701826
```

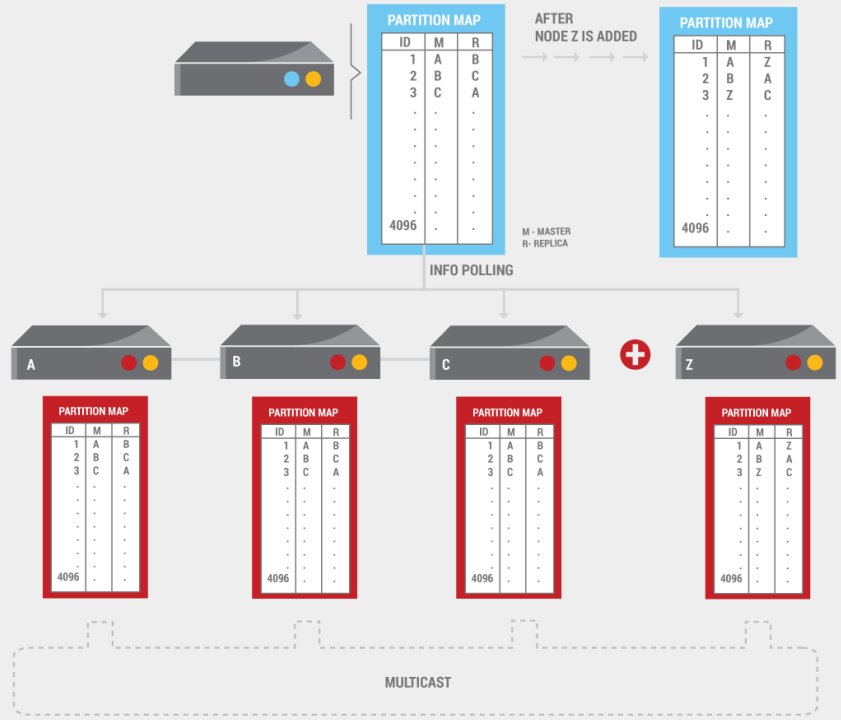

Analytical Data

Scalable Column Store - Great for complex multi-filter queries

- Total data volume **~250 TeraBytes**
- Average query latency **0.541 sec**
- Biggest tables **100 billion rows**
- **300 thousand** reporting queries per **day**
- **Ingest 6 billion** rows per **day**


```
SELECT avg(price) FROM tickstore
WHERE symbol = 'AAPL"  date =
'5/06/09'
```

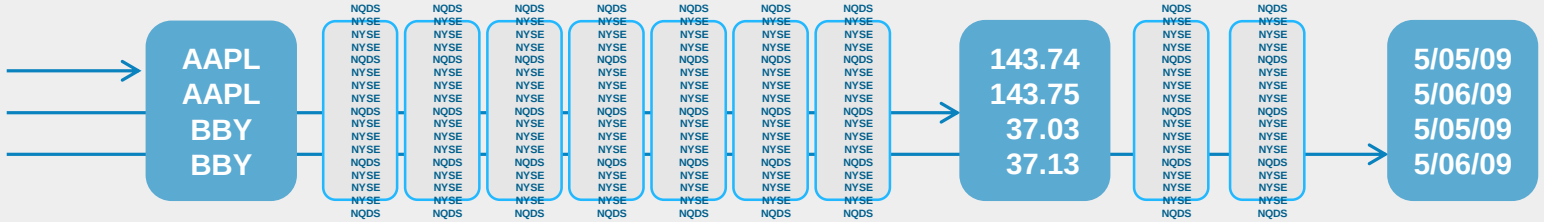

Column Store - Reads 3 columns

Big Data – Data Lake

Raw Data

For:

- Analytical data
- Extract Transform Load (ETL)
- Large data volumes
- Cold storage / data archiving

Example: **~10 Peta Bytes** client campaign data

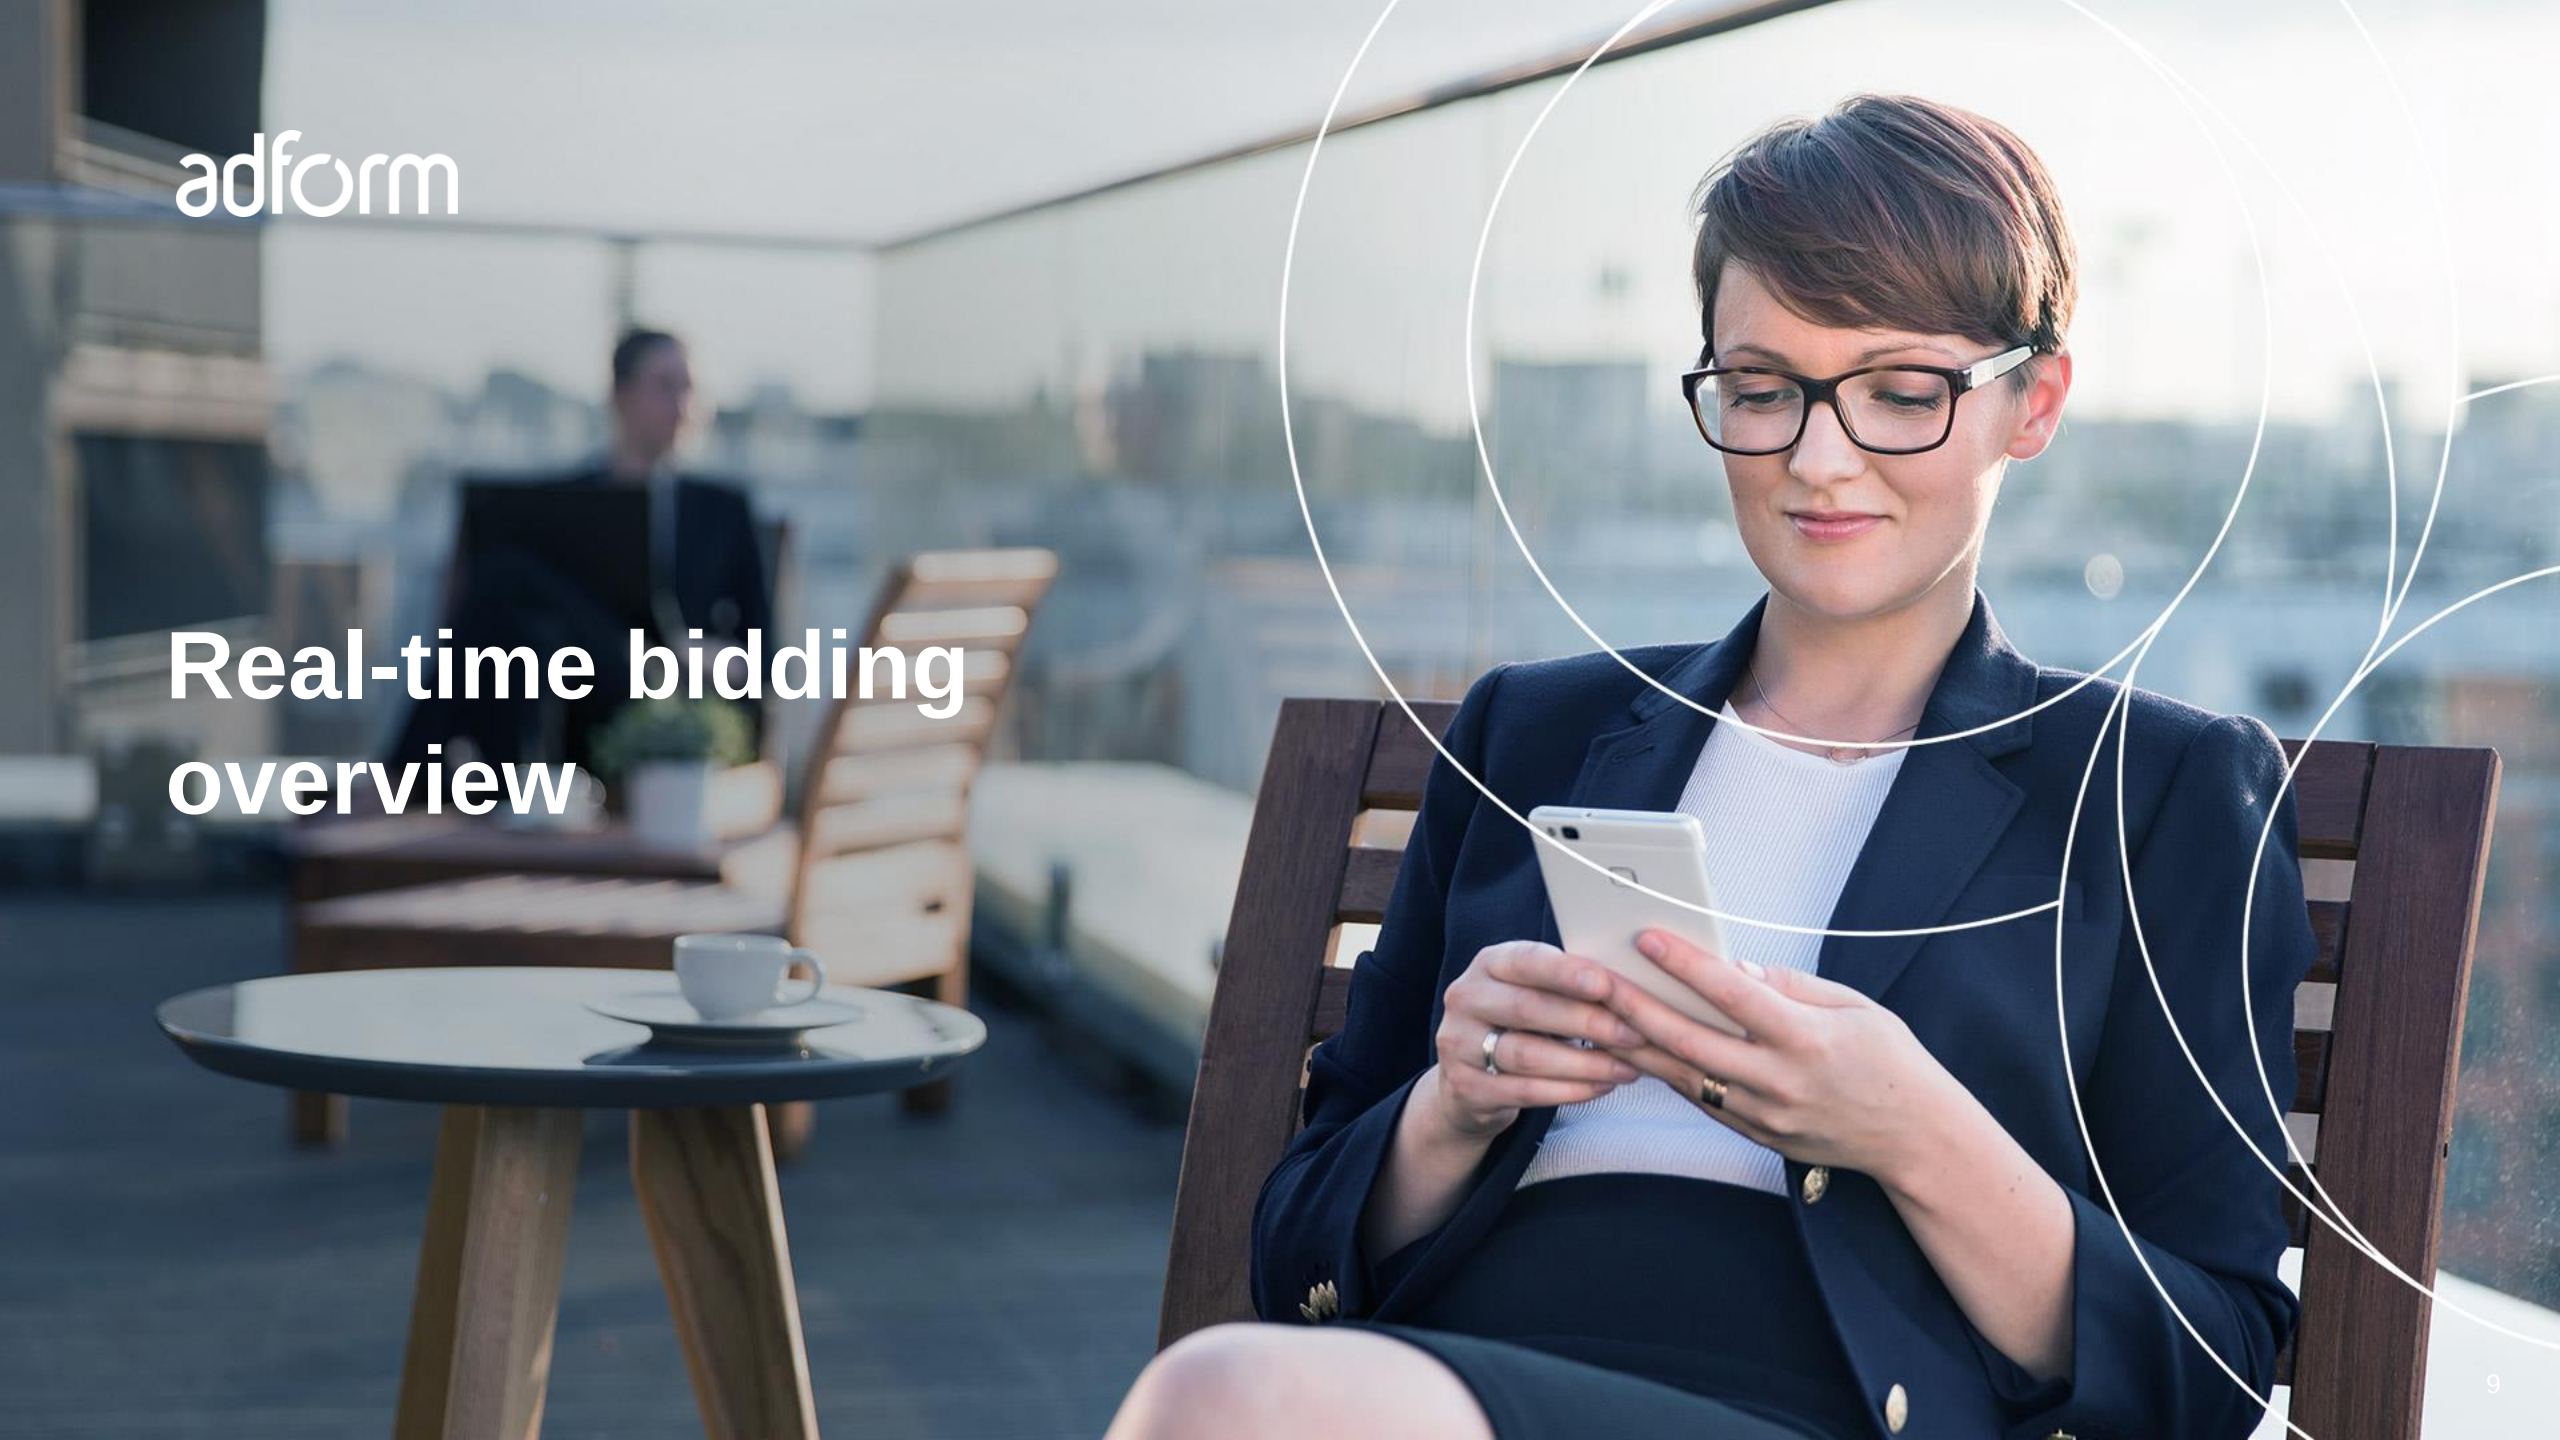

Real-time bidding overview

Digital Advertising

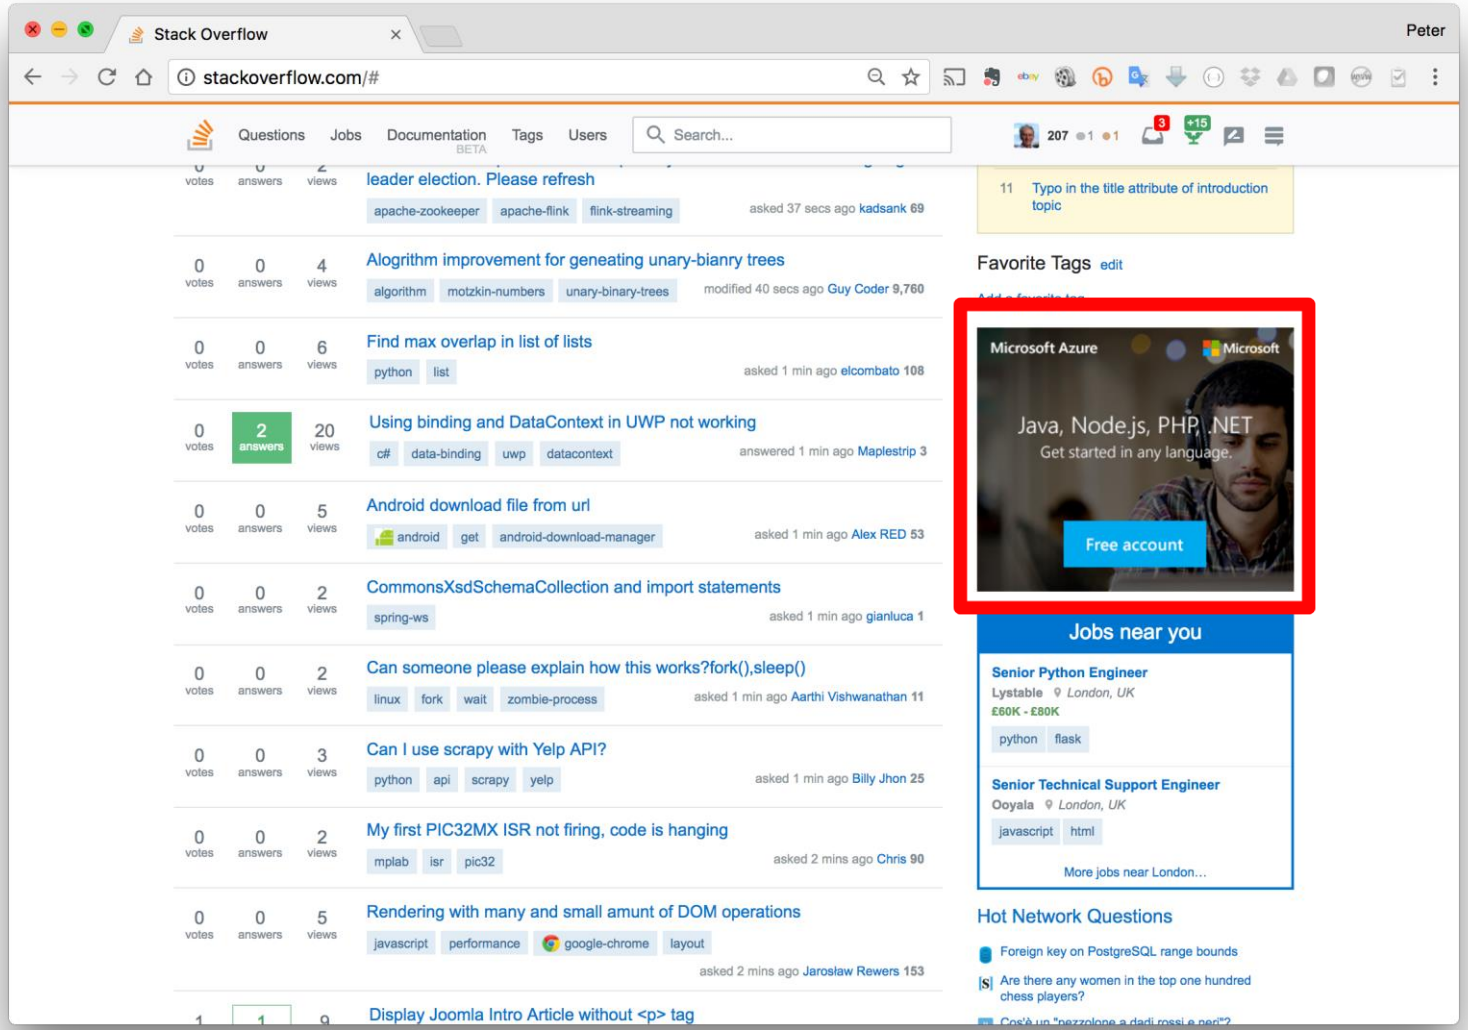

Real Time Biddings - Auction

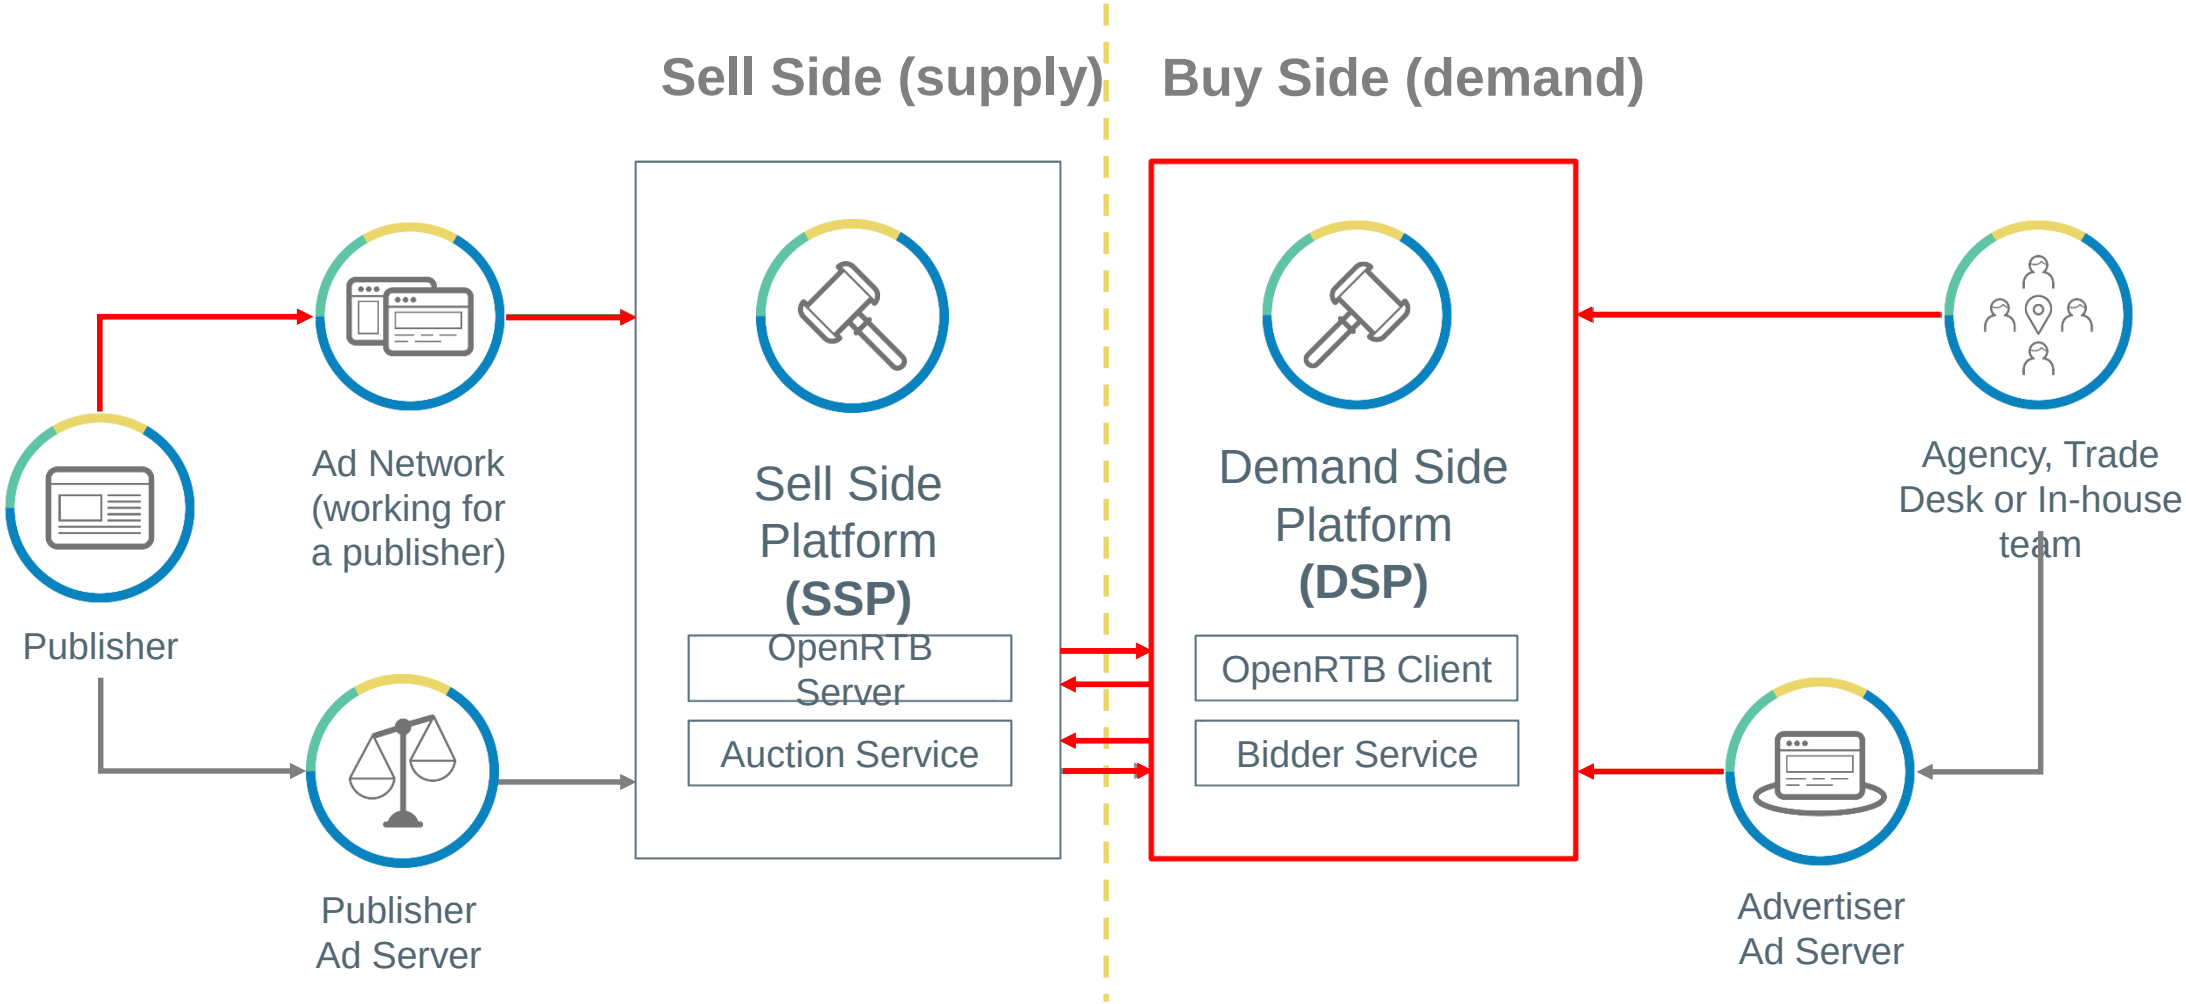

Real Time Bidding (RTB) is the kernel of programmatic advertising


adform
Real-time “ish”
Campaign Reporting

Campaign reporting

Fast reports

- **Real enough** Time for the human user
- Graphs and cards in the UI
- Quick visual indicators
- Latency: **< 1 second**

Slow reports

- **Periodic** reporting
  e.g Daily/Weekly/Monthly
- Statistical analysis
- Latency: **0.5 – 10 minutes**

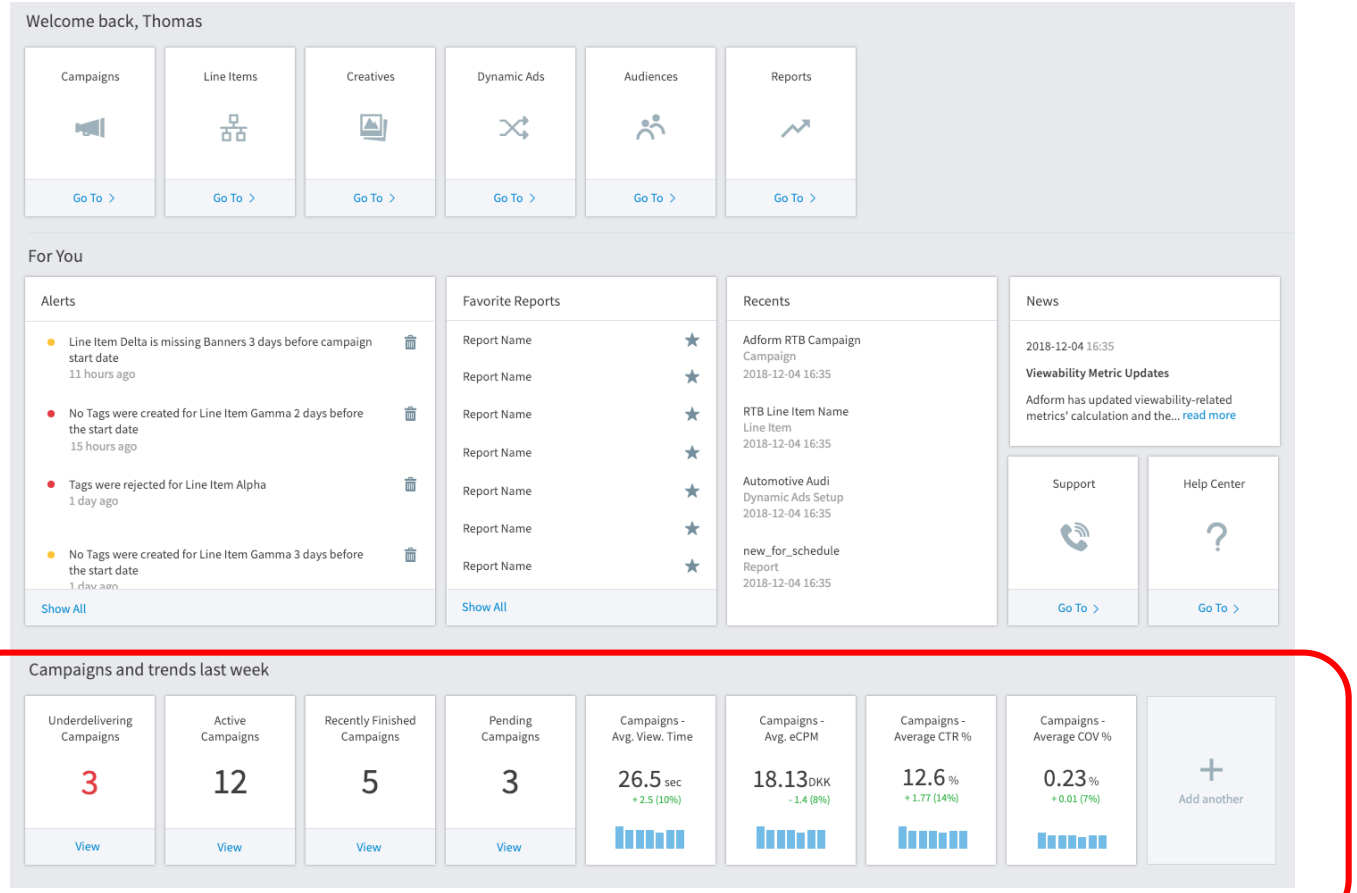

Ad events

Represent the **value** of the campaign a.k.a **Money**

- Measure the **effectiveness** of the **campaign**.
- **Insights** about the consumer - **Data Science**

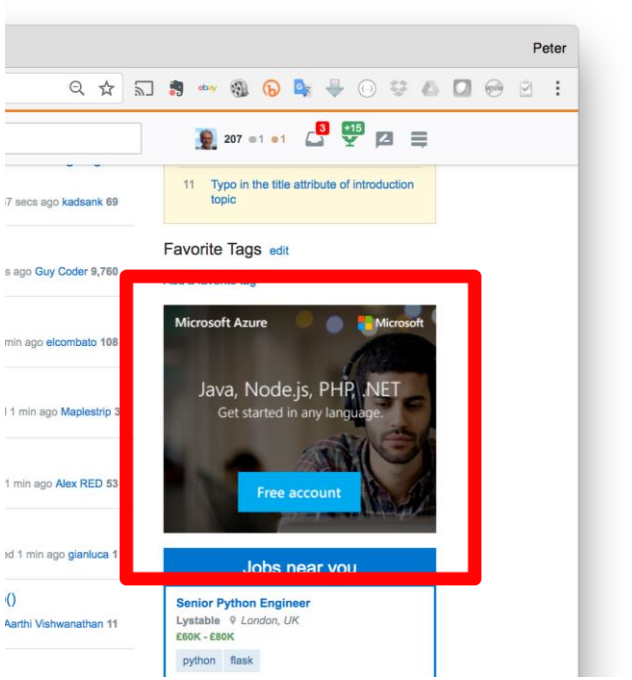

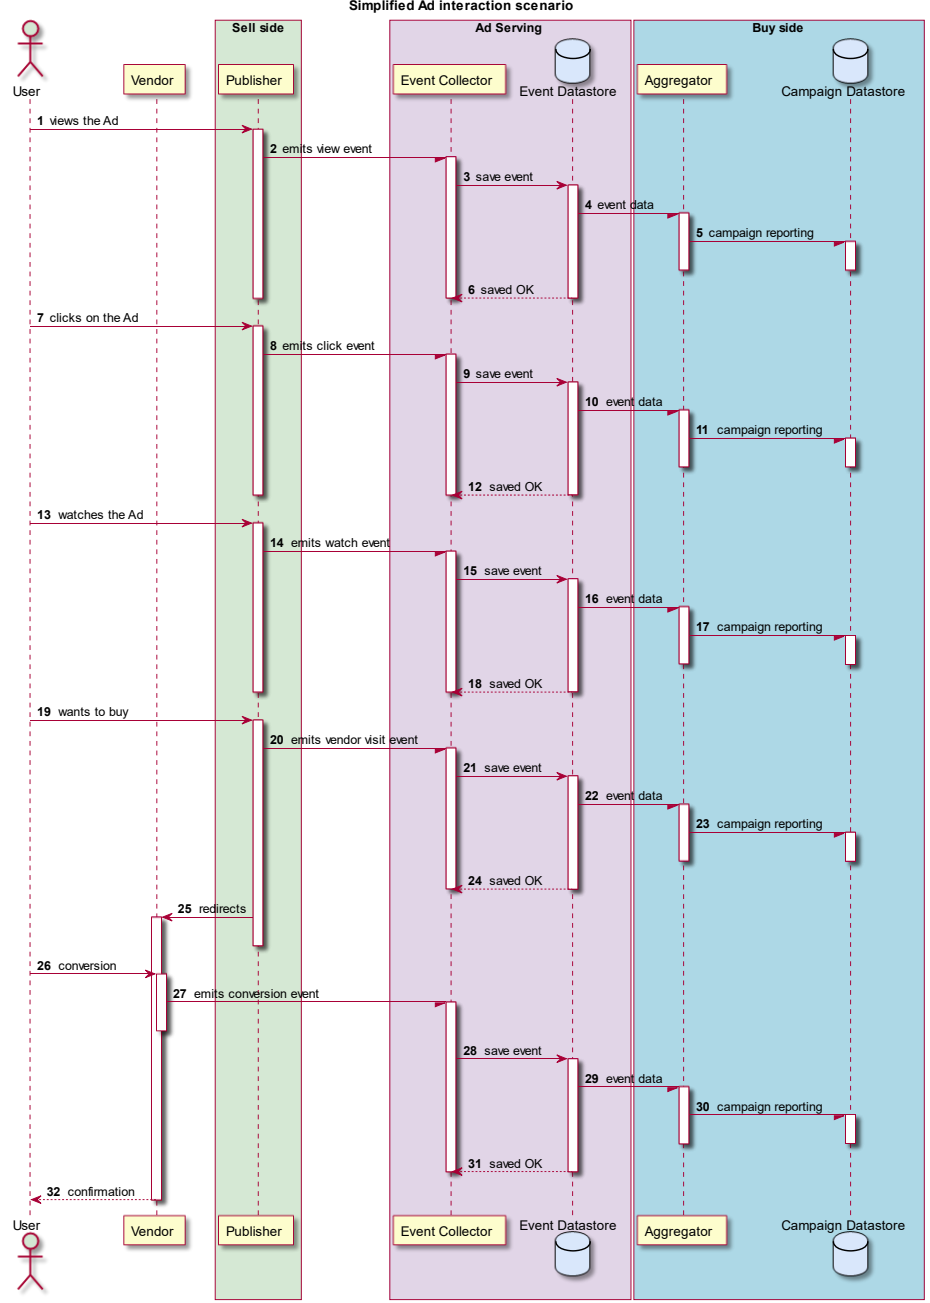

Event Collection

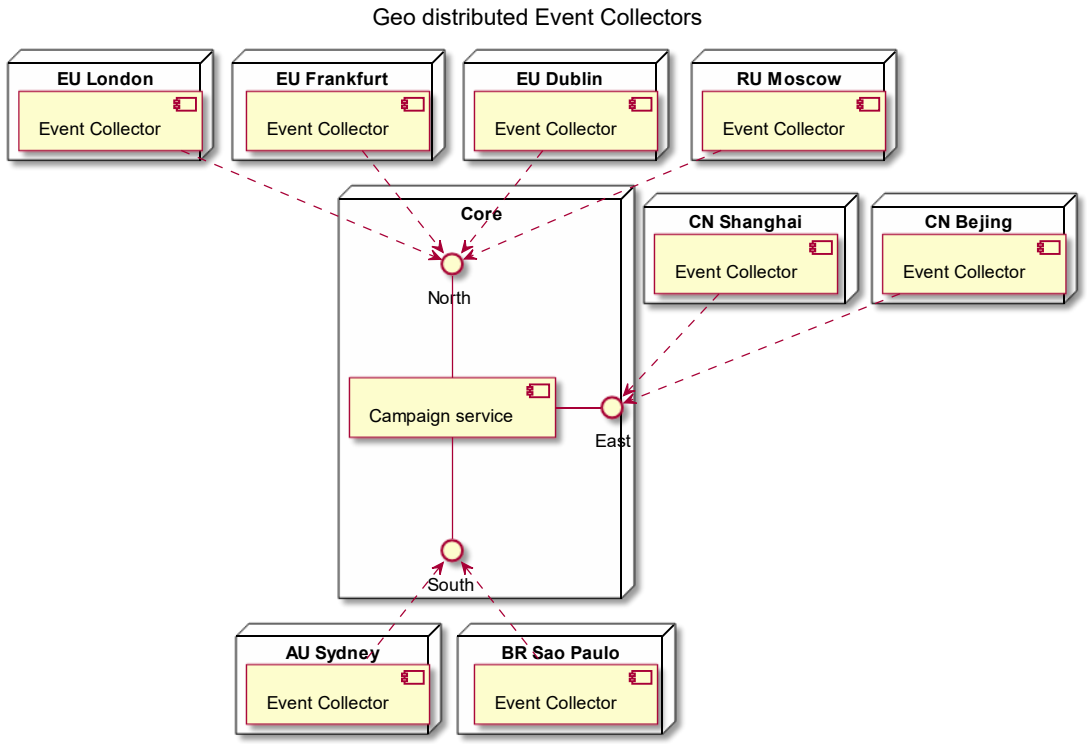

Event collector

- Super **simple web service** API
- **Collects a portion** of events per Geo location
- **Feeds to** the Campaign service
- Linear and elastic scalable
- No state

Event Collection with Ad Event datastore

- Datastore acts as buffer
- Lower latency for event collector instance
  - 1-2 ms
- Fewer Event collector instances
- Adds reliability

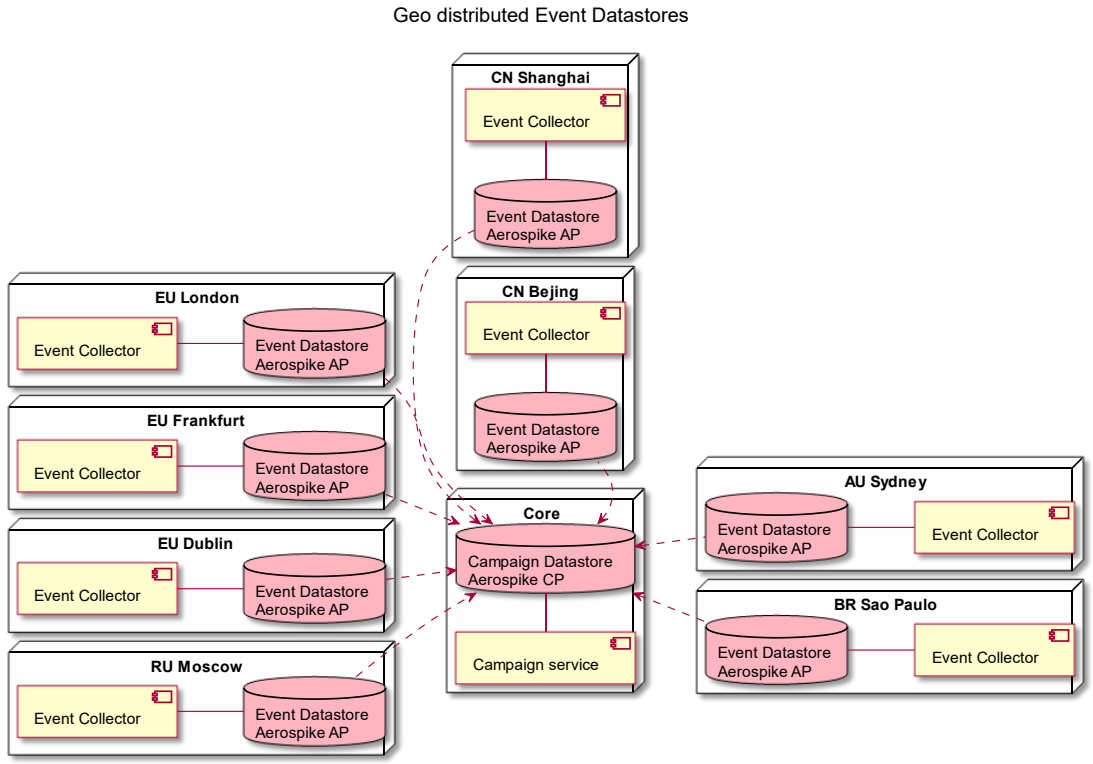

Real-time Aggregation and Reduction

Event collector

- **Receives** events
- Writes to the **local datastore** with
  - Tag id
  - Event type (click, etc)
  - Timestamp

Aerospike Kafka Connector

- **Publishes** events to **Kafka**

Aggregator/Reducer

- Receives **event messages** from the **Kafka** topic
- Maps the **Tag to** the **Campaign**/Execution Plan
- Atomically Increments **counters** (etc) in the **data cube** of the Campaign/Execution plan
- Uses **Aerospike CDT** operations

Ad Event Aggregation and Reduction

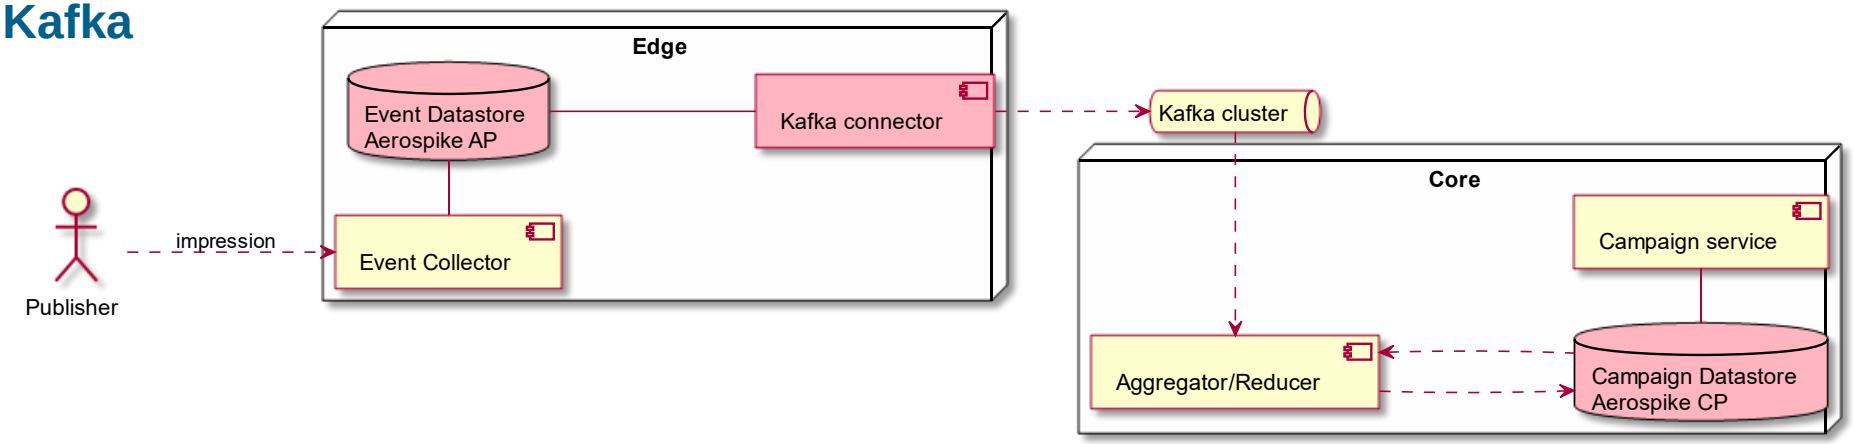

Aerospike Complex Data Type Operations

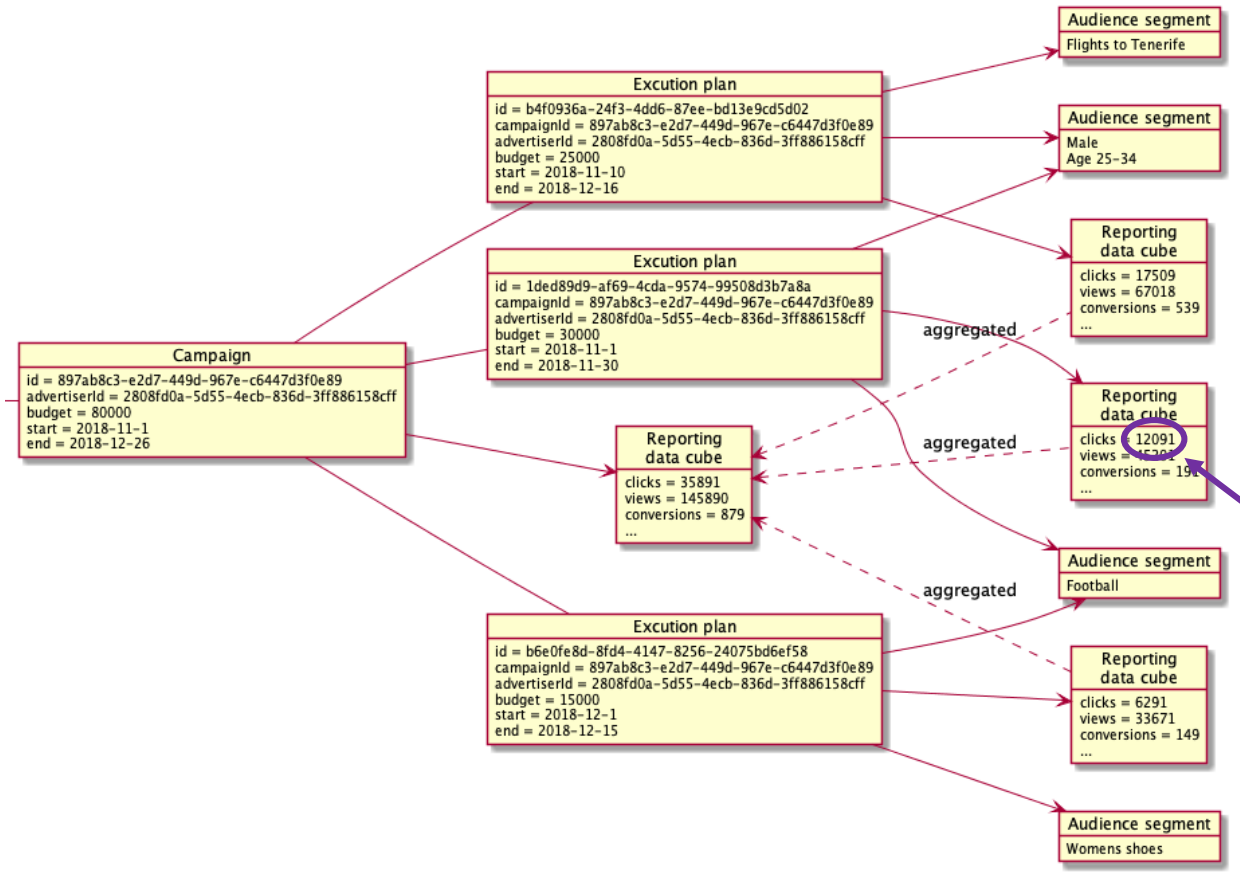

CDT is like a "**document**"

- Atomic **operation** on a **path** to the **data** element
- Latency < 1ms
- Durable

Example

- **Operation**: Increment (and read)
- **Path**:

```
campaign.execution-plan.data-cube.clicks
```

Aerospike datastore summary

User profile datastore

- Size: ~8-12 billion records
- Throughput: ~ 7 million / sec
- Latency: ~0.9 ms
- CAP: Availability, Partitioning
- Uptime: 24/7/365

Events datastore

- Size: 40-80 million records
- Throughput: ~2-3 hundred thousand / sec
- Latency: ~0.9 ms
- CAP: Availability, Partitioning
- Uptime: 24/7/365

Campaign datastore

- Size: 40-80 million records
- Throughput: ~2-3 hundred thousand / sec
- Latency: ~0.9 ms
- CAP: Consistency, Partitioning

Note: Actual sizes vary seasonally

Uncle Pete’s Advice

Why Aerospike

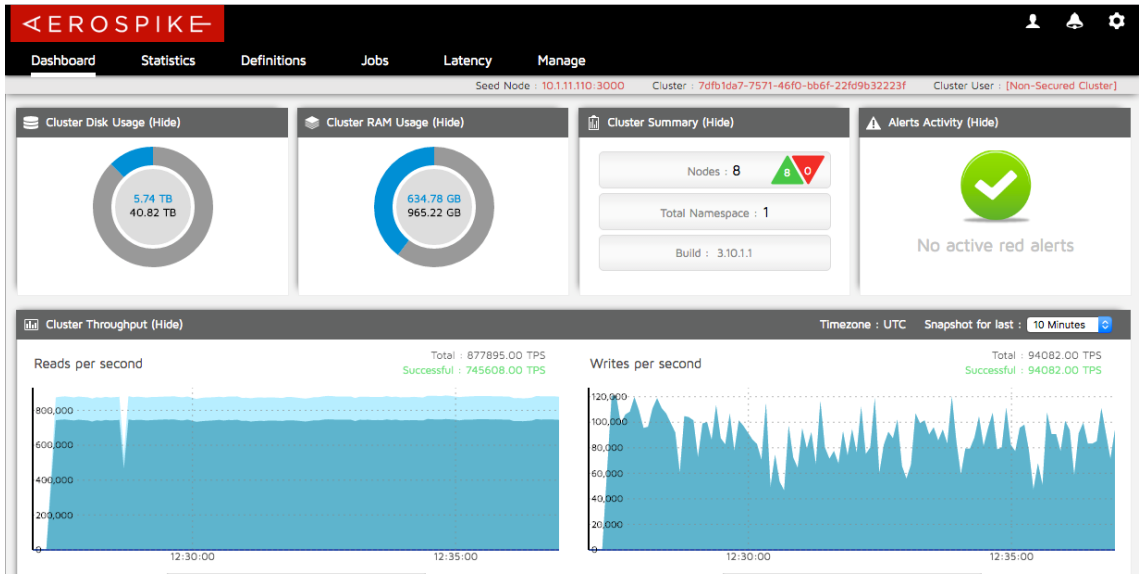

Secret weapon: **Aerospike**

- **8-9 million** reads/s, **1-2 million** writes
  (Over all clusters)
- **< 1ms latency** to read profile
- **~9 billion** cookies in 12 servers
  (largest cluster)
- **HA** with redundancy
- Moved **from Cassandra 5** years ago
  - **Easy** to use
  - **Works** as advertised
  - **Value** for money

  ~150 servers → **12** (largest cluster)

Open source & Enterprise products

Open source != Free

- Modern **escrow**
- Fork the **source repository**
- **Build** it yourself
- **Survive** on your own

Enterprise is value for money

- Support
- Training
- Consulting
- Contractual obligation - **A neck to choke**

adform

Questions?