

WEBINAR

Operationalizing ML & AI with MemSQL

July 25, 2019



Today's Agenda

Submit questions in the Question Box



Eric Hanson

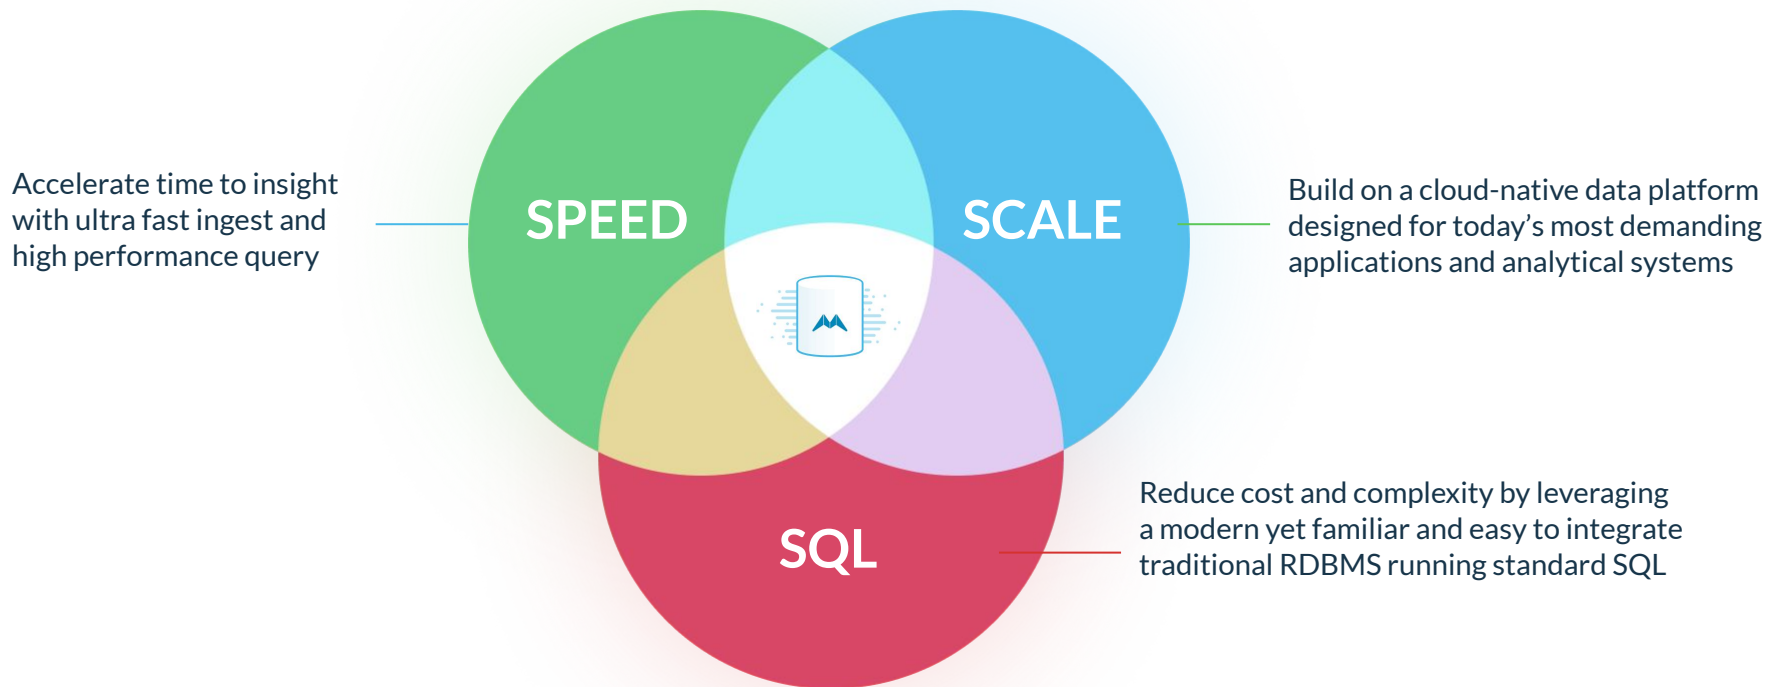
Principal Product Manager
hanson@memsql.com


- MemSQL Overview
- The Promise and Challenges Delivering Real-Time ML/AI Applications
- How MemSQL Operationalizes ML Applications
- Q&A

MemSQL Overview

MemSQL: The No-Limits Database

The cloud-native operational database built for speed and scale






ML/AI & MemSQL

The Promise of ML/AI

Artificial intelligence (AI) has the potential to create tremendous value across sectors



Share of AI impact in total impact derived from analytics, %

Source: McKinsey Global Institute analysis

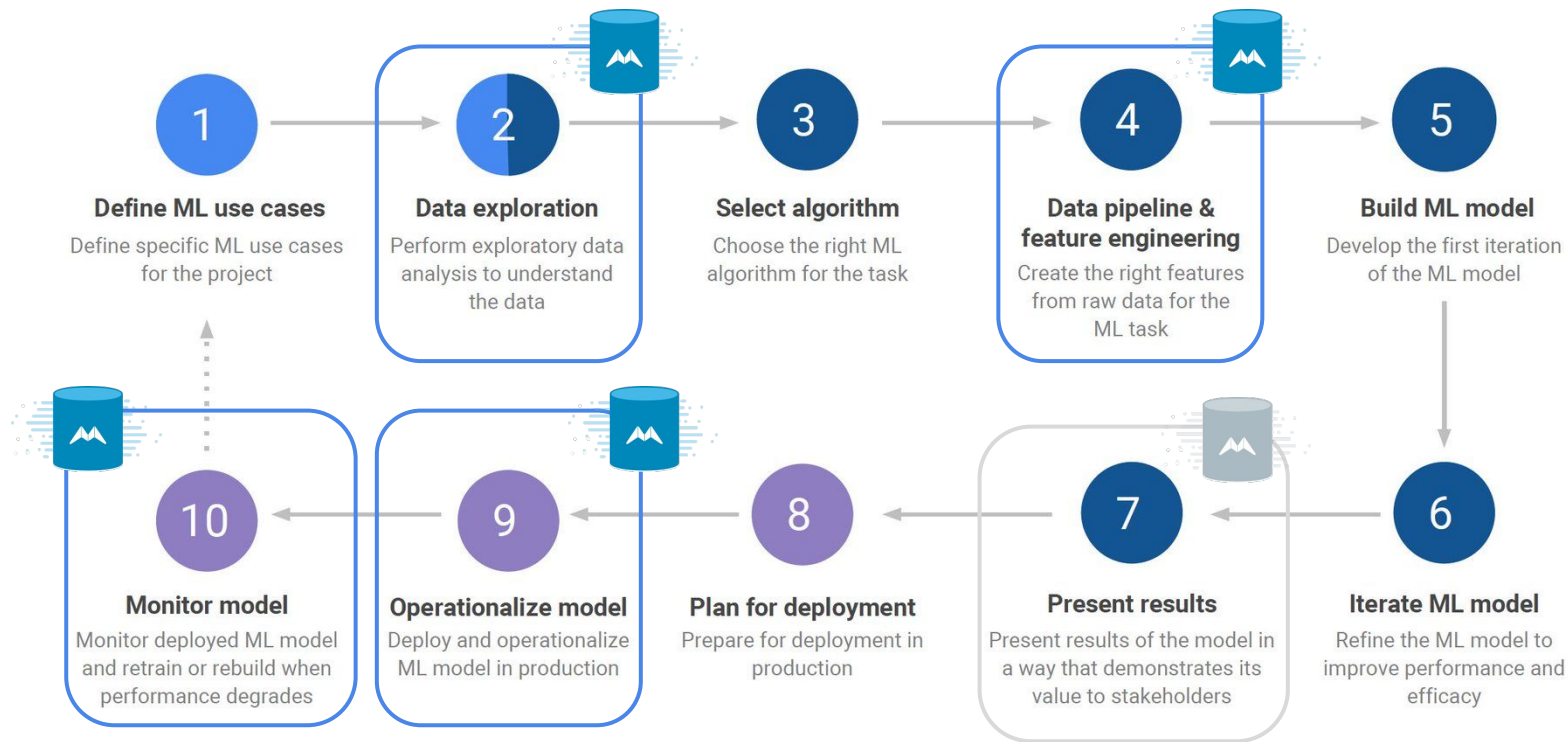
Challenges Delivering a Real-Time ML Application

- Data exploration & preparation
- Algorithm selection, model training & refinement
- Assembling data for scoring
- Scoring
- Model monitoring, refinement, & retraining

Challenges Delivering a Real-Time ML Application

- Data exploration & preparation
- Algorithm selection, model training & refinement
- Assembling data for scoring
- Scoring
- Model monitoring, refinement, & retraining

MemSQL Drives Improvement Across the ML Lifecycle



Operationalize ML/AI with Speed, Scalability, SQL, & Programmability

- Score data during streaming load with pipelines transforms
- Score fast in-database with
 - vector functions
 - compiled expressions and functions
 - scale-out & parallelism
- Assemble up-to-date feature records to score with an external app




How MemSQL Operationalizes ML Applications

AI/ML Tool Integration with MemSQL

- MemSQL as a source for training data
- Support for virtually all tools that can connect to MySQL, e.g.

pandas

$$y_{it} = \beta^t x_{it} + \mu_i + \epsilon_{it}$$



NumPy



TensorFlow



-  named connector

Use Loading Tools for Bulk or Real-Time Ingest

Real-Time Data



Transforms



ETL/ELT
Informatica, Talend




Data Lakes



Use MemSQL Pipelines for Scoring on Load

- MemSQL Pipelines
 - Automatic streaming loader
- Transform types
 - Python
 - Executable
- Use
 - Compute new "score" column from other columns during load
- Benefits
 - Use with any external scoring code or libraries



Pipeline Transform Example

```
CREATE PIPELINE mypipeline AS  
LOAD DATA KAFKA '192.168.1.100:9092/my-topic'  
WITH TRANSFORM  
( 'http://www.memsql.com/my-transform.tar.gz',  
'my-executable.py', '' )  
INTO TABLE t
```

More at <https://docs.memsql.com/memsql-pipelines/v6.8/transforms/>

Implementing Scoring With Extensibility

- Support for:
 - User-defined functions (UDFs)
 - Stored procedures (SPs)
 - User-defined aggregates (UDAFs)
 - Arrays
 - Records
 - Control structures
- Benefits
 - Combine with other SQL (joins, filters, etc.)
 - Use with SQL-compatible tools
 - Compiled code and scale-out performance
- Implement PMML or other models for in-DB scoring


Example: Regression Line

- $f(x) = 0.73x + 2.95$
- create or replace function `f(x double)`
returns double as

`begin`
 `return 0.73 * x + 2.95;`
`end;`
- `select a, b, f(x) from t;`

Scoring Using Vector Embeddings

- Produced by deep neural net hidden layer
- Embedding maps high-dimensional space to low e.g.
 - High: image (256 X 256 matrix)
 - Low: 1000-element vector
- Applications
 - face matching
 - product photo matching
 - document similarity
 - and on and on!



Real-Time Image Recognition Workflow

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

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

MemSQL Functions for Vector Similarity Matching

- `DOT_PRODUCT(vector, vector)`
- `EUCLIDEAN_DISTANCE(vector, vector)`

(fast SIMD implementations)

Helper functions:

- `JSON_ARRAY_PACK('[float [, ...]]')`
- `VECTOR_SUB(vector, vector)`

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
-----+-----
int     | 4KB binary
```

To find similar images, we use this SQL query

```
SELECT
    id, DOT_PRODUCT(feature, <input>) as score
FROM
    features
WHERE
    DOT_PRODUCT(feature, <input>) > 0.9
ORDER BY score DESC
```

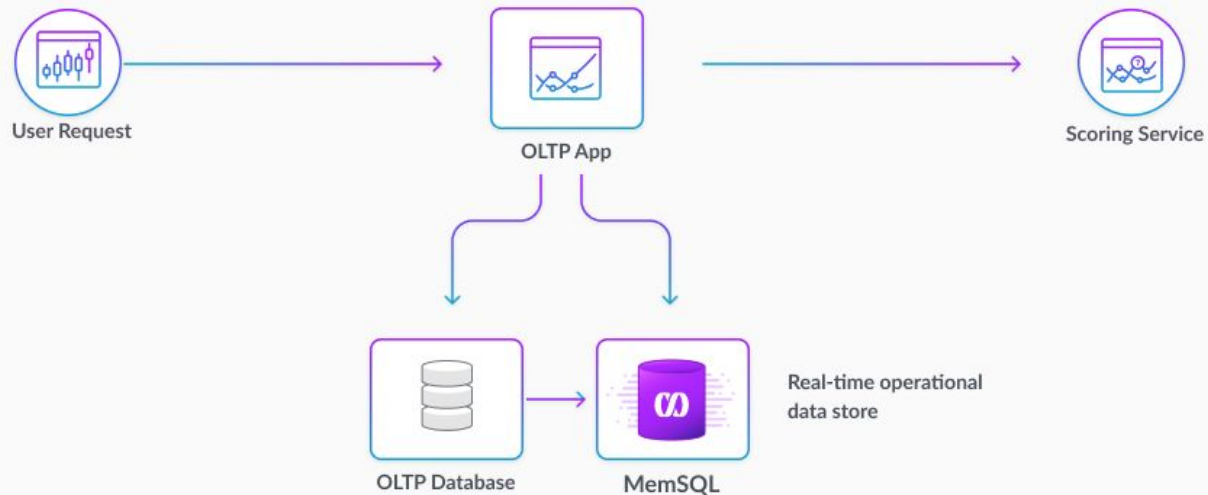
Face Matching with MemSQL

- Face matching, pre-trained networks
 - VGG-Face
 - Facenet
- Match millions of images to a query image in a fraction of a second
- Apps can use this combined with SQL join/filter
- New images can be continuously added and tagged with embedding

Real-Time Feature Retrieval For Scoring in OLTP Apps

- Example: credit card fraud detection
- 1 second budget from card swipe to approval
- Do fraud detection within this 1 second
- 50 msec budget
- 70-value feature record to score

Credit OLTP + Scoring App Architecture



Creating the feature record in real time

Example features:

has transacted with merchant
days since last transaction with merchant
max amount transacted with merchant in last 180 days
min " "
180 days total merchant transaction amount
card present transaction count last 1, 7, 30 days
online transaction count last 1, 7, 30 days
...
(70 or more features!)

Options

1. **Old way:** nightly batch job to accumulate feature record for each customer in traditional "operational data store"; look it up for scoring
2. **New way:** run up to 70 queries concurrently (one per feature or few features) on latest data to get features in < 50 msec *total* in real time operational data store

MemSQL enables #2!

Benefits:

- enables new features like "transactions in last 2 hours"
- fresh features & new kinds of features enable catching more fraud cases

Additional Resources

MemSQL Vector Functions

<https://msql.co/vector-functions>

Google Crash Course on ML: Embeddings

<https://msql.co/google-embeddings>


MemSQL Face Matching Blog

<https://msql.co/Face-matching>

Thank You! Questions?

Try at memsql.com/download

- No time limit
- Deploy to production
- Full featured
- Up to 4 Nodes, unlimited disk
- Get support at forum.memsql.com



Learn more at memsql.com/product

Questions? Email us at team@memsql.com