How Uber Tackled Massive Telemetry

Who am I?

Martin Mao

- **2019 - now: Chronosphere**
  - Co-founder/CEO
- **2015 - 2019: Uber**
  - Co-creator + Technical Lead of M3
  - Manager of M3 teams
- **2013 - 2015: AWS**
  - Technical Lead on EC2
  - Technical Lead on AWS Systems Manager

Agenda

Telemetry at Uber

Evolution of Telemetry Stack

M3 Architecture

Telemetry @ Uber

Uber Mission: Make transportation as reliable as running water.

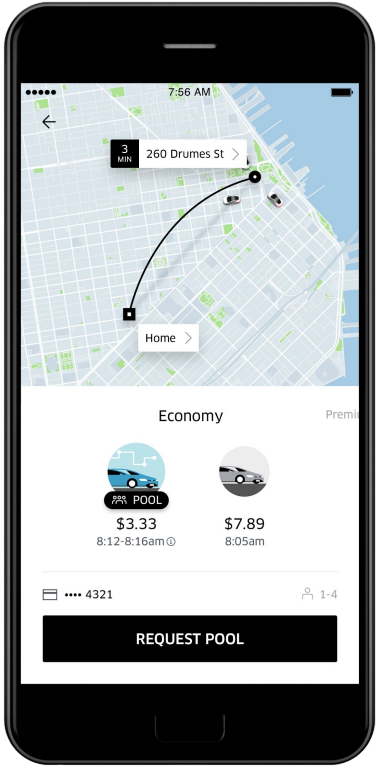

Telemetry @ Uber

Uber Mission: Make transportation as reliable as running water.

Single Company Wide Telemetry Platform:

- **System Monitoring:** Container, physical host/system and network.
- **Application Monitoring:** Every microservices' latency, success/error rates plus custom measurements.
- **Business Monitoring:** Every measurement for every product in every city.

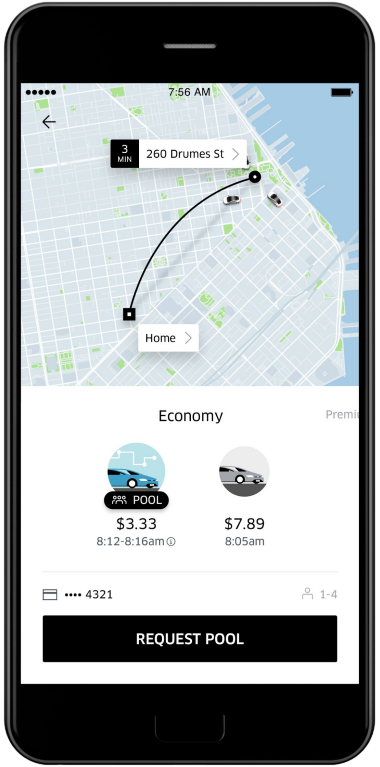

Evolution of the Telemetry Stack

Evolution of the Telemetry Stack

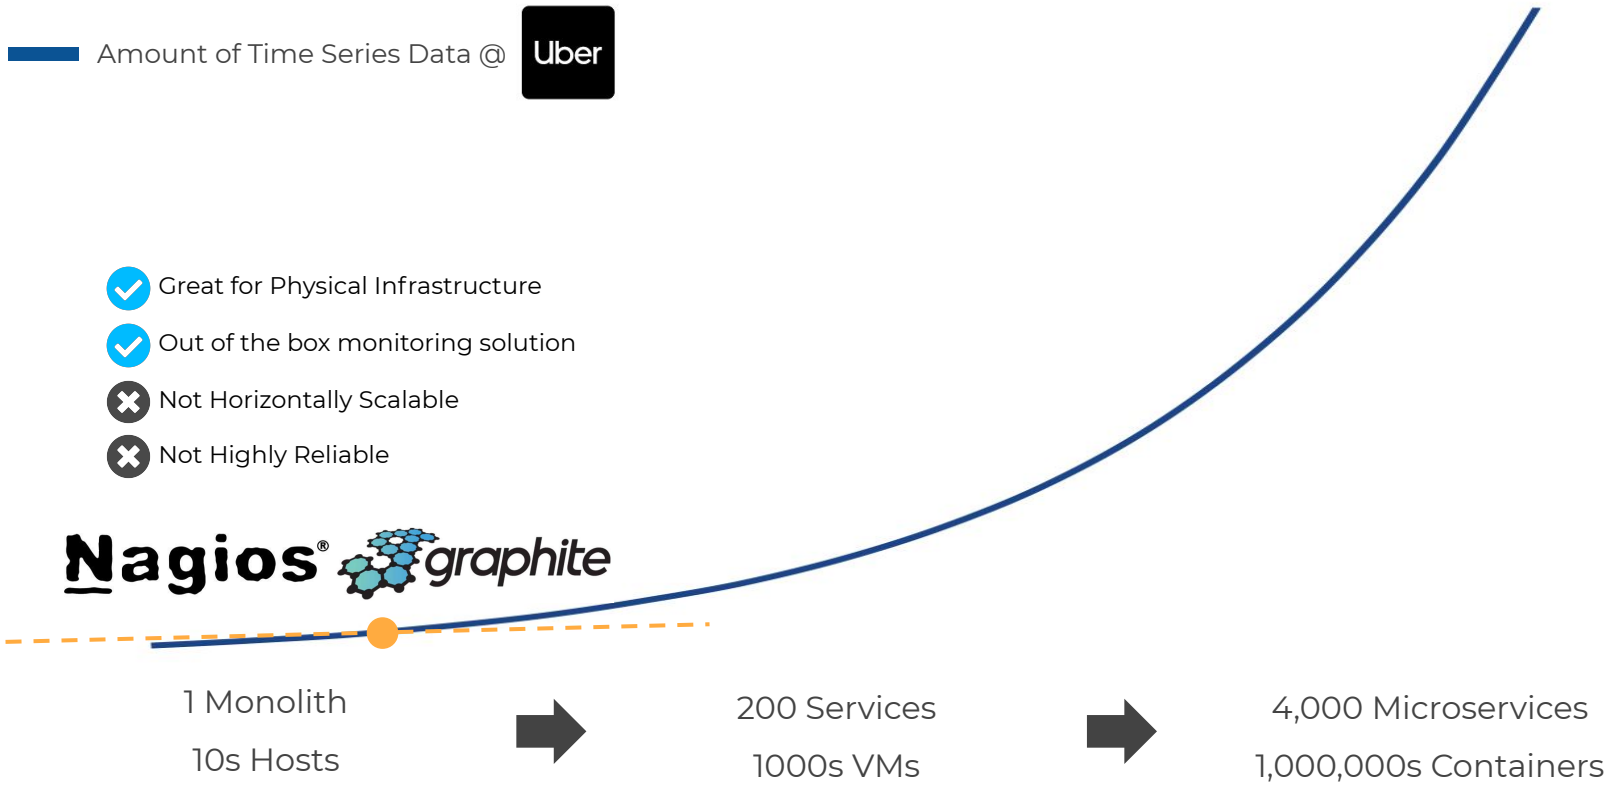

Evolution of the Telemetry Stack

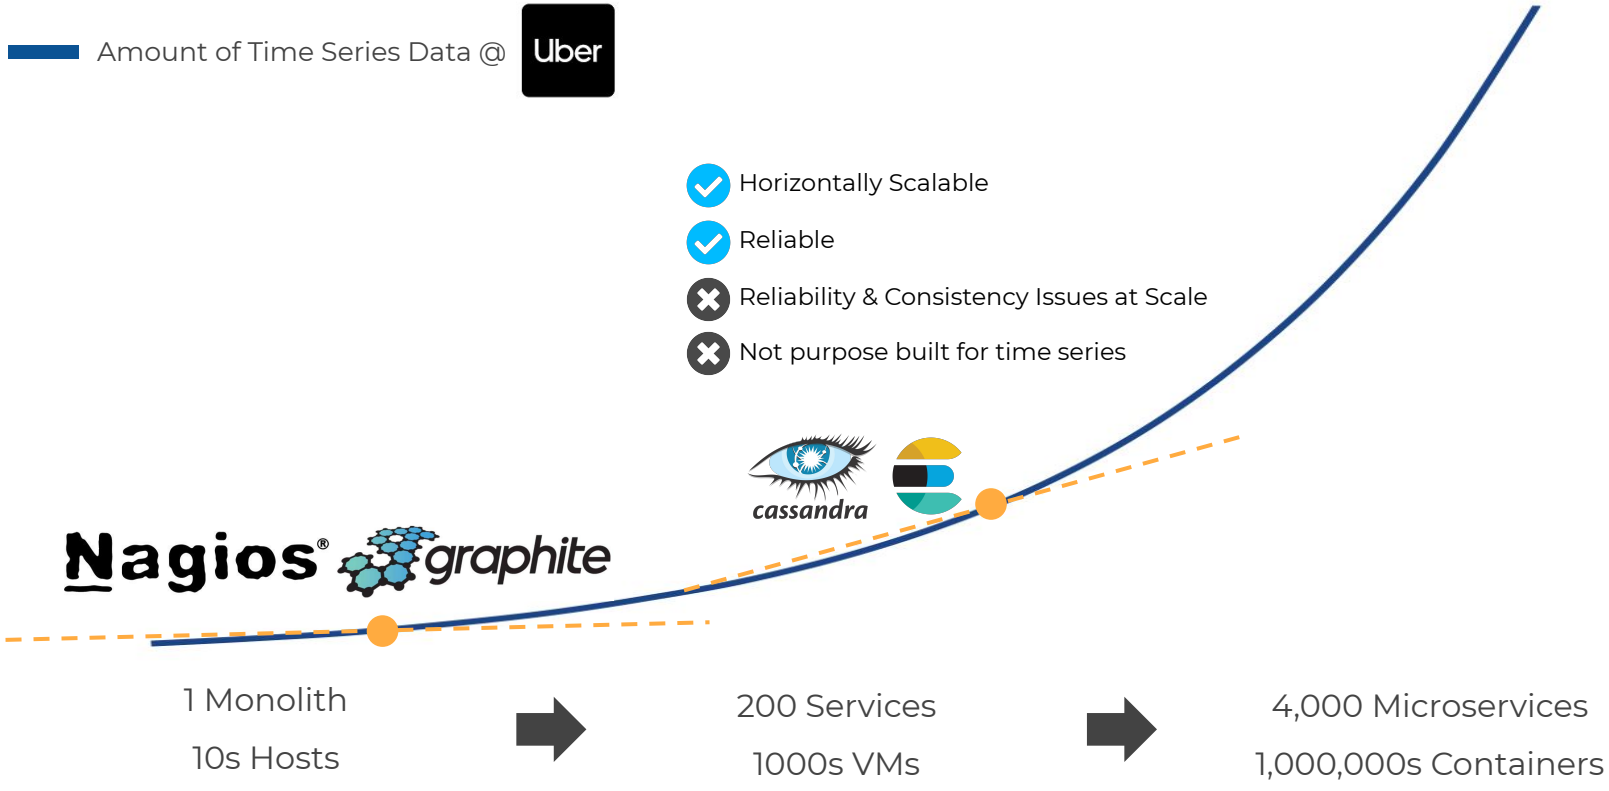

Evolution of the Telemetry Stack

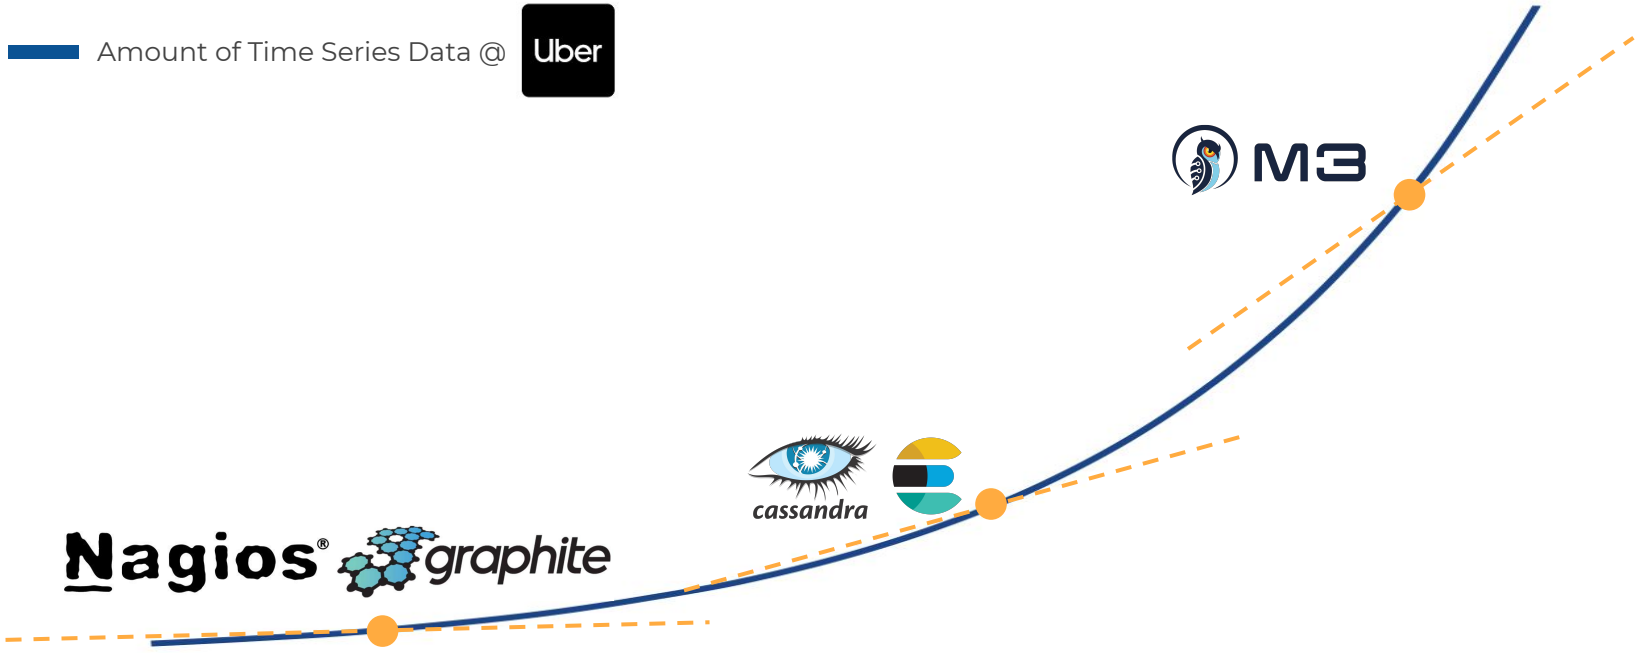

M3 Design Principles

- Highly Reliable & Consistent:
  - Three consistent copies of all data.
  - Tolerates node, AZ or region failures.

M3 Design Principles

- Highly Reliable & Consistent:
  - Three consistent copies of all data.
  - Tolerates node, AZ or region failures.
- Highly Scalable:
  - Horizontally scalable to store billions of metric time series.
  - Built with simple architecture & operation in mind.

M3 Design Principles

- Highly Reliable & Consistent:
  - Three consistent copies of all data.
  - Tolerates node, AZ or region failures.
- Highly Scalable:
  - Horizontally scalable to store billions of metric time series.
  - Built with simple architecture & operation in mind.
- Purpose Built & Efficient:
  - Custom built time series database with metric index.
  - Optimized compression algorithm for time series data.
  - Built in downsampling and aggregation.

M3 Architecture

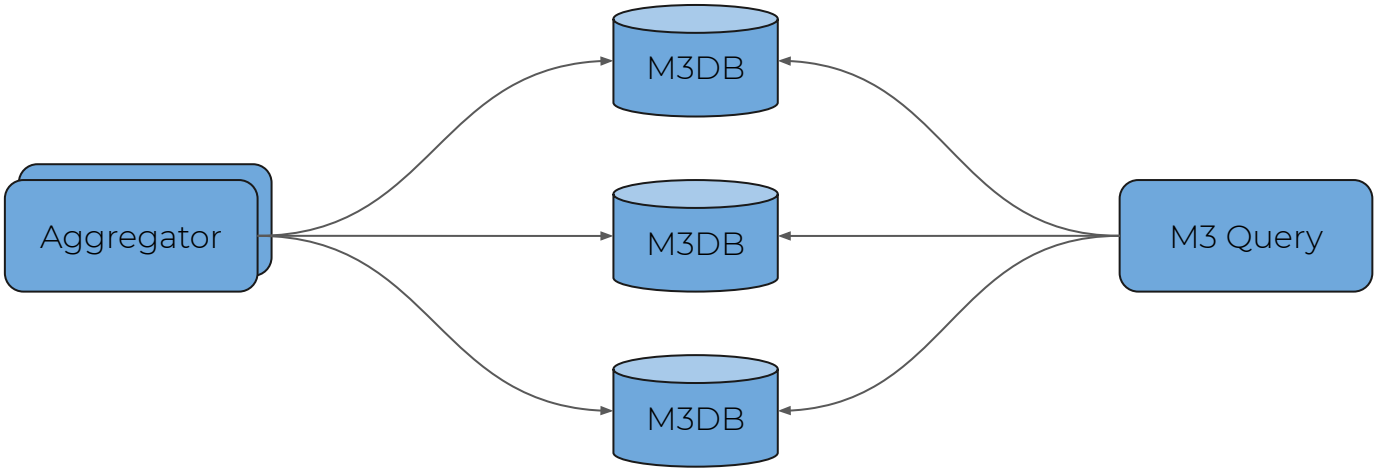

M3 Architecture

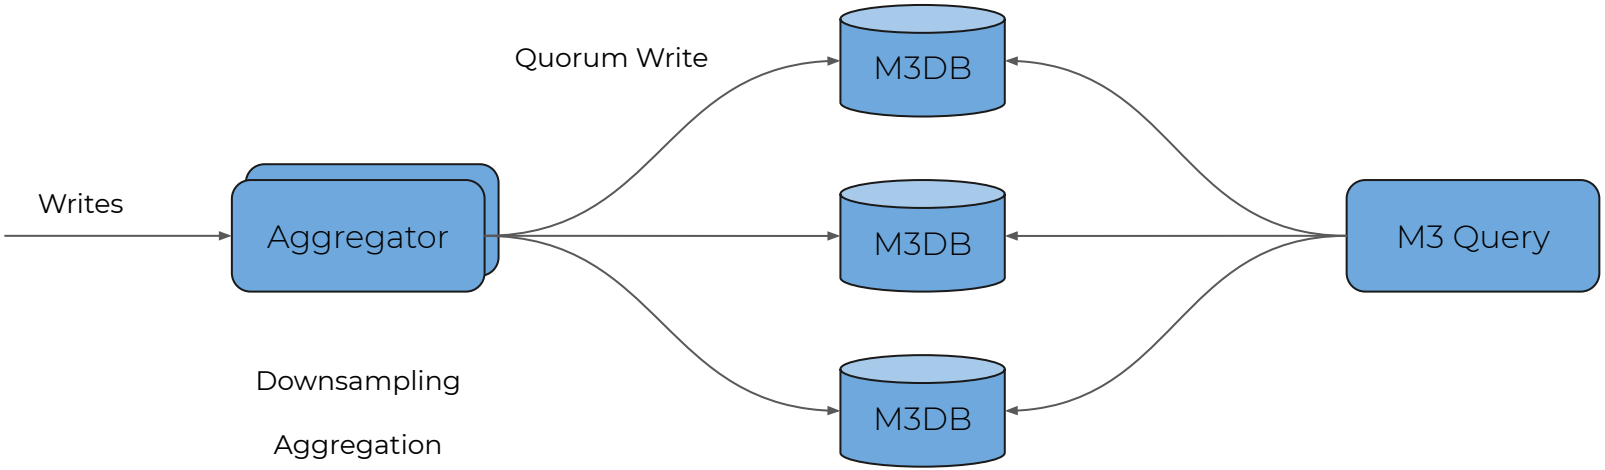

M3 Architecture

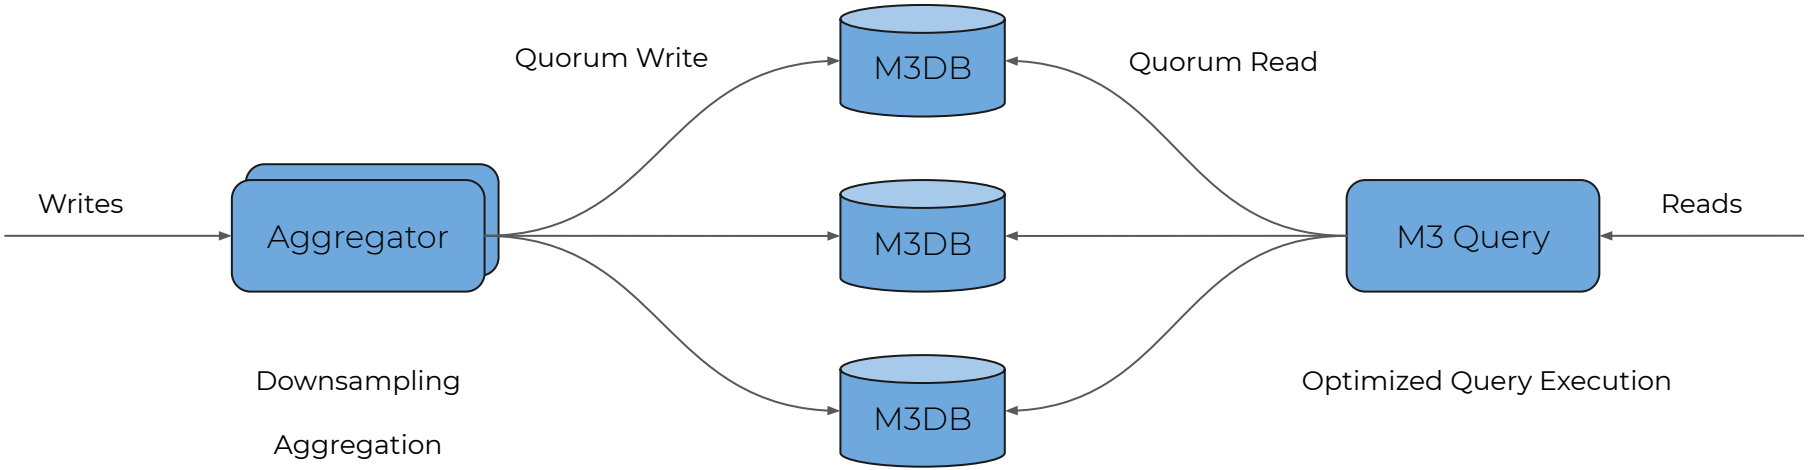

M3 Architecture

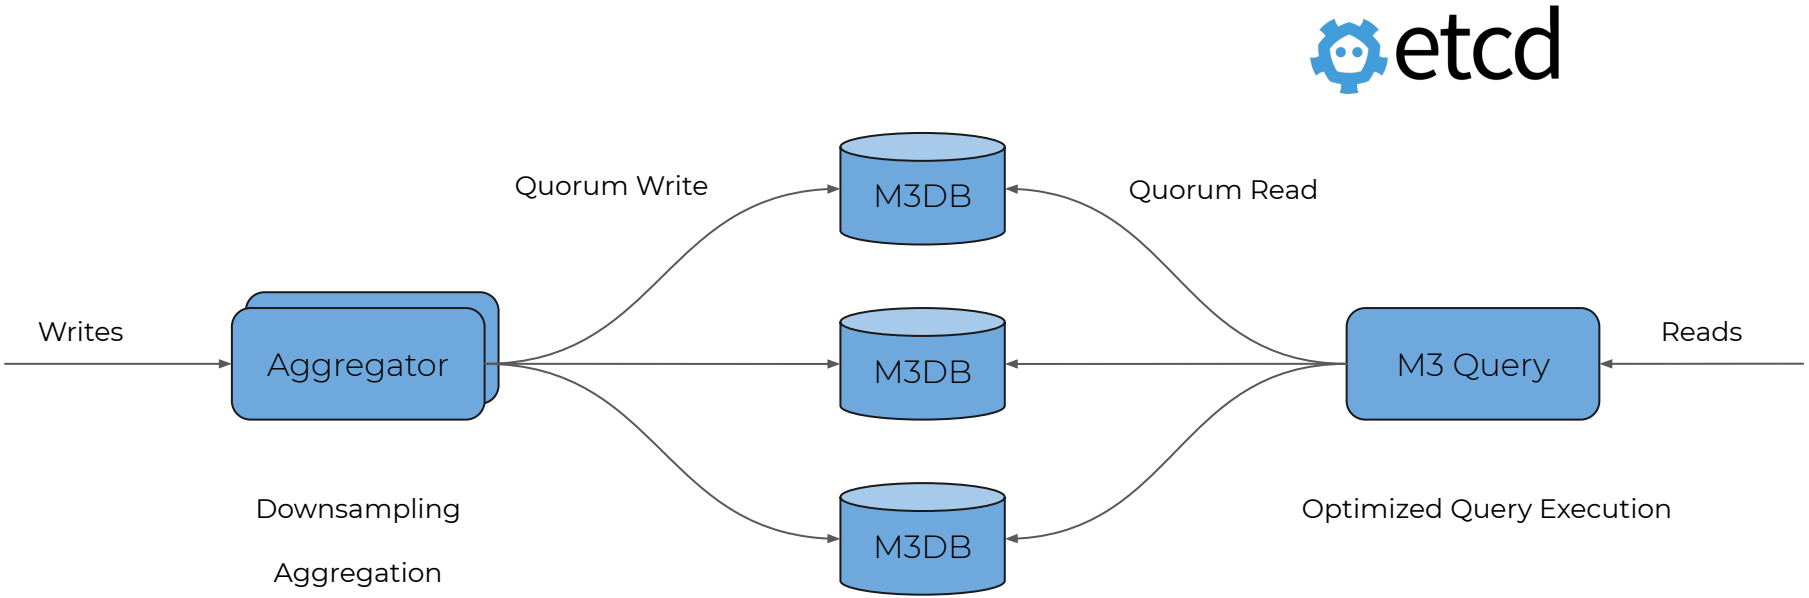

M3DB

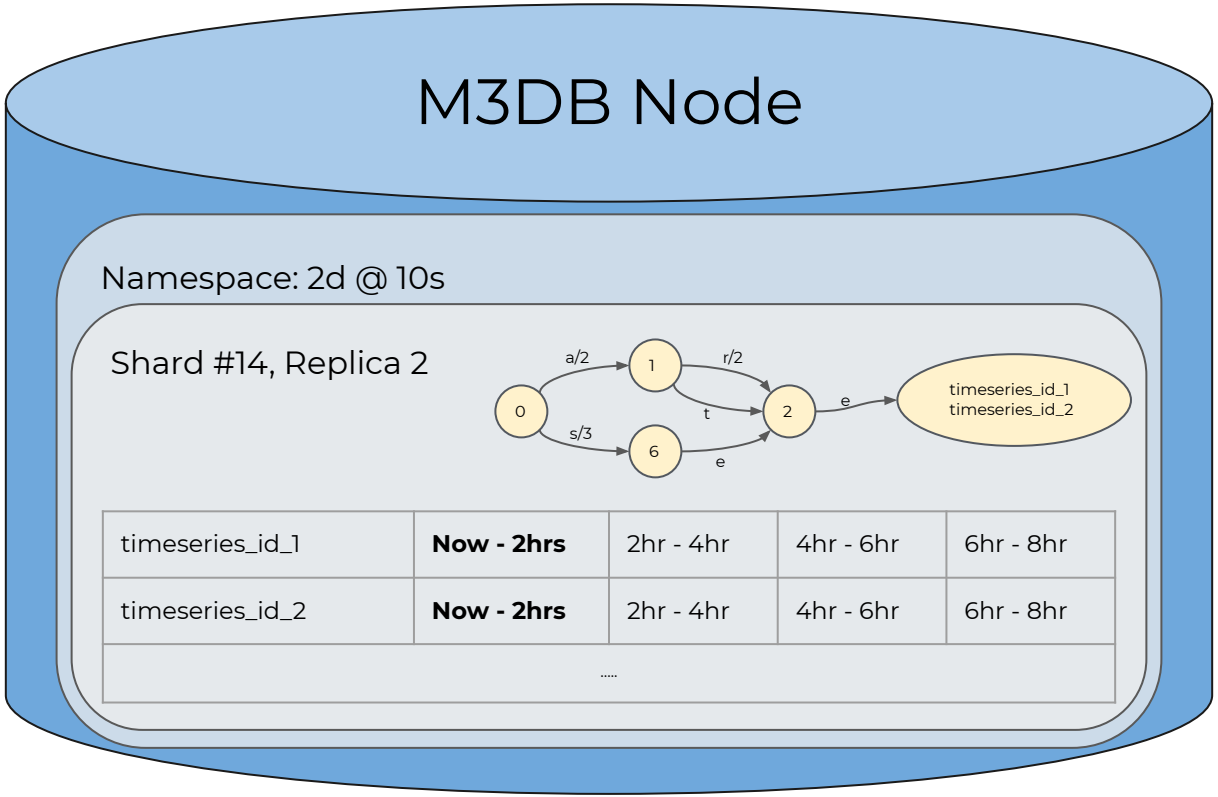
M3DB Node
Namespace: 2d @ 10s
Shard #14, Replica 2
a/2
r/2
s/3
t
e
e
0
1
2
6
timeseries_id_1
timeseries_id_2
timeseries_id_1 | Now - 2hrs | 2hr - 4hr | 4hr - 6hr | 6hr - 8hr
timeseries_id_2 | Now - 2hrs | 2hr - 4hr | 4hr - 6hr | 6hr - 8hr
.....

M3 Multi-Zone

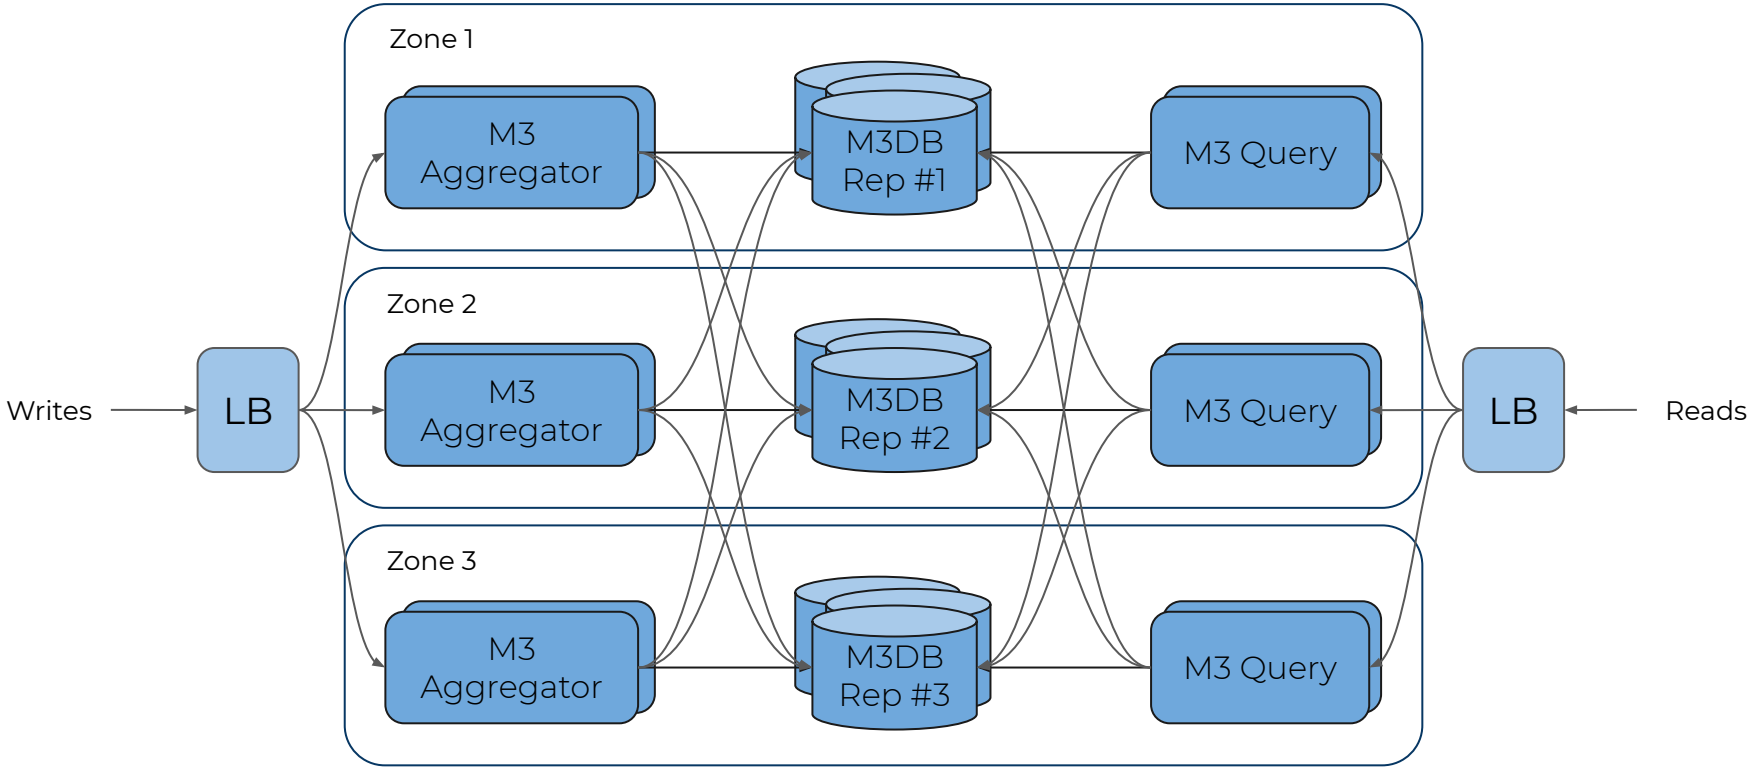

M3 Multi-Zone

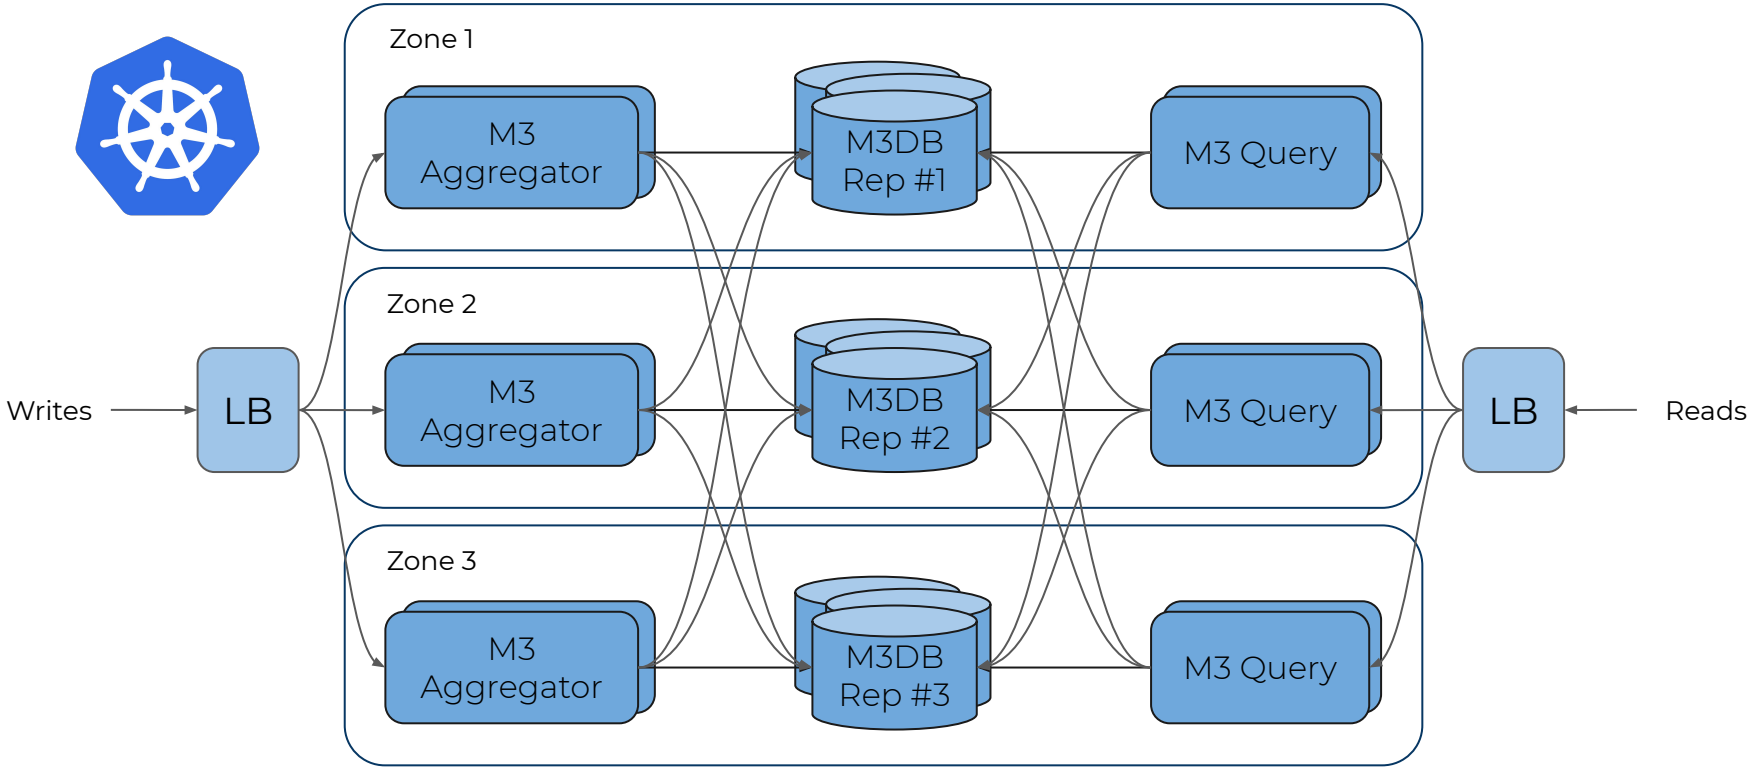

M3 Numbers

Not a STAC benchmark

Performance

Datapoints Written & Read per Second

1.5B+ 3B+

Collected by the platform and visible within seconds to serve 500,000 alerts, 4,000 dashboards, anomaly detection and analytics

Scalability

Total Metric Time Series Stored

12B+

One of the largest production telemetry platforms in the world - similar scale to Google's in-house platform Monarch

Efficiency

Cost Savings

10X

Reduction in hardware spend when compared to previous version of metrics platform built on Cassandra

Reliability

System Uptime

99.99%

System was always up as it was relied upon to monitor not only the infrastructure and applications, but business operations as well

M3 Community & Resources

- Website: https://m3db.io
- Slack: https://bit.ly/m3slack
- GitHub: https://github.com/m3db/m3
- Documentation: https://docs.m3db.io/
- Office Hours (every 3rd Thursday, 11-1pm EST): sign up for slot here
- M3 Community Meetup: https://www.meetup.com/M3-Community/

Thanks + Q&A