DATA+AI SUMMIT
BY databricks

From Snowflake To Enterprise-Scale Apache Spark™

Nic Jansma + Amir Skovronik

Akamai

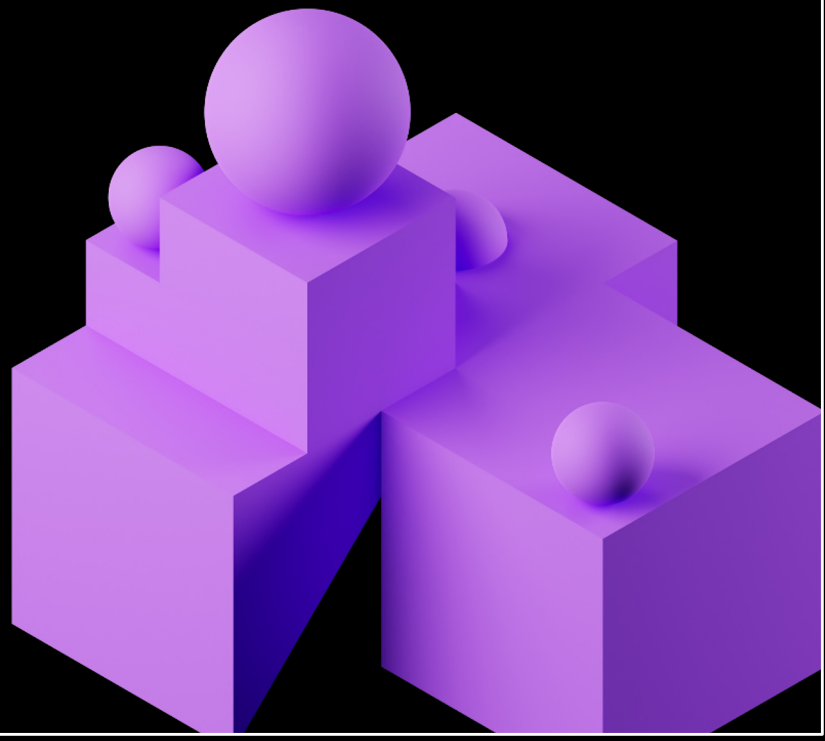

Databricks
2023

From Snowflake To Enterprise-Scale Apache Spark

Nic Jansma
njansma@akamai.com
Sr. Principal Lead Engineer (mPulse)

Amir Skovronik
askovron@akamai.com
Distinguished Engineer (Asgard)

mPulse: Real User Monitoring

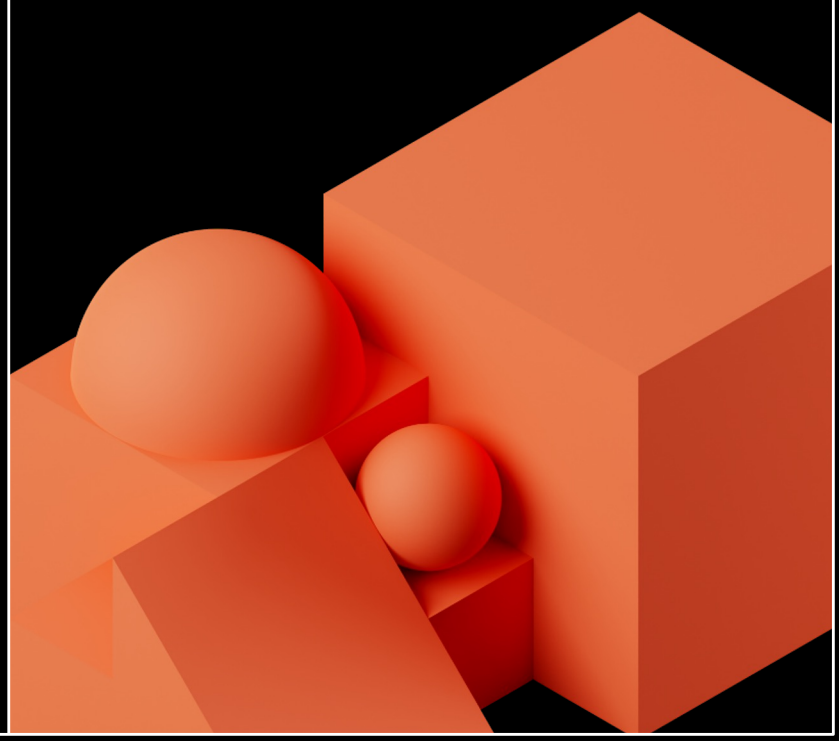

What is mPulse?

Real User Monitoring (RUM)

mPulse provides real-time user experience and performance analytics, and maps those results to business goals and outcomes.

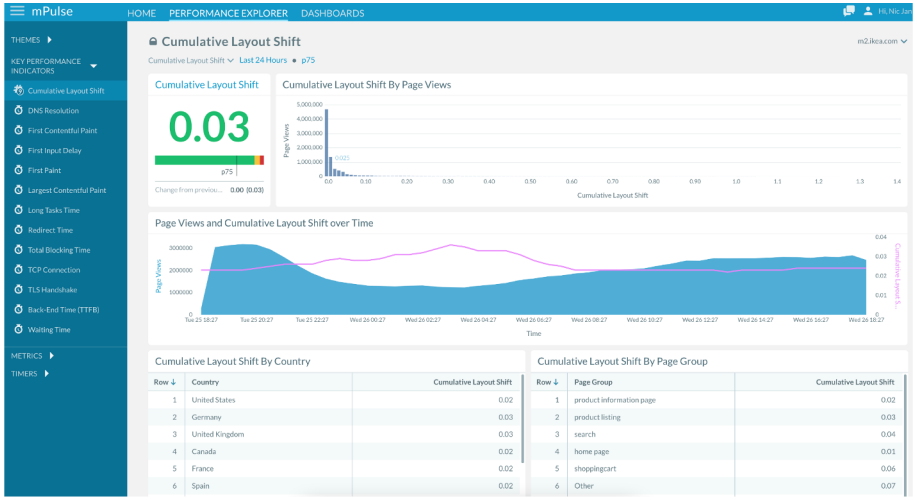

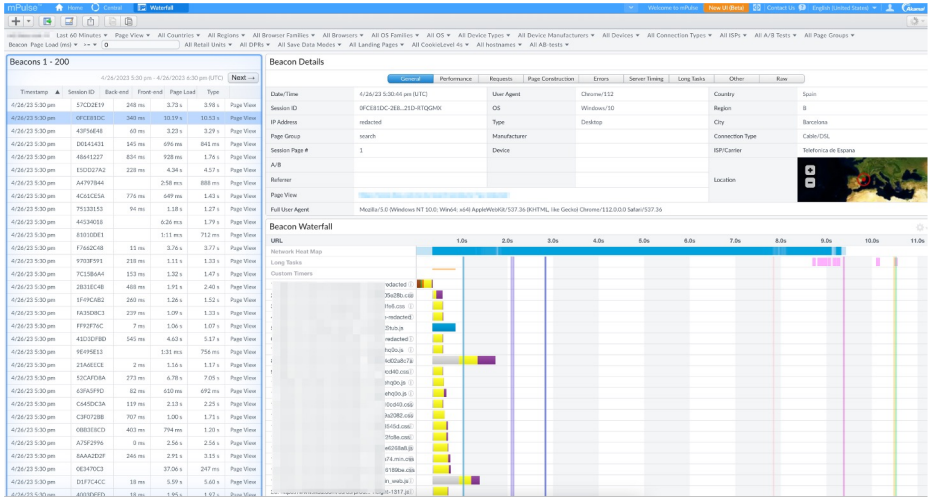

Scale

- > 2 billion beacons / day (no sampling!)
- Real-Time (aggregate) dashboards: User experiences are reflected within **5-10s**
  - 7 TB raw data / day
- Waterfall (individual) dashboards: Full debug trace of **every** page load + beacon available within **5 minutes**
  - 4 TB raw logs / day

Scale

- 13 months retention
- 50 fact / dimension tables
  - >1 T rows
  - >1 PB storage
- 60 QPS

Goals of Migration

- Early Snowflake adopter but needs have changed
- **Highest cloud cost** for mPulse (>$10m / year)
- New Akamai internal team (Asgard) dedicated to providing a data warehouse solution for all of Akamai
  - mPulse was to be one of the first large products to transition to Asgard
- Unique technical challenges
- Equal-or-better performance
- Customers shouldn't notice a difference

Challenges

- Years of assumptions built into mPulse from Snowflake dependency
  - Snowflake made it easy to “throw $ at the problem” by just up-sizing warehouses so we never focused on optimization
- Needed a comprehensive query inventory, and discussions and plans for how to transition each workload
- Other internal teams depend on mPulse data, and they need their own migration paths and hand-holding
- New tooling needs
- Organizationally, two sibling teams (mPulse, Asgard) needed to figure out how to work together and support each other

Asgard: Enterprise-Scale Apache Spark

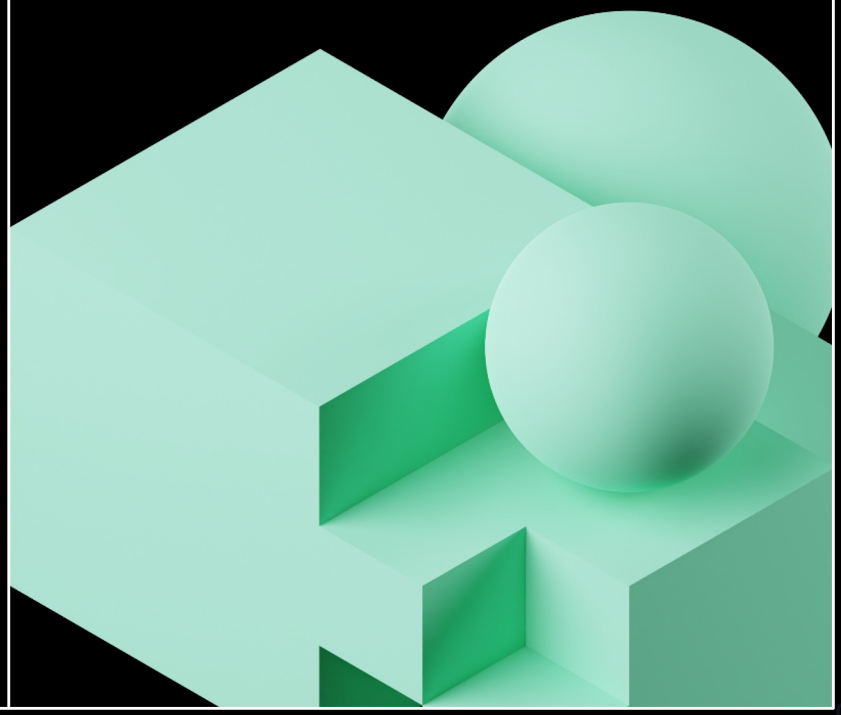

What is Asgard?

- An homegrown cloud based Data Warehouse
- Snowflake like deployment model (S/M/L/XL WH)
- Snowflake like ingest API (COPY INTO)
- Spark SQL query API
- Spark SQL API for ETL execution
- Infrastructure
  - Compute: AKS
  - Storage: Azure Gen2

Asgard Secret Sauce

- Customized & enhanced Spark version
- Unique partitioning model
- State of the art columnar format

Customized & Enhanced Spark Version

- Internal code optimization
  - Optimize synchronized blocks, data structures & SQL functions
  - Improved cached Data frame in memory compression
  - Improved driver stability protection
- Custom Strategy, Rules & Filters in order to enable better push down capabilities
- Cached data Locality awareness
- AZ locality awareness

Unique Partitioning Model

- Inspired by delta lake.
- Internally a file is split according to table partition keys.
- A file footer points to the start/end offsets of each partition.
- Meta service: A custom service which exposes per table metadata on all files in a SQL queryable format.
  - Metadata includes: column value ranges, size, cached location, etc...
  - The data is stored in an in memory database which able to serve X10 q/s at < 200ms.
  - The service can be scale out easily to accommodate with high query rate.

State of the art columnar format

- Code name: Padawan
- Extend pushdown predicates capabilities (Regex, UDF etc …)
- Extend pushdown aggregation capabilities.
- Support delete pushdown.
- Support “explode” pushdown.
- Support optimized data encryption

In local benchmark compared to Parquet:

10 - 15% storage footprint reduction.
Same write time.
20 - 80% improved query time.

Migration Results

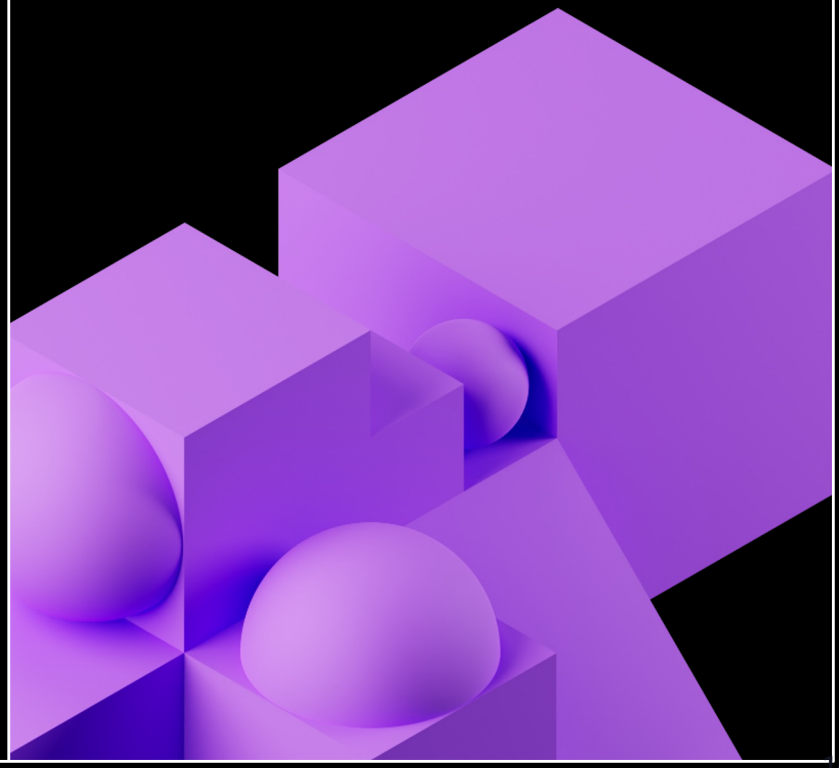

$10m / year savings (80% cost reduction)

Better Performance (20-80%)

Adaptable for Our Needs

The fine print...

Migrations have a cost:

- Years to complete
- Developer fatigue
- Opportunity cost vs. other features

Whats next?

- mPulse
  - Further infrastructure optimizations
  - Migration to Akamai Cloud - Linode
  - Opportunities to build new features

- Asgard
  - Migration to Akamai Cloud - Linode
  - Auto scaling & self management WH
  - Spark jobs on demand
  - Research tools

Thanks!

(Q&A)