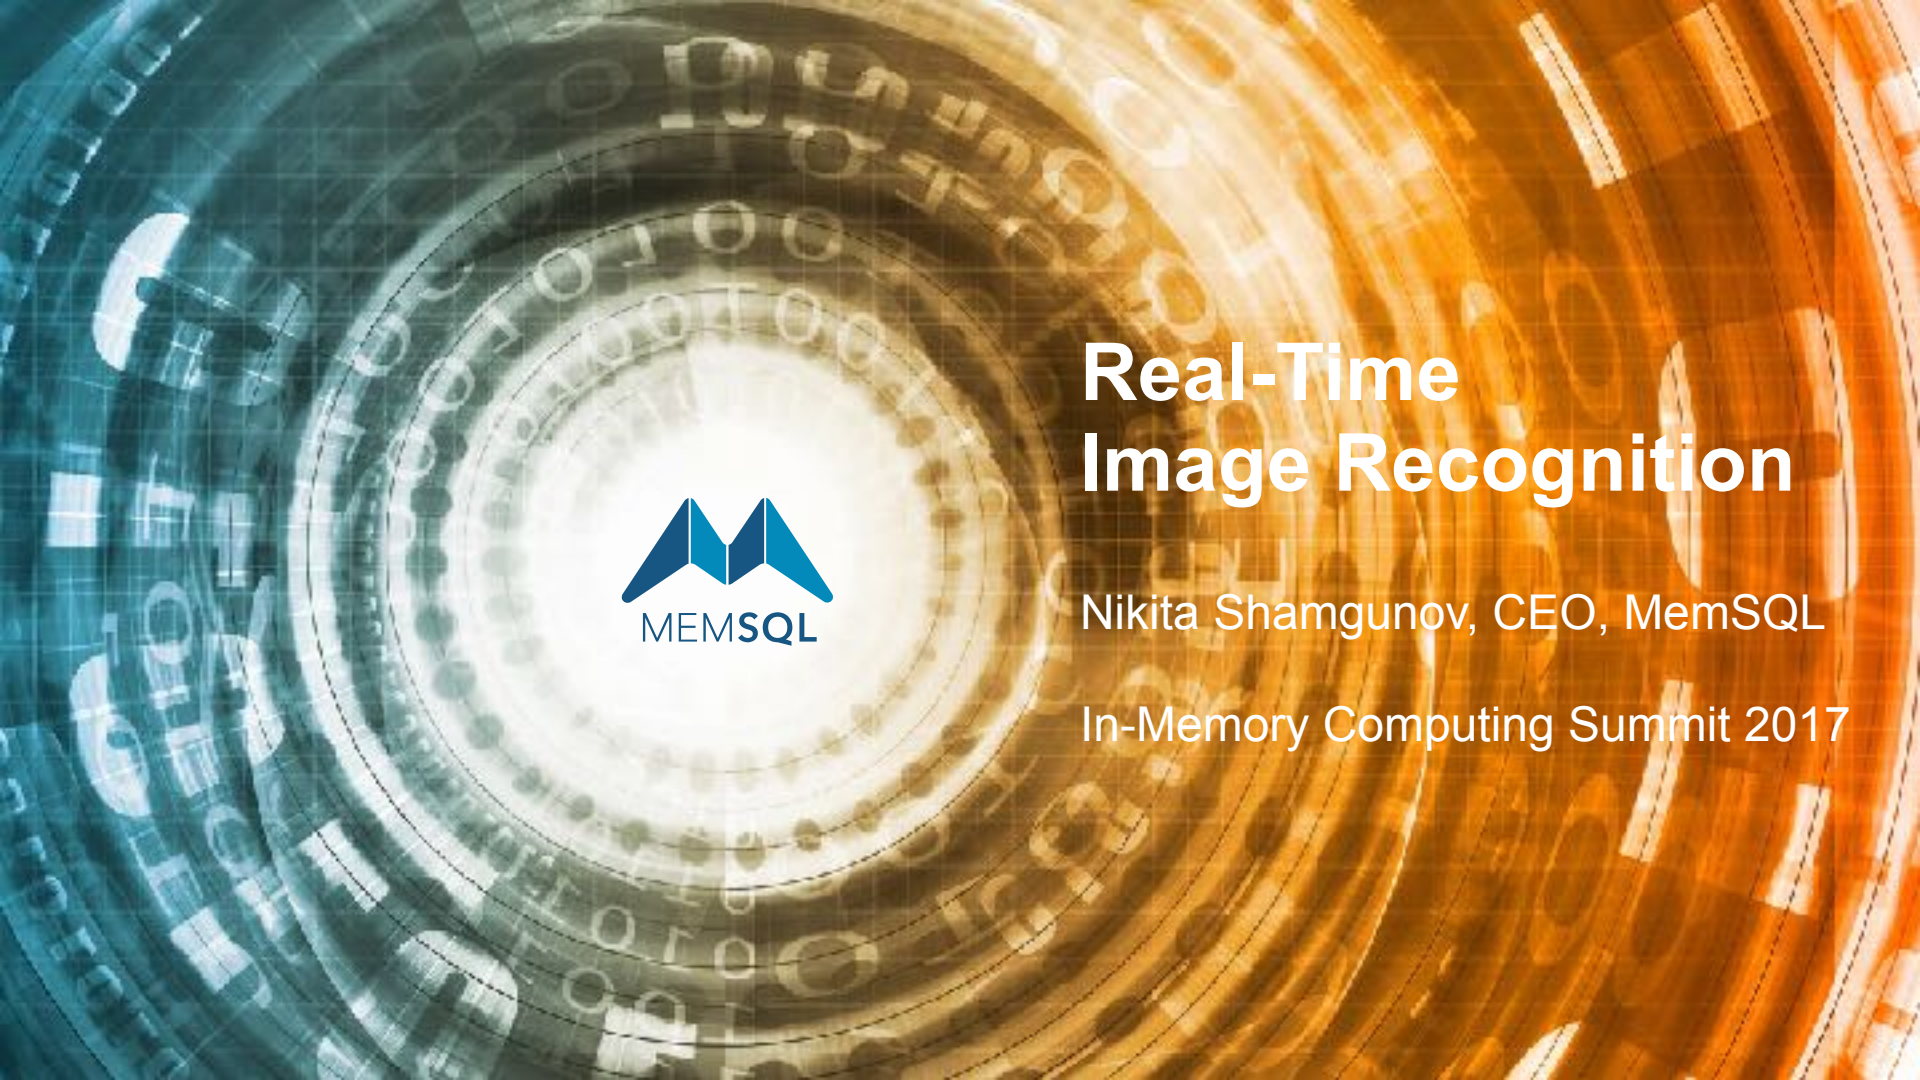

Real-Time Image Recognition

Nikita Shamgunov, CEO, MemSQL

In-Memory Computing Summit 2017

The future of computing is visual


and also
numerical :)


You're almost there!
STOP NAVIGATION


11:36
Search

Putting image recognition to work today

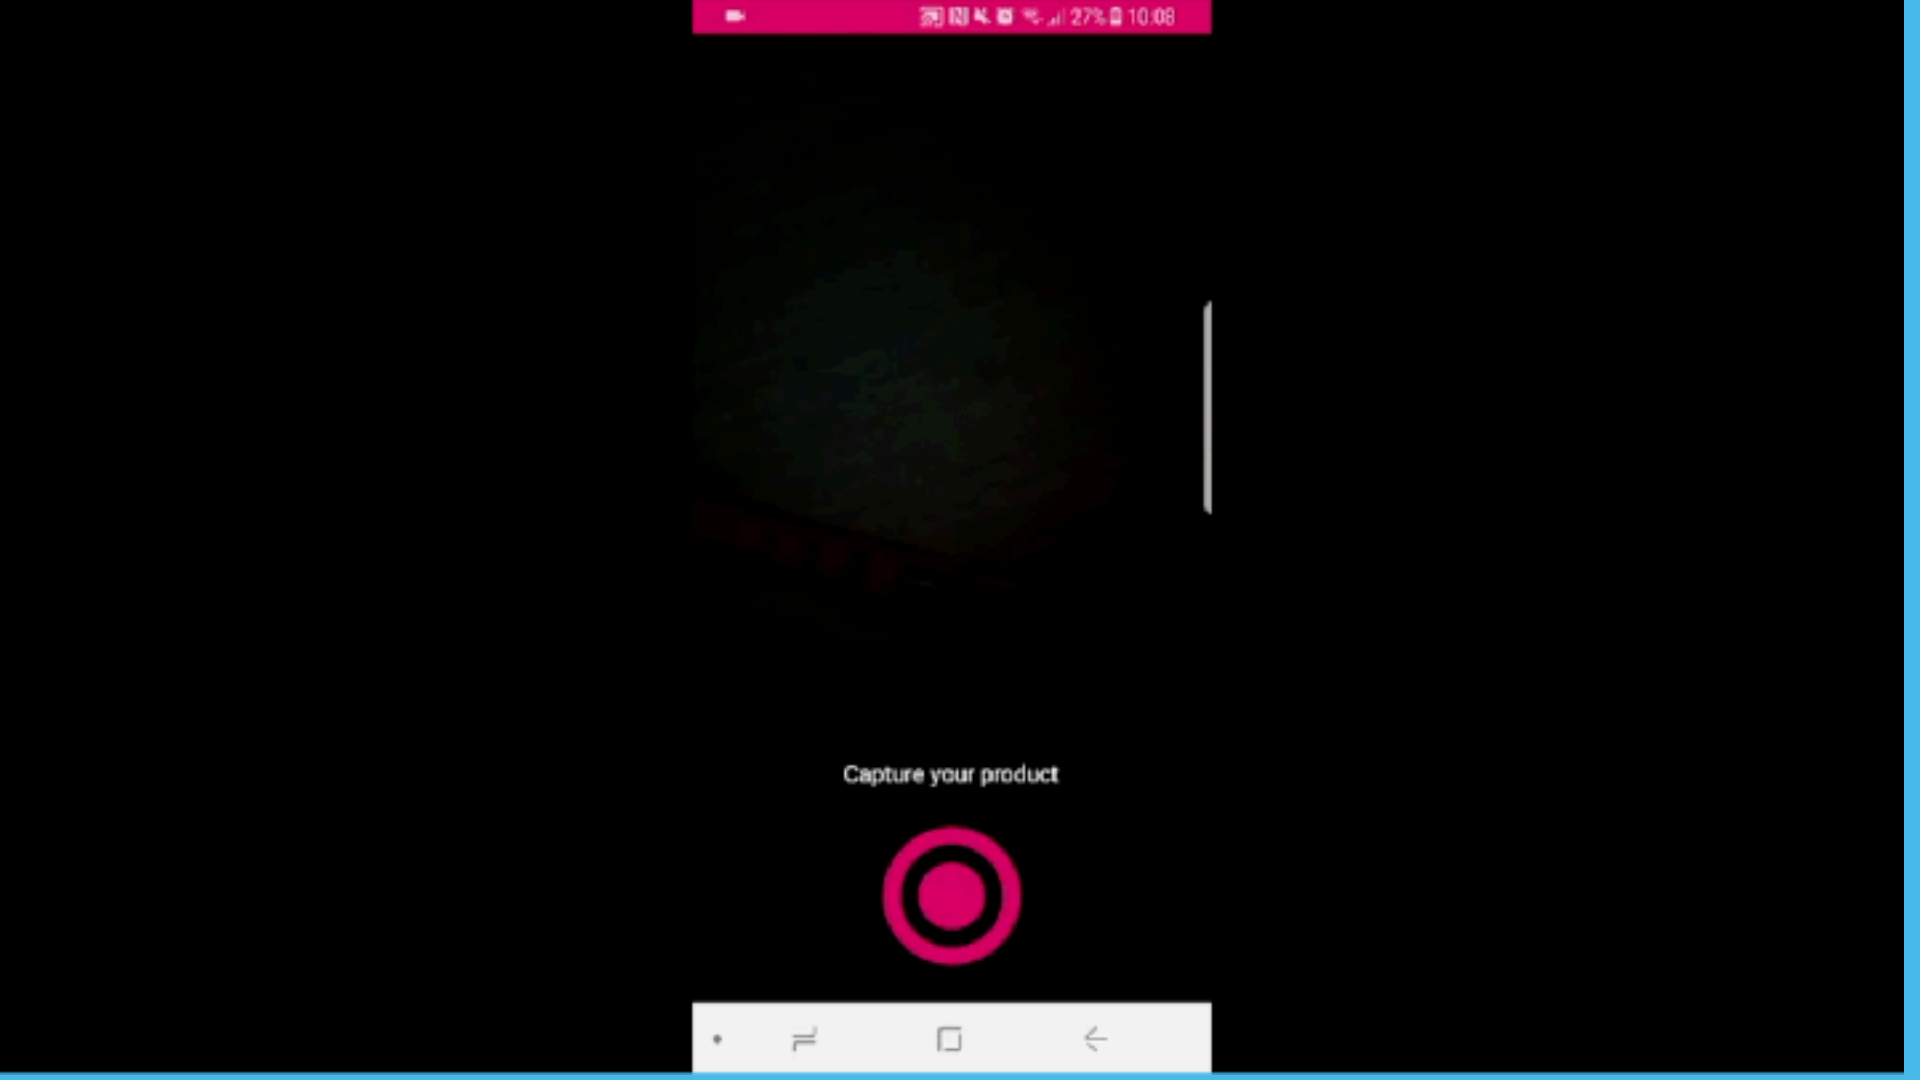
27% 10:08
Capture your product

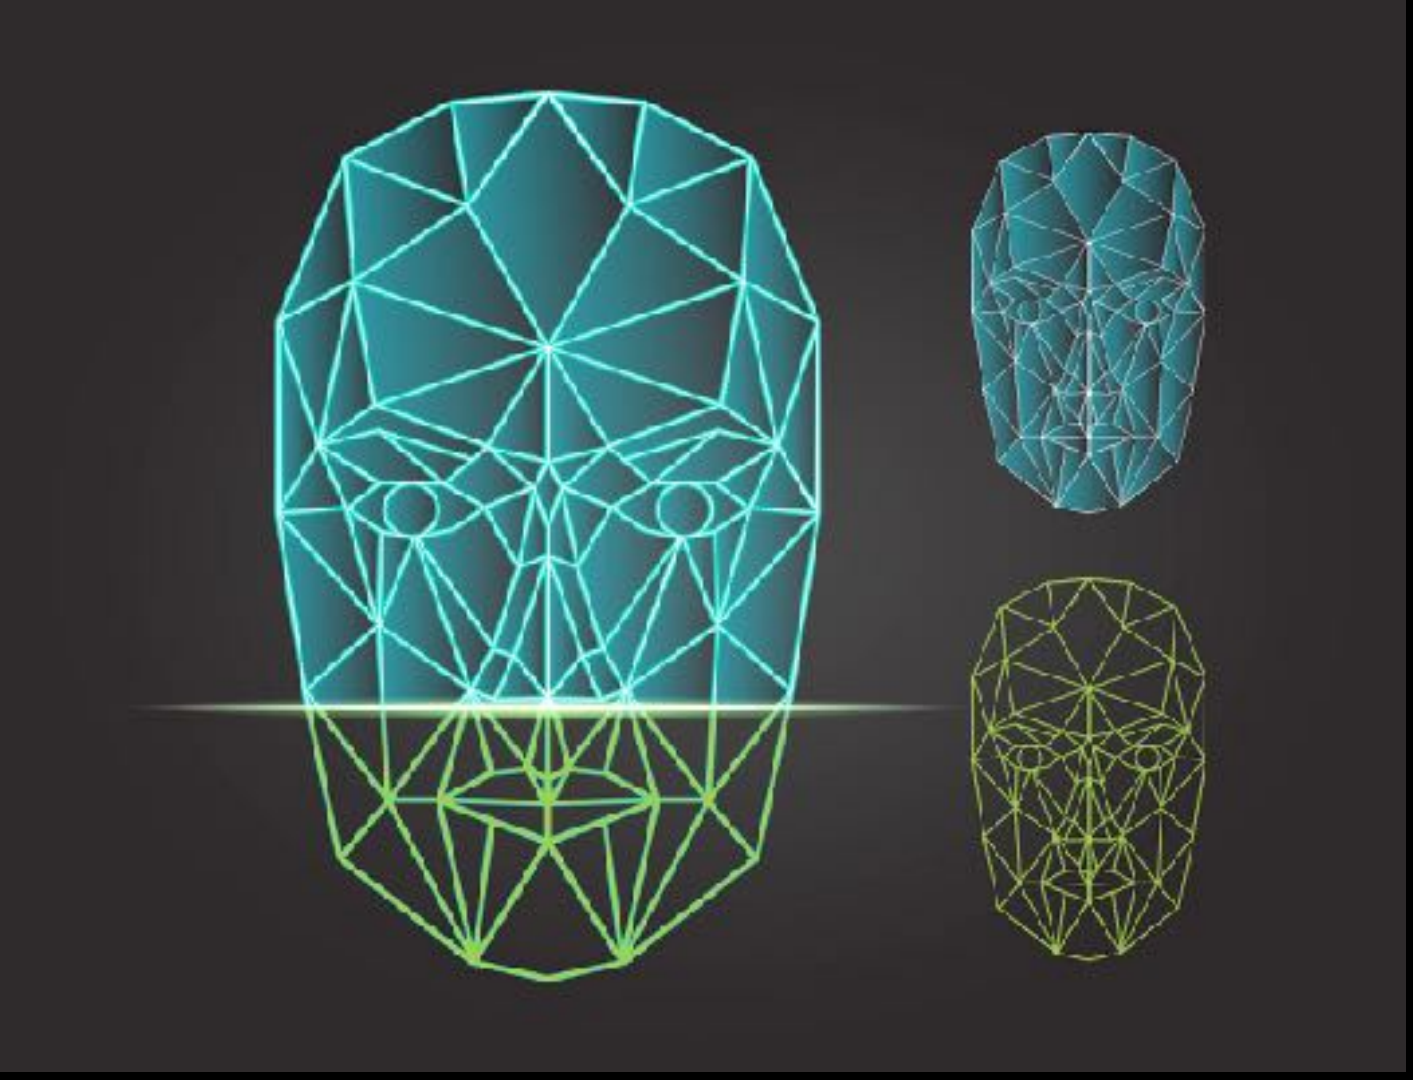

How It Works

Real-Time Image Recognition Workflow

- Train the model with Spark, TensorFlow, and Gluon
- Use the Model to extract feature vectors from images
  - Model + Image => FV
- You can store every feature vector in a MemSQL table

```
CREATE TABLE features (
  id bigint(11) NOT NULL AUTO_INCREMENT,
  image binary(4096) DEFAULT NULL,
  KEY id (id)USING CLUSTERED COLUMNSTORE
)
```

Working with Feature Vectors

For every image, we store an `ID` and a normalized feature vector in a MemSQL table called `features`.

```
ID      |       Feature Vector
x       |       4KB
```

To find similar images, we use this SQL query

```
SELECT
 id
FROM
 features
WHERE
 DOT_PRODUCT(feature * <input>) > 0.9
```

Understanding Dot Product

- Dot Product is an algebraic operation
  - SUM(Xi*Yi) TODO: Put a formula
- With the specific model and normalized feature vectors DOT PRODUCT results in a similarity score
  - The closer the score is to 1 the more similar are the images

Performance Enhancing Techniques

Achieving best-in-class Dot Product implementation

- SIMD-powered
- Data compression
- Query parallelism
- Scale out

- Result: Processing at **Memory Bandwidth Speed**

Performance Numbers

- Memory Speed: 50GB/sec
- Each vector 4K
- **12.5 Million Images a second** per node
  or
- **1 Billion images a second** on 100 node cluster

Demo

Demo Architecture

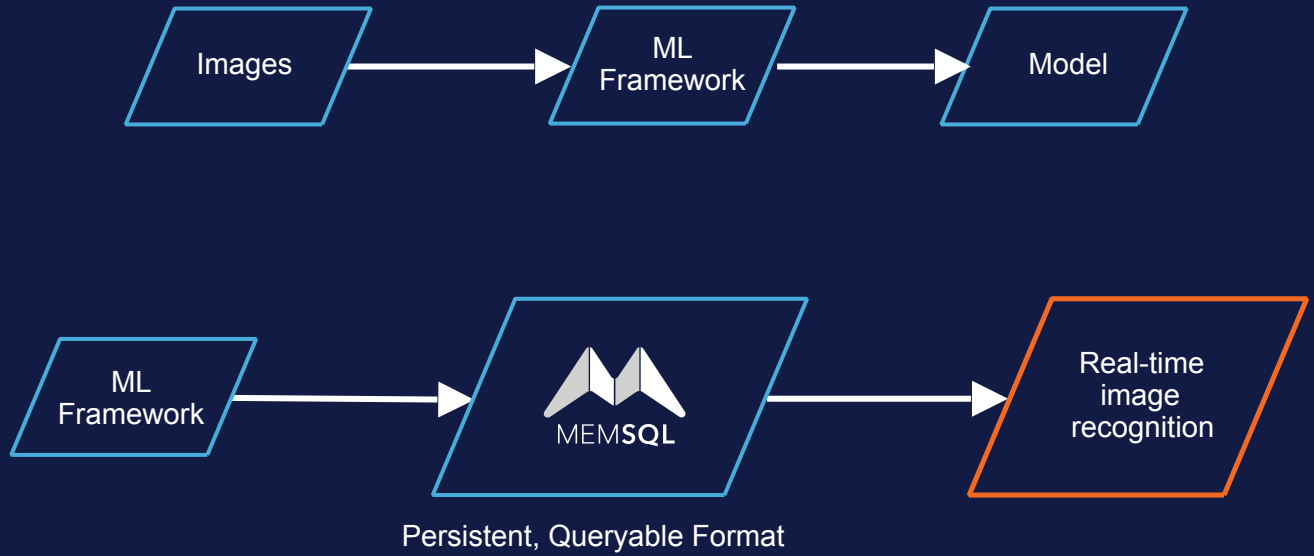

```
SELECT
    id
FROM
    features
WHERE
    DOT_PRODUCT(image, 0xa334efa…)
```

About MemSQL

MemSQL: The Real-Time Data Warehouse

- Scalable
  - Petabyte scale
  - High concurrency
  - System of record
- Real-time
  - Operational
- Compatible
  - ETL
  - Business Intelligence
  - Kafka
  - Spark

- Deployment
  - MemSQL Cloud
  - Any public cloud
  - On-premises
- Developer Edition
  - Unlimited scale
  - Limited high availability and security features

2017 Magic Quadrant for Data Management Solutions for Analytics

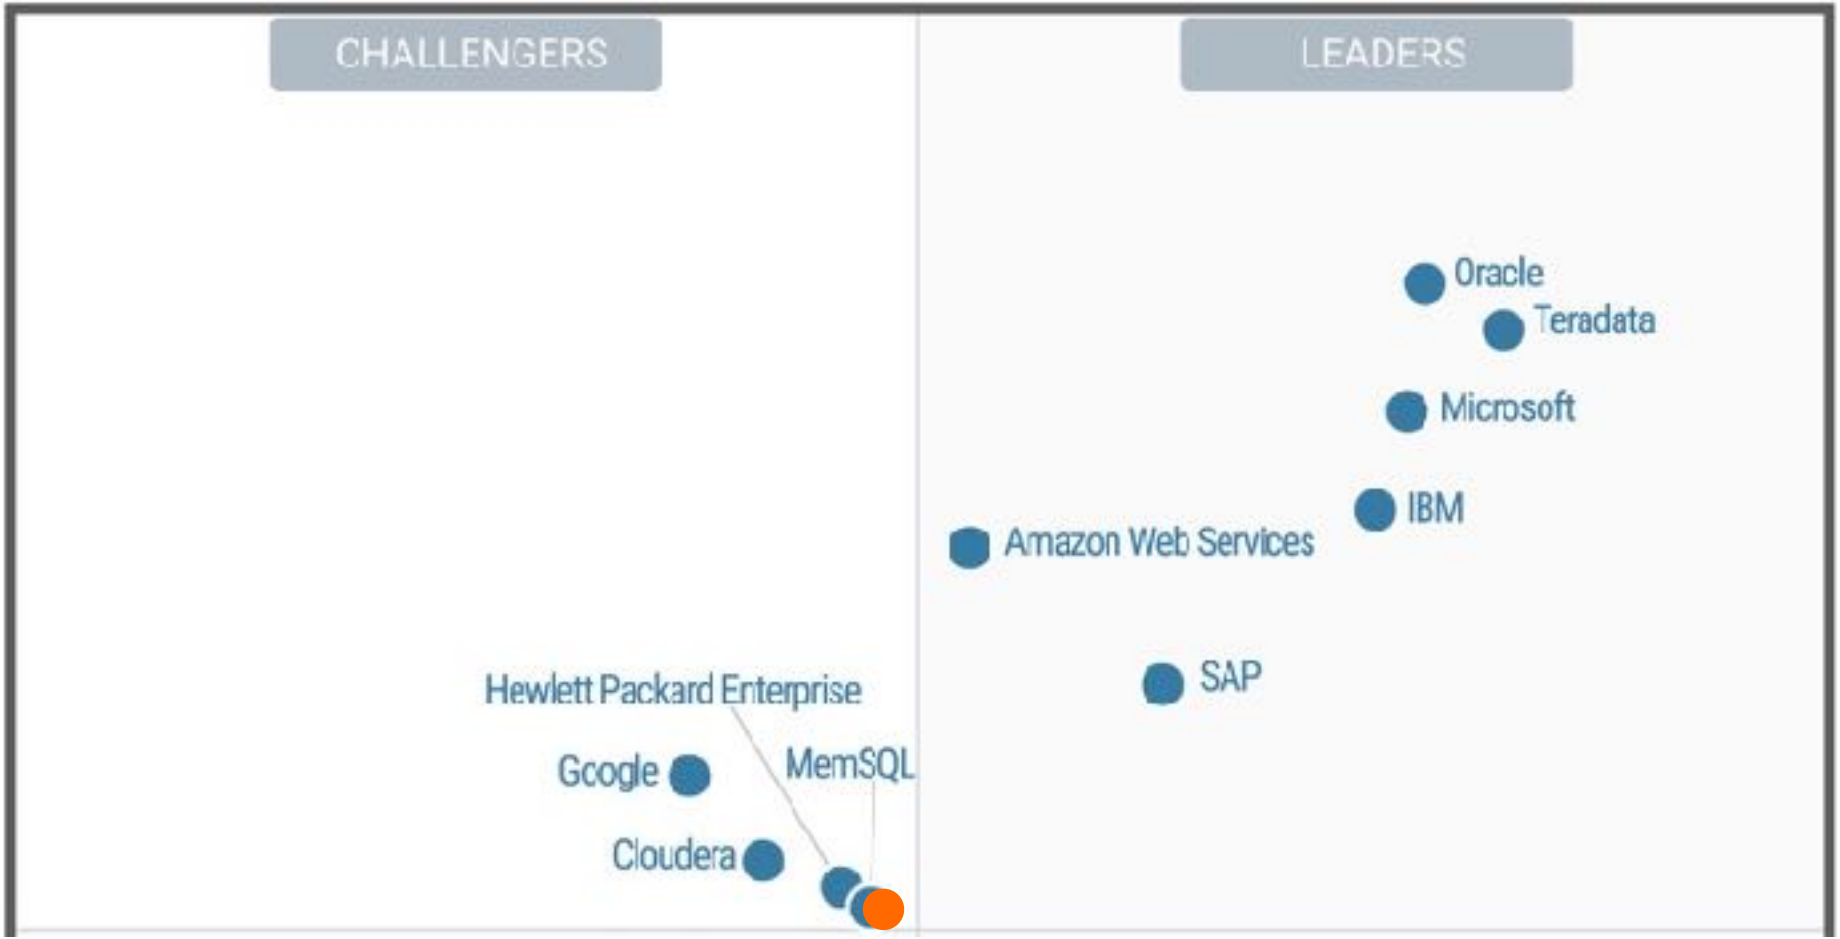

About ML Training

ML training is available through a variety of frameworks, including Spark MLlib, TensorFlow, Gluon, and Caffe.


APACHE
Spark
TensorFlow ™
GLUON
Caffe
+
MEMSQL

Understanding ML Frameworks and MemSQL

ML Frameworks

Fast, large scale
General processing engines
Great for training

MemSQL

Fast, large scale
Real-time data warehouse
Great for real-time scoring

Example: MemSQL Spark Connector

Highly parallel, high throughput, bi-directional

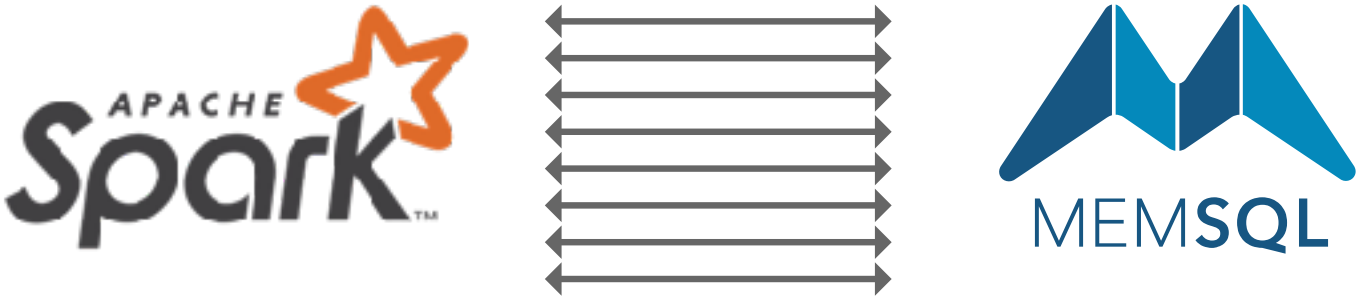

Thank you!

@NikitaShamgunov

www.memsql.com