KalDB: Log analytics with DuckDB

Suman Karumuri
DuckCon San Francisco
June 2023

About me

Principal Engineer at Airbnb.

I build large scale Observability systems.

Motivation

A tale of 2 services

- Service A: Structured logs + analytics queries.
- Service B: Semi-structured logs + Search.

But, one logging solution doesn’t fit all use cases.

In practice, we use two separate systems

- Data fragmentation.
- Not easy of use: different query languages and UI.

Log Analytics vs Log Search

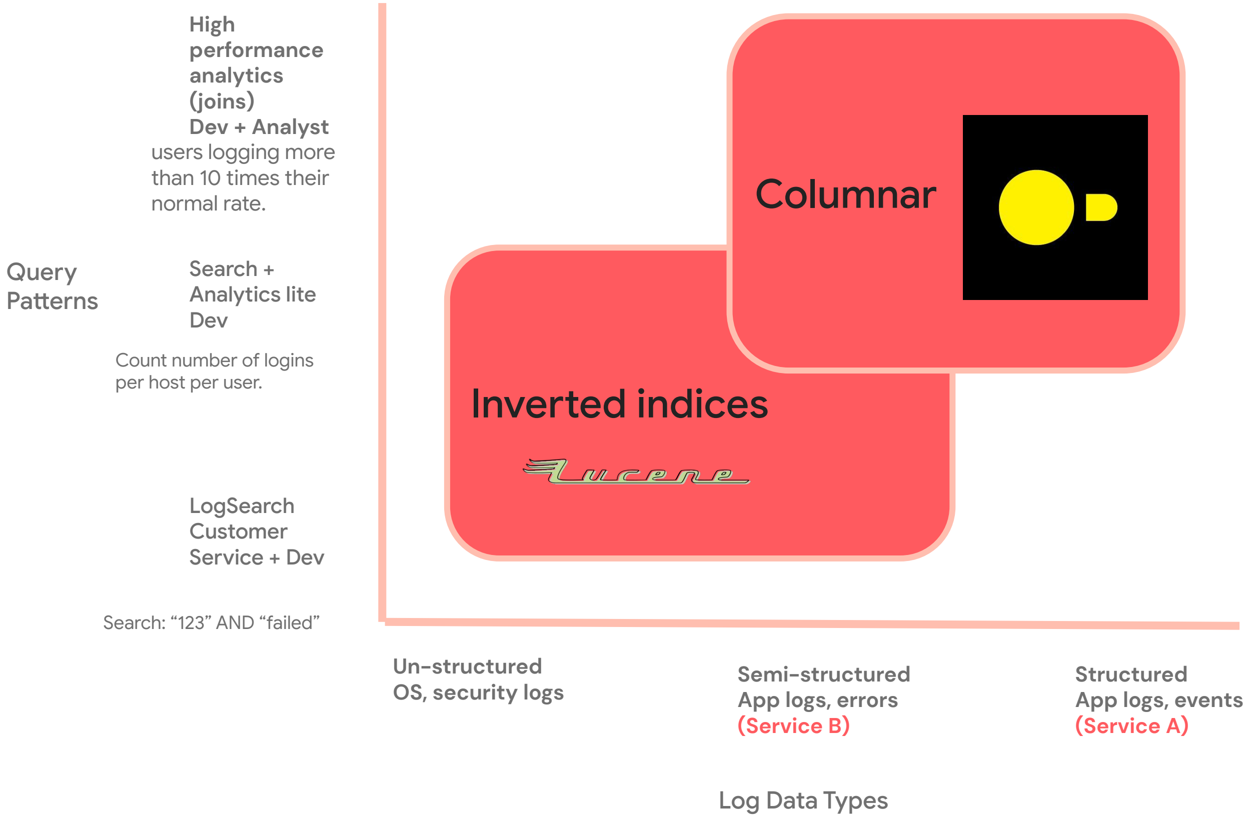

Query Patterns

High performance analytics (joins)
Dev + Analyst

users logging more than 10 times their normal rate.

Search +
Analytics lite
Dev

Count number of logins per host per user.

LogSearch
Customer Service + Dev

Search: “123” AND “failed”

Columnar

Inverted indices

Un-structured
OS, security logs

Semi-structured
App logs, errors
(Service B)

Structured
App logs, events
(Service A)

Log Data Types

Use one system for both log search and log analytics?

KalDB polystore.

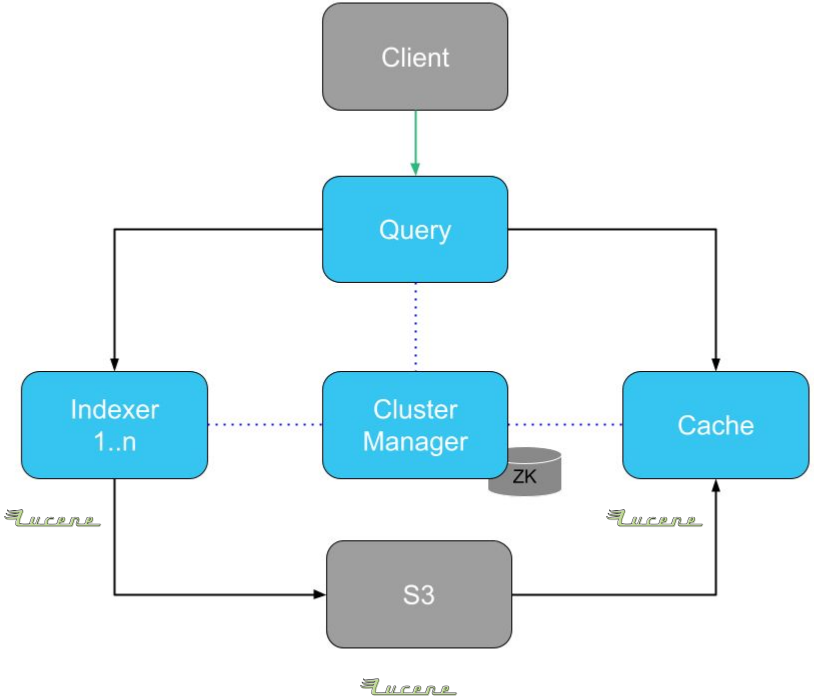

KalDB

KalDB is the ONLY lucene based cloud native observability database.

Disaggregated storage architecture.
Low operational overhead and k8s native.
Open Source

Drop in replacement for OpenSearch.

Designed for PB scale workloads.

Faster and up to 10x cheaper than OpenSearch.

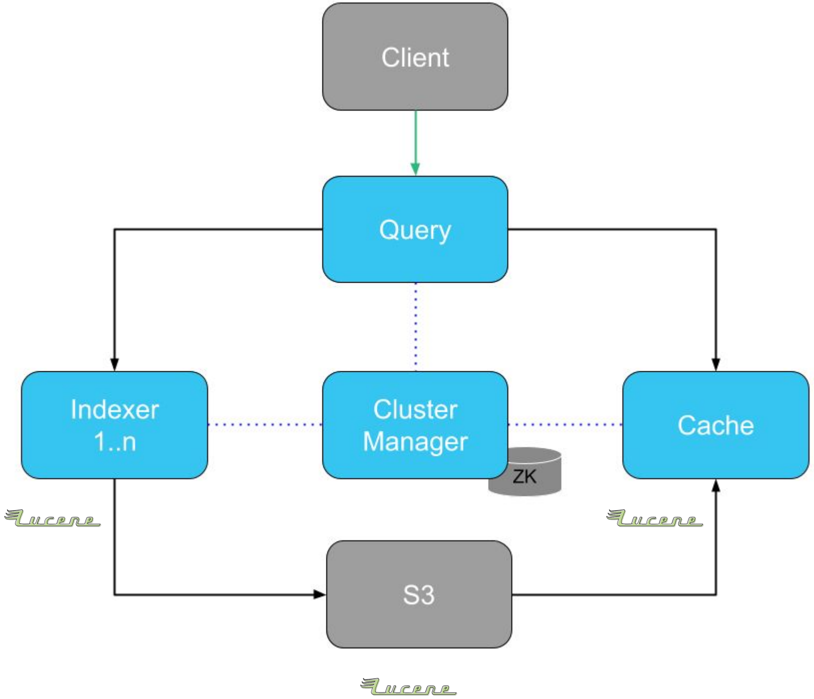

KalDB in production (Slack, Sep 2022)

~1PB of data. 7 days of retention.
70MBps peak ingest.
7M msgs/sec.

Link to talk

Great for Search, Not so great and expensive for analytics.

KalDB Polystore

One system for log search and analytics.

Pluggable storage engine.

Lucene or Duckdb.

Same query language and UI for both use cases.

Disaggregated storage architecture for duckdb.

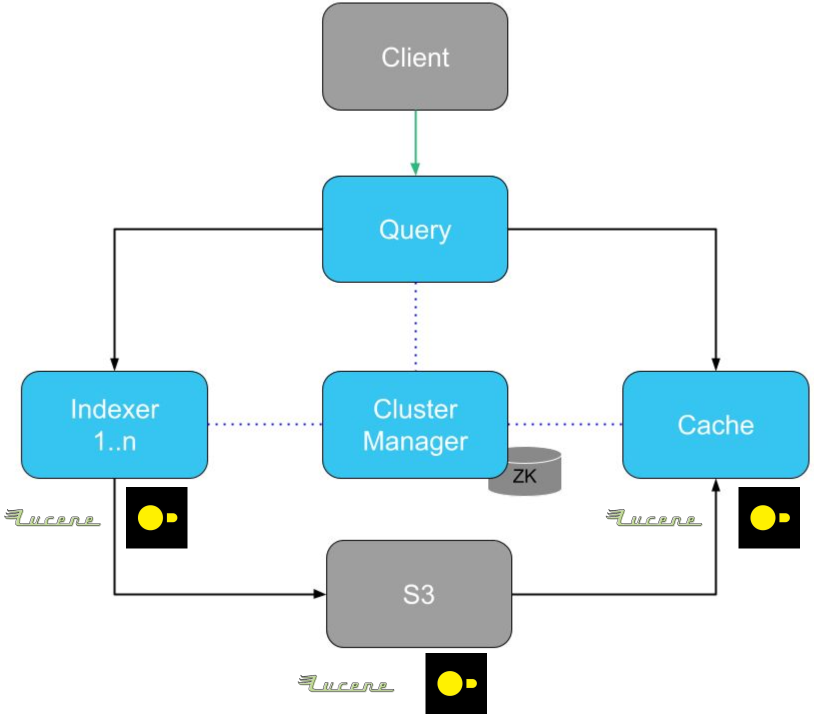

Challenges integrating DuckDB with KalDB

Merging Lucene semantics with Duckdb.

Schema-less ingestion.

Duckdb’s ingestion throughput exponentially falls with increased number of columns.

Implementing distributed scatter-gather for queries using Lucene aggregators.

Help needed!

Thank you!

All product names, logos, and brands are property of their respective owners. All company, product and service names used in this presentation are for identification purposes only. Use of these names, logos, and brands does not imply endorsement.