



# Infrastructure Monitoring



# Agenda

- Challenges
- InfluxDB Overview
- InfluxDB for Infrastructure Monitoring
- Use Cases

# Infrastructure monitoring use cases “speak” time series



2017-11-17 00:00:00.50	Temp: 129.4	Press: 0.50	Rot: 1,004.00	Part: 0
2017-11-17 00:00:01.00	Temp: 123.0	Press: 0.70	Rot: 1,106.00	Part: 0.50
2017-11-17 00:00:01.50	Temp: 121.5	Press: 0.50	Rot: 1,180.00	Part: 0.32
2017-11-17 00:00:02.00	Temp: 125.3	Press: 0.40	Rot: 1,050.00	Part: 0.78
2017-11-17 00:00:02.50	Temp: 126.7	Press: 0.80	Rot: 1,082.00	Part: 0.63
2017-11-17 00:00:03.00	Temp: 124.9	Press: 0.60	Rot: 1,194.00	Part: 0.99
2017-11-17 00:00:03.50	Temp: 127.3	Press: 1.00	Rot: 1,070.00	Part: 0.69
2017-11-17 00:00:04.00	Temp: 121.3	Press: 0.90	Rot: 1,092.00	Part: 0.55
2017-11-17 00:00:04.50	Temp: 124.0	Press: 0.40	Rot: 1,135.00	Part: 0.86
2017-11-17 00:00:05.00	Temp: 127.7	Press: 0.60	Rot: 1,099.00	Part: 0.93
2017-11-17 00:00:05.50	Temp: 128.3	Press: 0.50	Rot: 1,033.00	Part: 0.66
2017-11-17 00:00:06.00	Temp: 129.7	Press: 0.70	Rot: 1,130.00	Part: 0.67
2017-11-17 00:00:06.50	Temp: 120.4	Press: 0.70	Rot: 1,096.00	Part: 0



- Measuring change over time – status, performance, problems
- Events and activity encapsulated by time interval
- Time is often the magnifying glass for operation

# Challenges with managing time series data

## Massive Scale

---

Data is continuously arriving at high speed and volume

## Real Time Action

---

Applications must analyze data within streams and act in real time

## Data Cardinality

---

Higher number of tags collected cause high cardinality impacting performance

---

**Most tools  
(including relational databases)  
simply cannot handle these  
challenges**

# Relational datastores are not a fit for purpose

- **Scale** – Cannot ingest large volumes of data in short time period
- **Slow queries** – Not architected for real-time querying
- **Lifecycle management** – Sharding & data retention are not built-in

---

# What makes InfluxDB Special

# Uniquely equipped for time series data



## Scale

Designed to scale for large volumes of time series data

---

## Distributed

Non-blocking high volume writes and reads

---

## Availability

Write and read availability are prioritized

---

## Management

Data lifecycle management with built-in data retention

---

## Flexible

Schema on write



# InfluxDB 3.0: Columnar database for high performance & low cost



## Real Time

---

**Hot data in memory**

*Sub-second responses  
for recent data*

Optimized for low latency  
**SQL** or InfluxQL queries



## Lowest cost storage

---

**Cold data in object store**

*Superior compression &  
reduced TCO*

Optimized for lowest  
cost long term storage



## Unlimited Cardinality

---

**Optimized writes & reads**

*One datastore for all  
time series data*

Optimized for ingest  
scale & speed

# InfluxDB 3.0: an Open Data Architecture

## Flight

Transport columnar data at high speeds based on Arrow format

## DataFusion

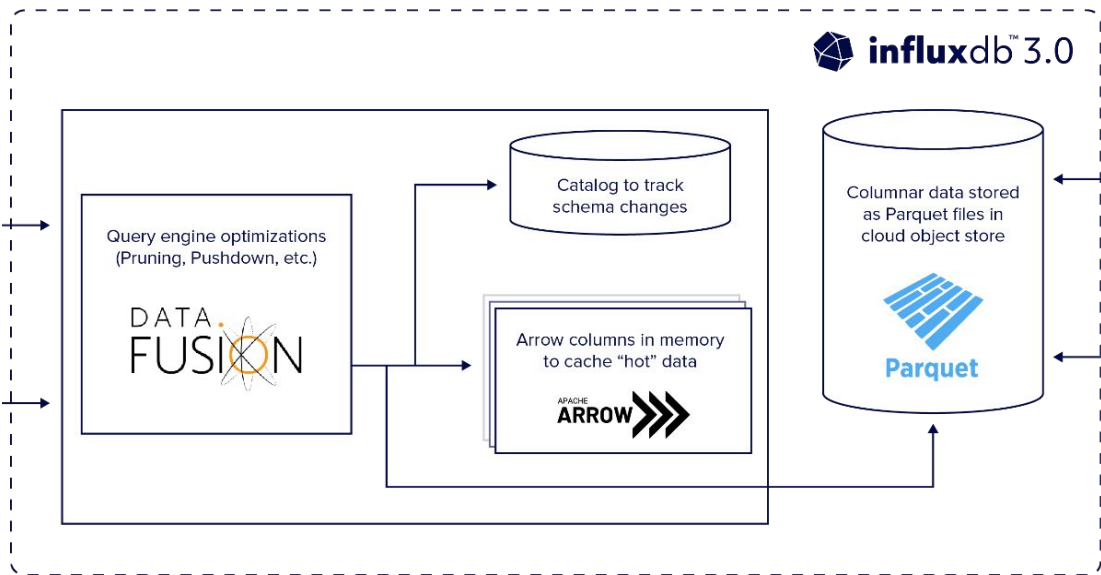
Fast query execution engine written in Rust

## Arrow

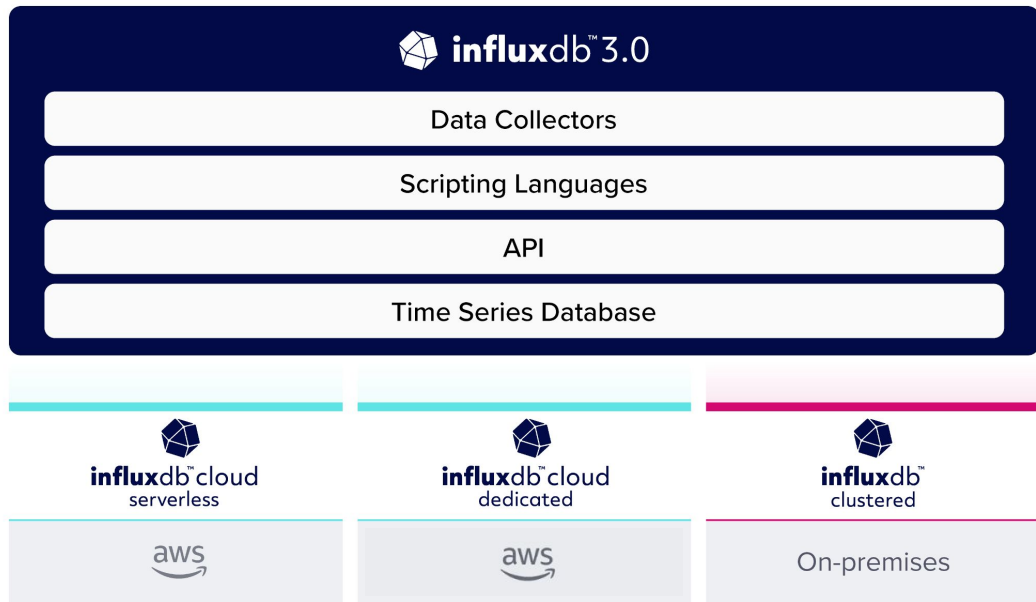
Optimized for running large analytical workloads

## Parquet

Open column-oriented file format designed for efficient data storage and retrieval



# InfluxDB 3.0: Run on cloud & on-premises



## Cloud Serverless

- managed service for small & medium workloads

## Cloud Dedicated

- managed service for large enterprise workloads

## Clustered

- software for large enterprise workloads in self-managed environments

\* We strongly encourage that customers evaluate the InfluxDB Edition that is desired for production

---

# InfluxDB for Infrastructure Monitoring

# Why InfluxDB for Infrastructure Monitoring?

- ✓ Manage high cardinality data across thousands of devices and interfaces
- ✓ Reduced operational complexity by consolidating data sources into one unified view
- ✓ Anomaly detection and faster troubleshooting with real-time access to granular data
- ✓ Reduced total cost of ownership (TCO) by storing data efficiently with high compression rates

# Partnerships & Integrations

## VISUALIZATION & ORCHESTRATION



## PHYSICAL & VIRTUAL

Applications

Services


API

Infrastructure

Monitoring

Sensors/Devices

## DATA TRANSPORT



## REAL-TIME ANALYTICS DATABASE



Edge

On-Perm

Cloud

## AI & ML

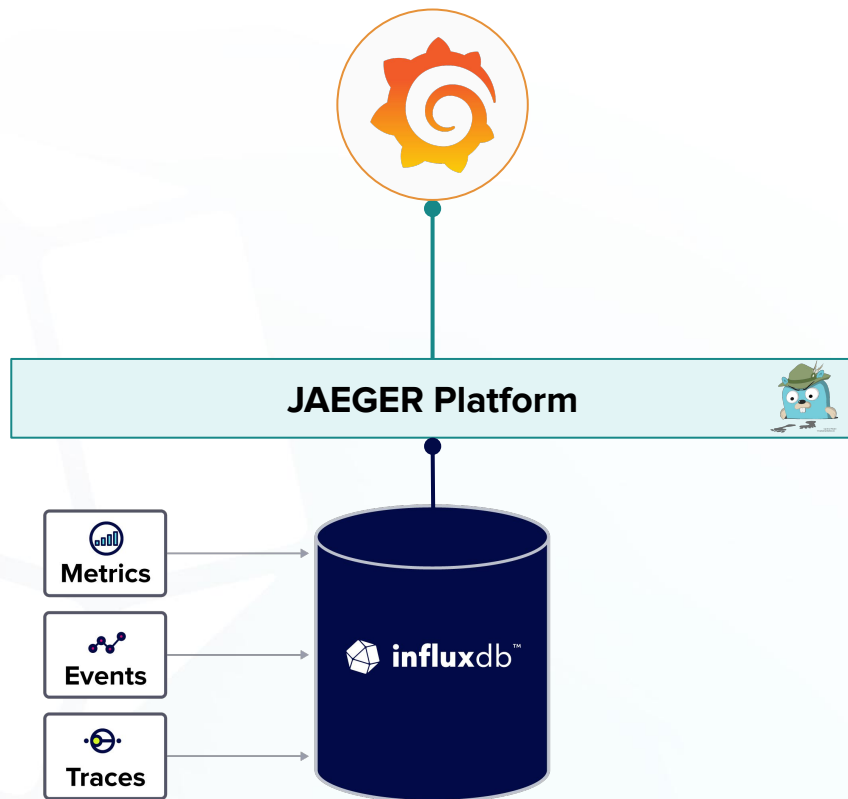


# Jaeger Compatible

**Jaeger: open source, distributed tracing platform. (Originally built by Uber)**

InfluxDB 3.0 is a compatible remote backend store for Jaeger. This means InfluxDB can be bolted to your current trace infrastructure.

<https://killercoda.com/influxdata/course/demos/otel>



---

# Customer Use Case



# Data Center Solution



## Problem:

Unable to collect and analyze data center metrics into a centralized TSDB. Resulting in lack of real time access to data center health and a single time series data pipeline for advanced analytics (ML/AI)



## InfluxDB Value:

Able to see data center metrics in a single pane to begin analysis high volumes of data across multiple sites



**influxdb**™




# Advertising Service

## Problem

Infrastructure monitoring crashes during high traffic events like the Superbowl which directly impact their revenue.

## InfluxDB Value

- Availability and horizontal scaling issues solved for the 96 databases that store their monitoring data
- High performance querying even with cardinality in their dataset at 8M



---

# Useful Links

[InfluxDB 3.0](#)

[Request a Proof of Concept](#)

[sales@influxdata.com](mailto:sales@influxdata.com)