

Composing Data Systems At Datadog



Alex Bianchi

Software Engineer, Cross-Product Queries, Datadog



Wendell Smith

Staff Engineer, Cross-Product Queries, Datadog



DATADOG

Agenda

01 What is Datadog?

02 Why are we Interested in Datafusion?

03 What is Datafusion?


04 How are we using Datafusion?



- Watchdog
- Events
- Dashboards
- Infrastructure
- Monitors
- Metrics
- Integrations
- APM
- Notebooks
- Logs
- Security
- UX Monitoring

Search... Saved Views \$Role \$env staging \$az

Service Level Agreements




Shop.ist Checkout Latency




Worst 4 groups sorted by 7d status

Past 7 Days



aws:ap-northeast-1	99.99%
aws:eu-west-2	100.00%



Business KPIs





RUM - Frontend Metrics



Top URLs

@VIEW_URL_DETAILS.PA
/
/cart
/department/sofas
/department/so
/checkout

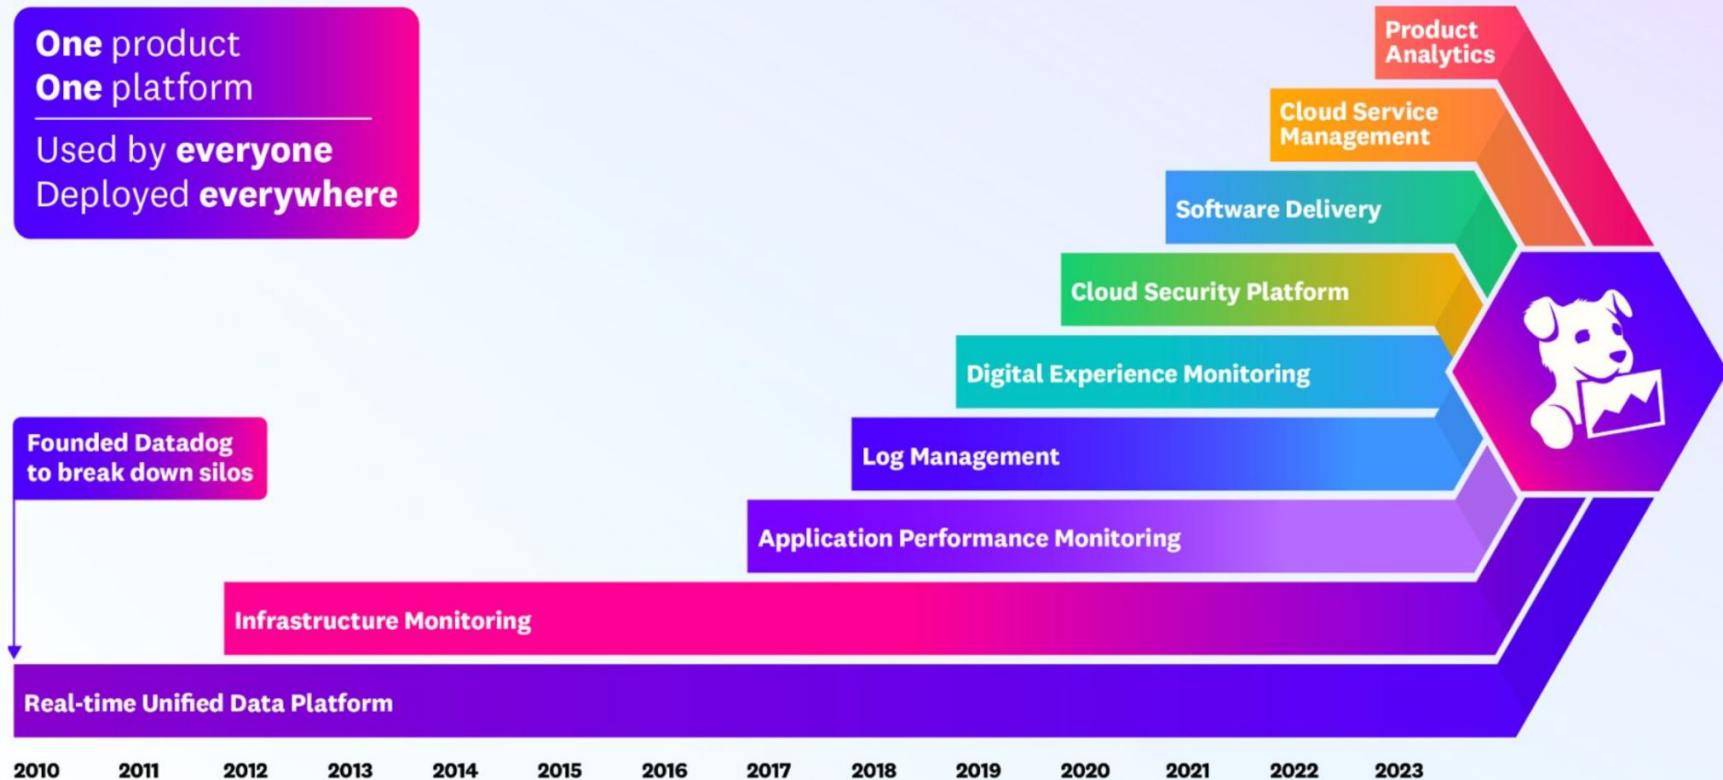
Performance



All of your monitoring and security tools in a single, unified platform

Visit datadoghq.com

Datadog Through The Years



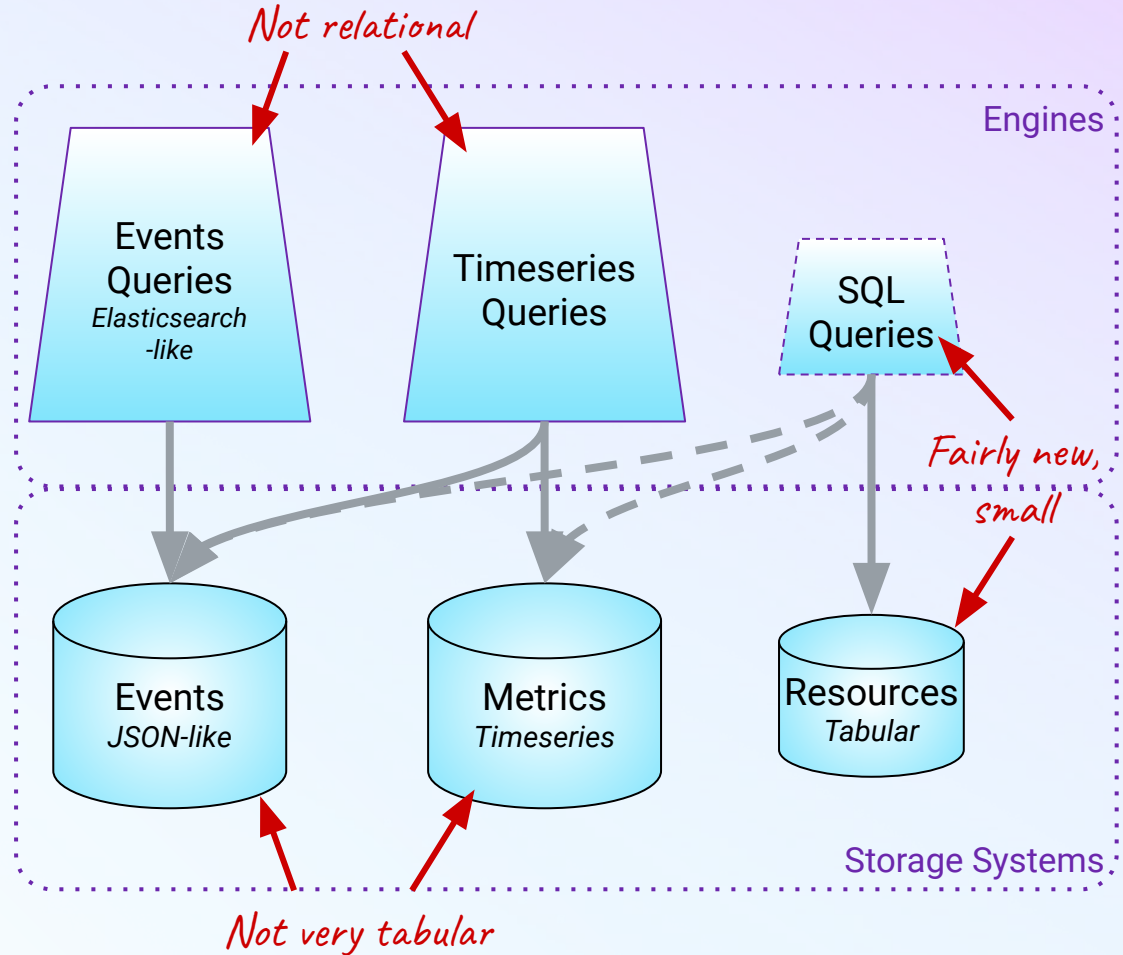
Our Current Tech Stack

Challenge

We have multiple non-relational engines from years of growth, acquisitions, and optimizations.

Goal

Enable more complex queries from multiple data sources



DDSQL Editor

Single access point

SQL style syntax

Joins across data stores

The screenshot shows the DDSQL Editor interface. At the top, there are several tabs: "Database Version ...", "Agentless Azure vulner...", "extract...", and "DDSQL Findings Funnel (...)". The active tab is "Database Version Query". Below the tabs, there is a sidebar with icons for "Schema", "Queries", and "Docs". The main area displays a SQL query:

```
1 SELECT DISTINCT p.database_name,  
2 d.database_version  
3 FROM postgres_logical_database p  
4 JOIN database_instance d USING (database_instance_key)  
5 ORDER BY d.database_version DESC;
```

Below the query, there is a status bar indicating "SUCCESS 161 rows". The results are displayed in a table with two columns: "database_name" (text) and "database_version" (text).

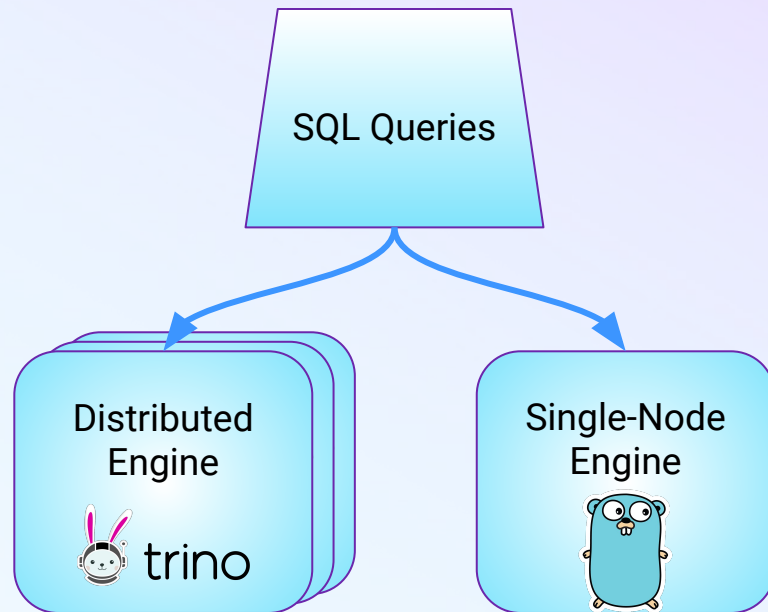
database_name	database_version
juliatian	v14.10.0
kafka_metastore_test	v14.10.0
chat	v14.13.0
kafka_metastore_general	v14.10.0
incidents	v14.13.0
rapid_petshop_go	v14.13.0
dep_versions	v14.10.0
side_scanning	v14.10.0
apm_trace	v14.10.0
atlas	v14.10.0

Cross Product Queries

Today, we use the distributed SQL engine Apache Trino; a tool designed to efficiently query vast amounts of data.

Trino was not designed for high throughput, real-time queries on small datasets

For these queries, we have our own execution engine. But building out complete functionality has had ups and downs.




98%

Of our cross product queries
could be executed in single
node.



**We're looking to Datafusion for
high throughput, low-cost
queries.**

What is



A P A C H E

DATAFUSION™

“DataFusion is an **extensible query engine** written in Rust that uses **Apache Arrow** as its in-memory format. DataFusion’s target users are developers **building fast and feature rich** database and analytic systems, **customized** to particular workloads.”

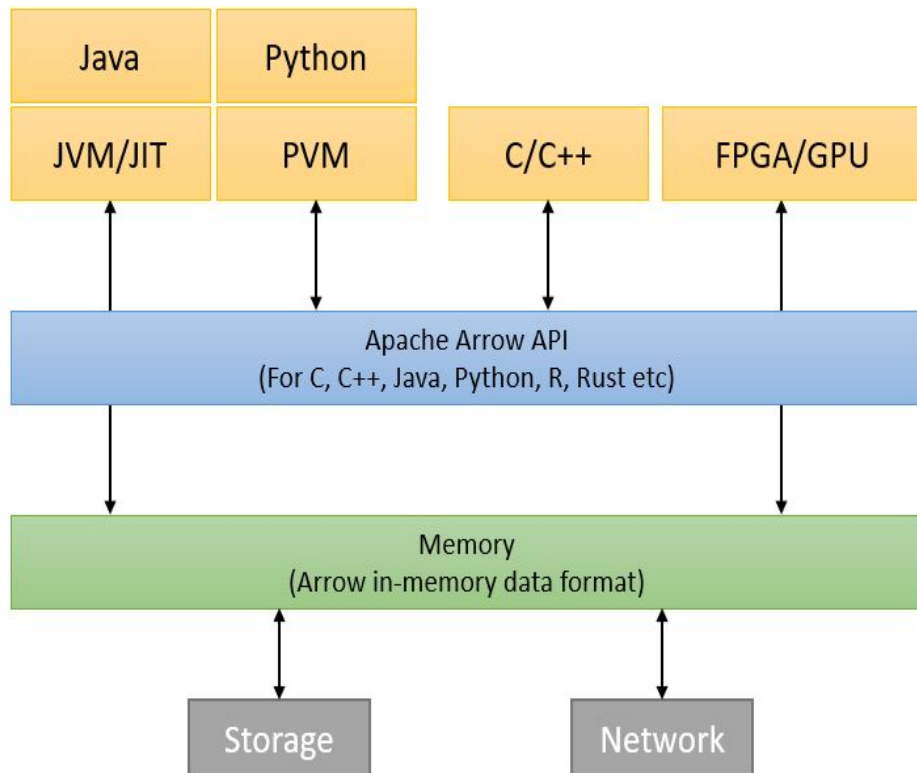
datafusion.apache.org

The website is pretty good, you should take a look!


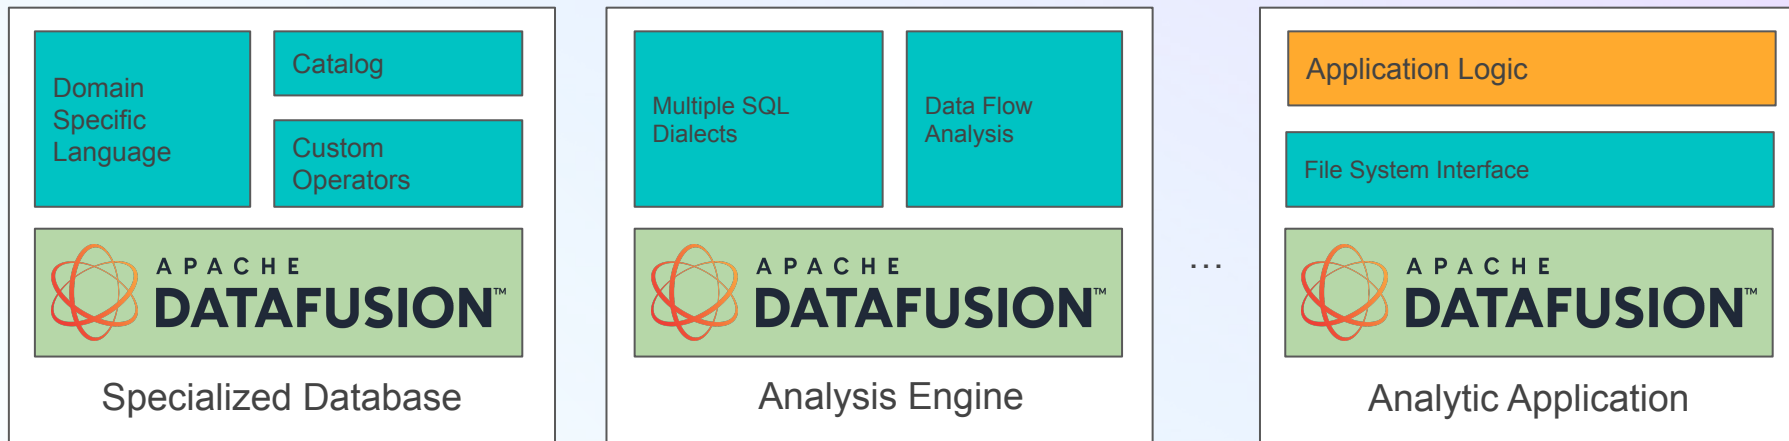
Apache Arrow

- Language-independent
- columnar memory format
- Efficient over the wire communication (Flight)
- Built for efficient analytic operations on modern hardware (SimD)

Libraries are available for **C**, **C++**, **C#**, **Go**, **Java**, **JavaScript**, **Julia**, **MATLAB**, **Python**, **R**, **Ruby**, and **Rust**



Building fast and feature rich database and analytic systems: DataFusion is LLVM for Databases



Query Engine Extensibility

Relational

SQL Operations

DataFrame Operations

Multiple API's

Composable

Parsing

Query-planning

UDF's

Table Providers

Performant

Columnar

Vectorized

Streaming

Multi Threaded

No Garbage Collector

Customized to Particular Workloads

Domain-Specific Database Systems

- Time series databases (e.g. InfluxDB 3.0 and Coralogix)
- Streaming SQL platforms (e.g. Synnada and Arroyo).

Run-times for specialized query front-ends

- Comet for Apache Spark
- Seafoal for PostgreSQL
- Vega
- InfluxQL

SQL analysis tools

- Dask-sql
- SDF

Table formats

- Rust implementations of...
- Delta Lake
 - Apache Iceberg
 - Lance

Today

Segmented Query Stack

Proprietary memory layouts
increase conversion overhead

Different execution engines lead
to inconsistent user experiences

Multiple query representations
limits interoperability

Goal

Composable Data Systems

Apache Arrow minimizes
conversion + marshaling costs

Datafusion's extensibility let it
match user's expected behavior

Substrait can be used across our
downstream data sources

How we are using Datafusion at Datadog

Event Reads

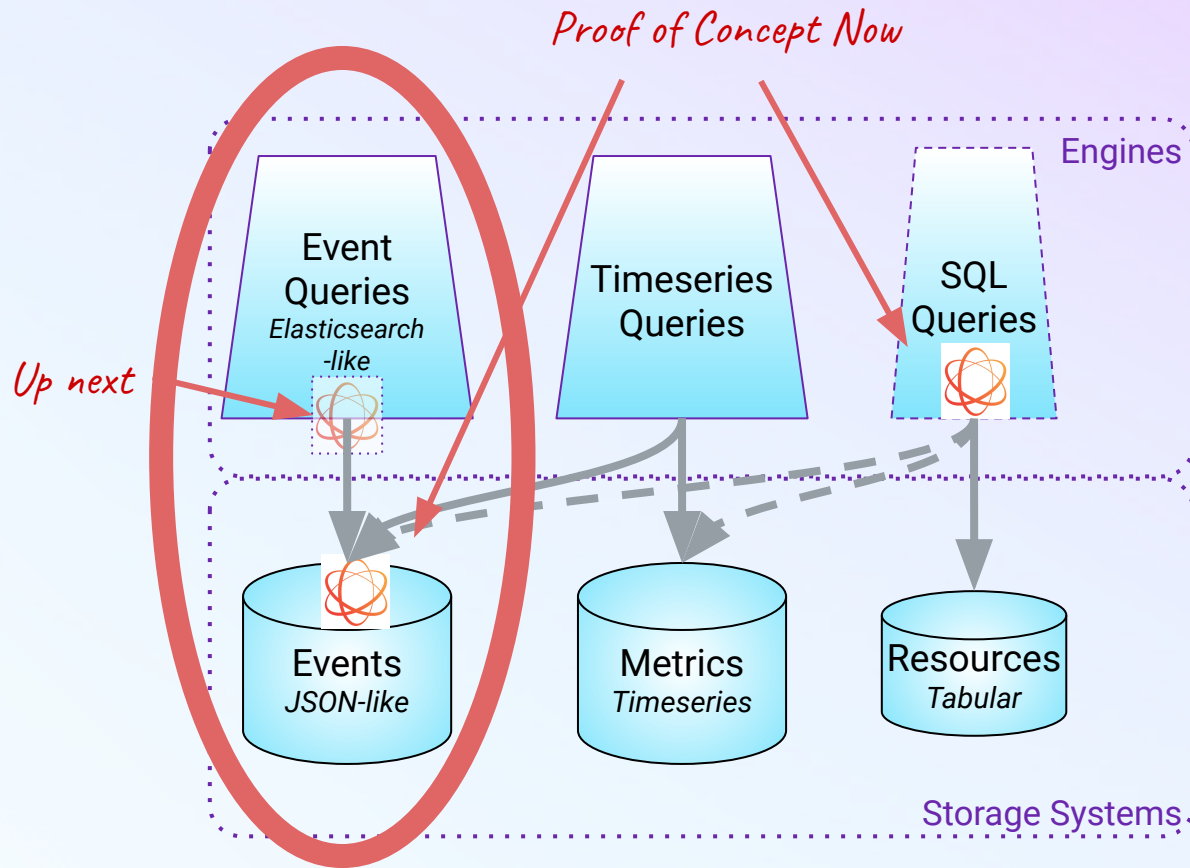
Filter and project on read

Event Queries

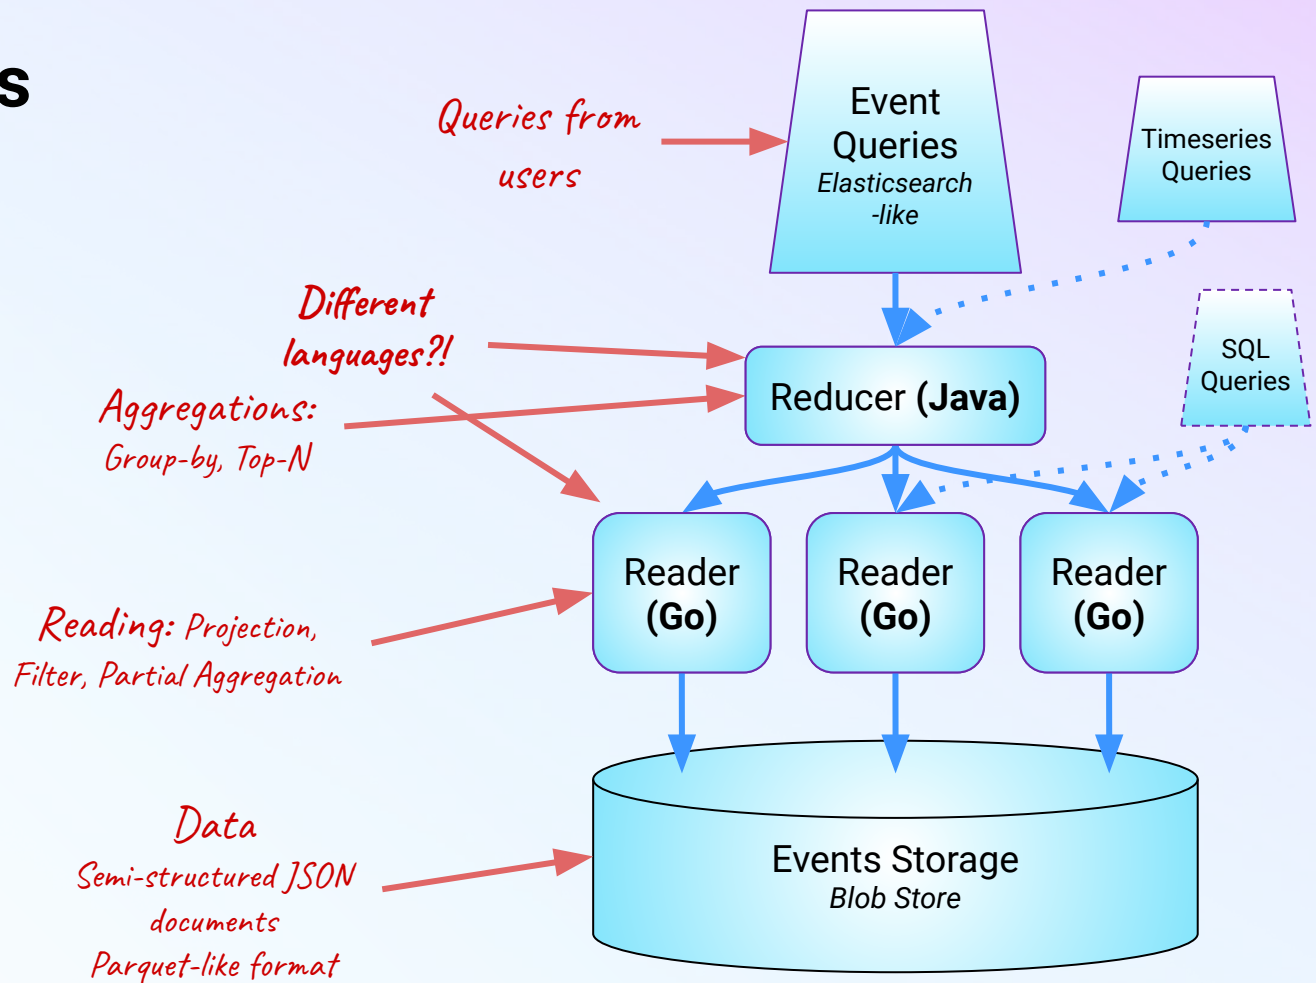
Aggregation, sort, limit

SQL Queries

JOINs, subqueries, CTEs



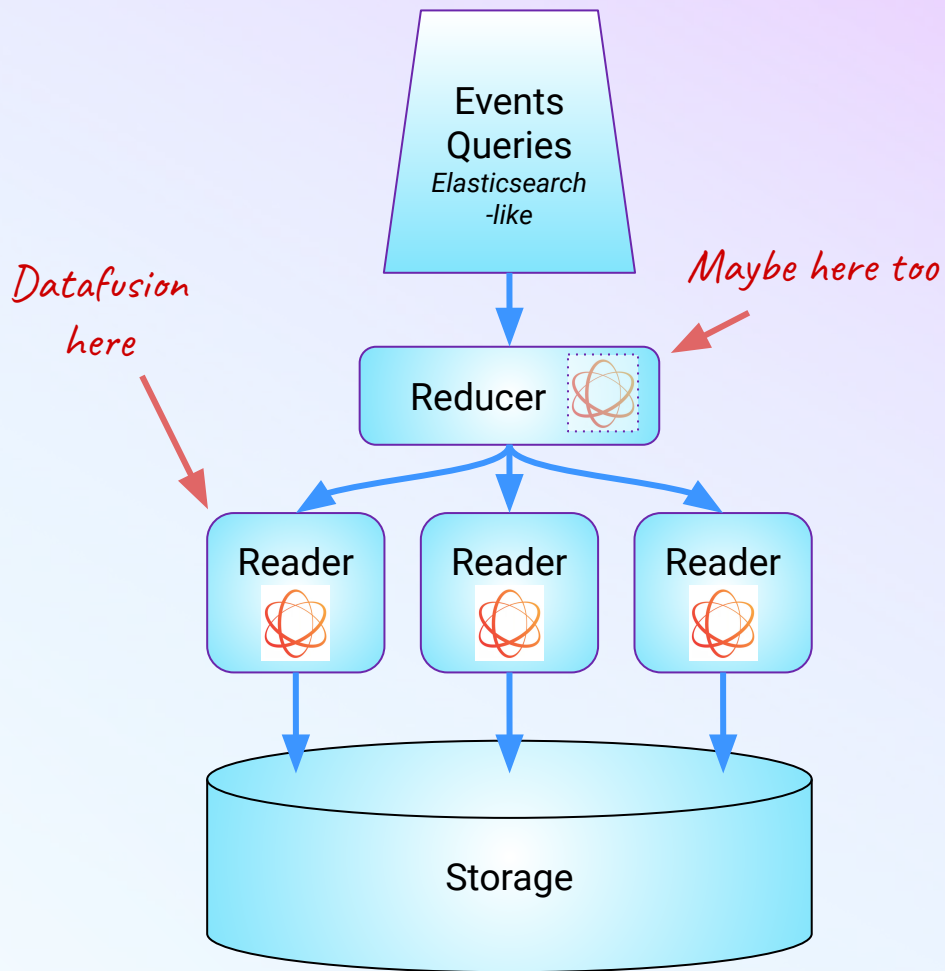
Events Queries



Goal

Datafusion as a shared common engine across services

- Out of the box industry-supported “framework”
- Consolidate contracts and behaviors
 - Single IR, single engine, single format
- Free time to focus on higher level problems
 - Text search, interactive queries, approximate operators, shuffling, etc



Events Queries and Datafusion

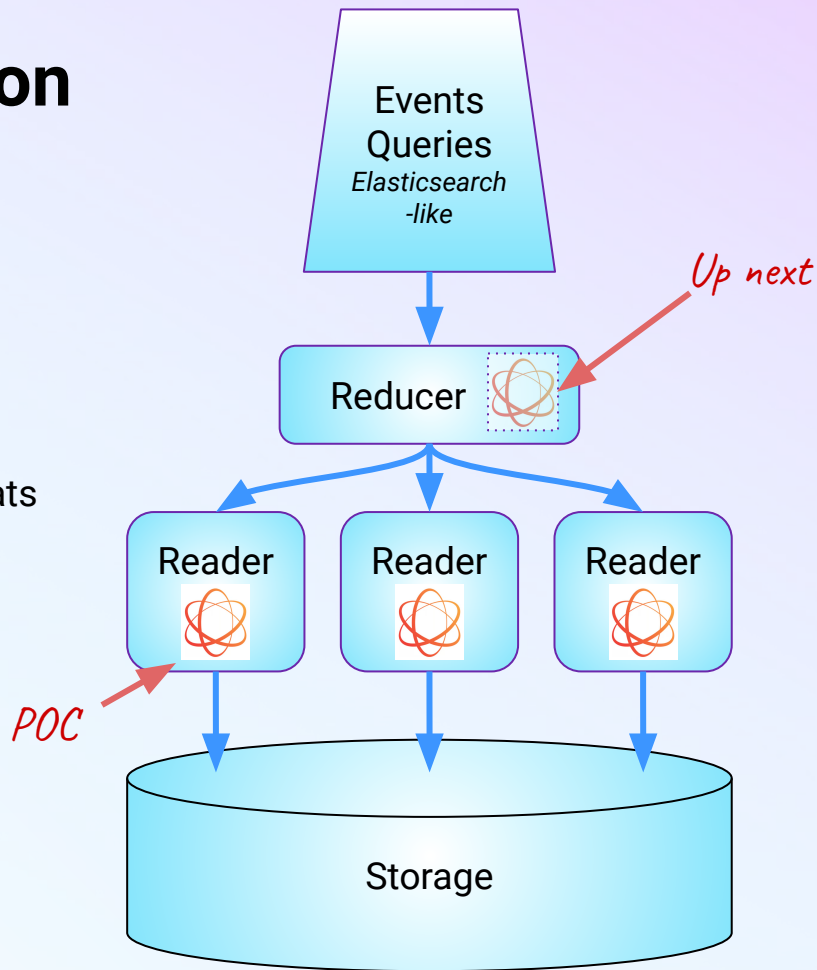
How is it going?

Challenge: Matching Behaviors

- Requires flexibility: extensions to match existing functionality
 - Custom coercions, approximations, exotic operators
- Efficient integration with our various storage formats
 - e.g. late materialization / Arrow translation: not all our encodings are cheaply translatable

Status

- Currently **shadowing** queries on readers
- Seeing **performance improvements** and discovering bugs in our existing engines
- **Contributing back** to community: [Decoupling logical/physical types](#)



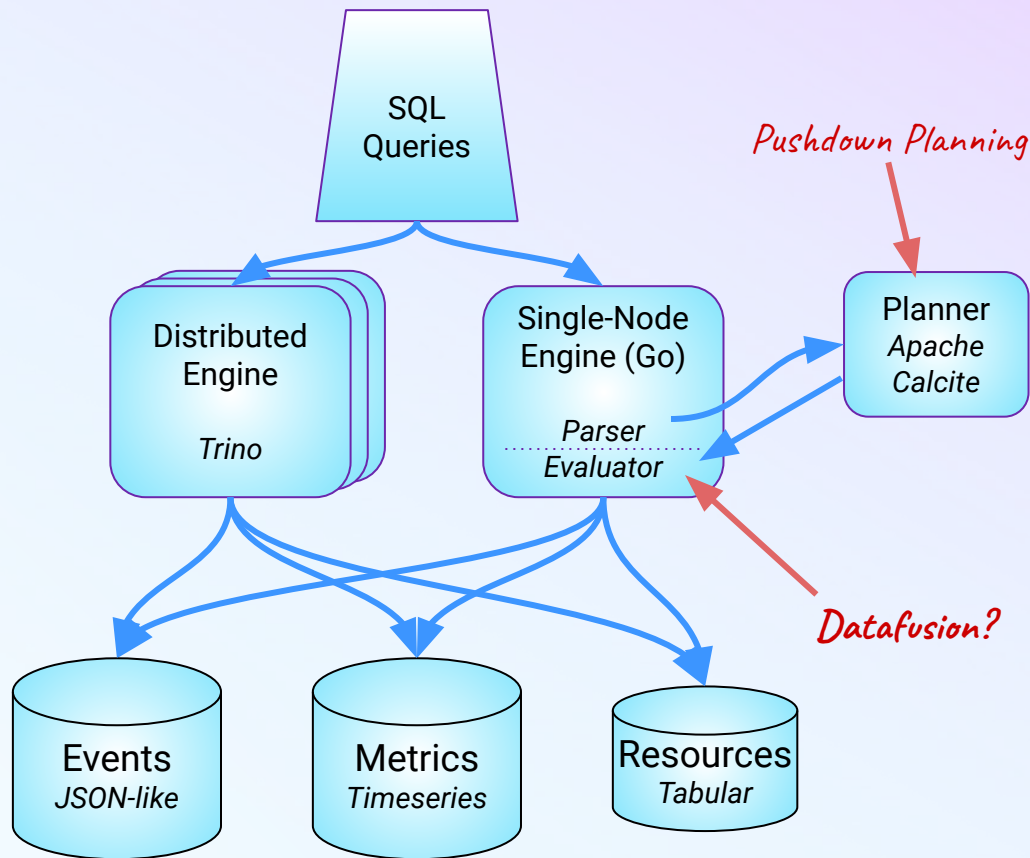
SQL Queries

Distributed Engine

For large operations

Single Node Engine

For high-throughput,
Low-cost queries



SQL Queries and Datafusion

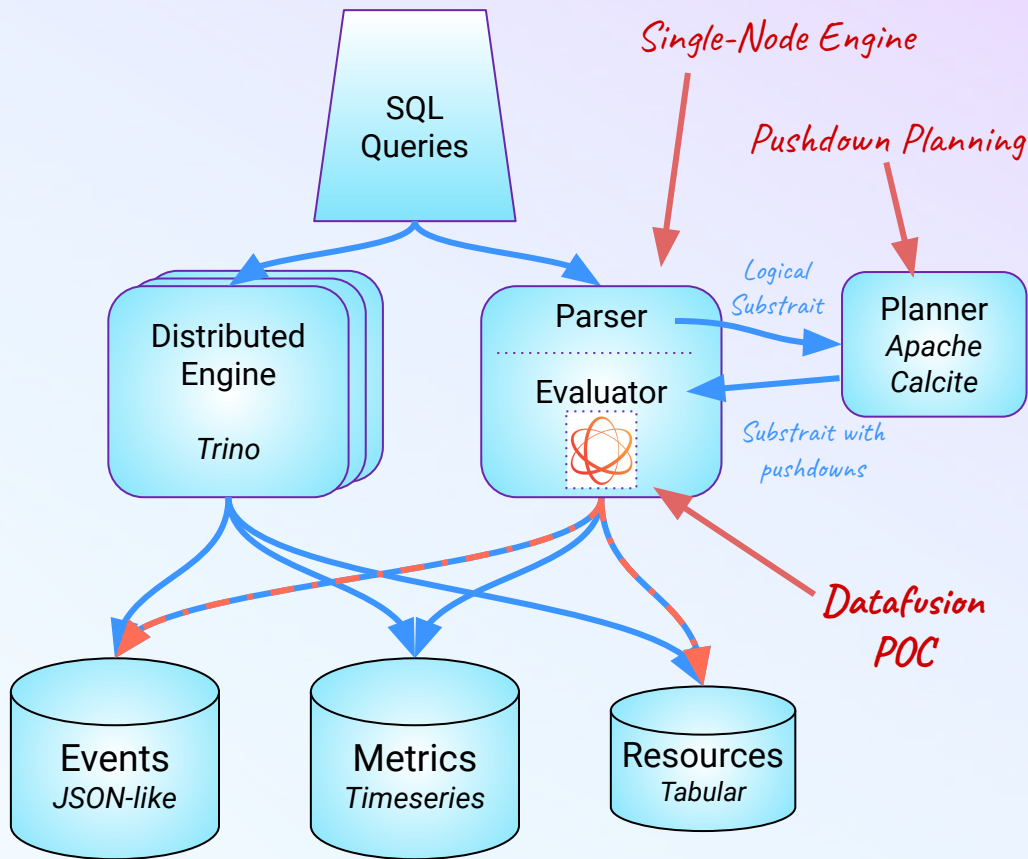
How is it going?

Challenges

- Our own **SQL dialect**
 - ⇒ Separate parser
- Extension **types**
 - String sets (for tags), timeseries
- Efficiency requires **pushdown** operations
 - ⇒ Separate planner
- Different **downstreams** to integrate with
 - ⇒ Substrait ([Contributions](#))

Status

- **Proof of concept**
 - Embedded library in Go service
 - Planner integration via substrait
- **Basic SQL operations** functional
- Events and Resources as **sources** (no Metrics)



Ongoing Work



Integrating and Extending via Substrait

Using Substrait requires many extensions




Bridging Rust to Go

Embedding eases the transition, but FFI work is finicky



Planner Integration

Calcite plans \neq Datafusion plans



Consolidation

Standardizing across our different engines:
Substrait extensions, connectors, function behavior, ...

TODAY

Specialized Engines

with converging functionality

Duplication of effort

(Reimplementation)

Custom IR / Formats

Different protobufs for each source

Behavior discrepancies

Re-inventing commodity operations

VISION

Convergence

of implementation and behavior

Library Reuse

Datafusion for same functions,
different places

Adaptability

Add or replace components, sources

Simplify Contracts

Subtrait for plans

Arrow for data

Thank you!