

# GOODBYE, DATA LAKE:

WHY CONTINUOUS ANALYTICS YIELD HIGHER ROI

STAC SUMMIT LONDON, 2018  
ORI MODAI, VP R&D @IGUAZIO





# Business Challenges in The Digital World

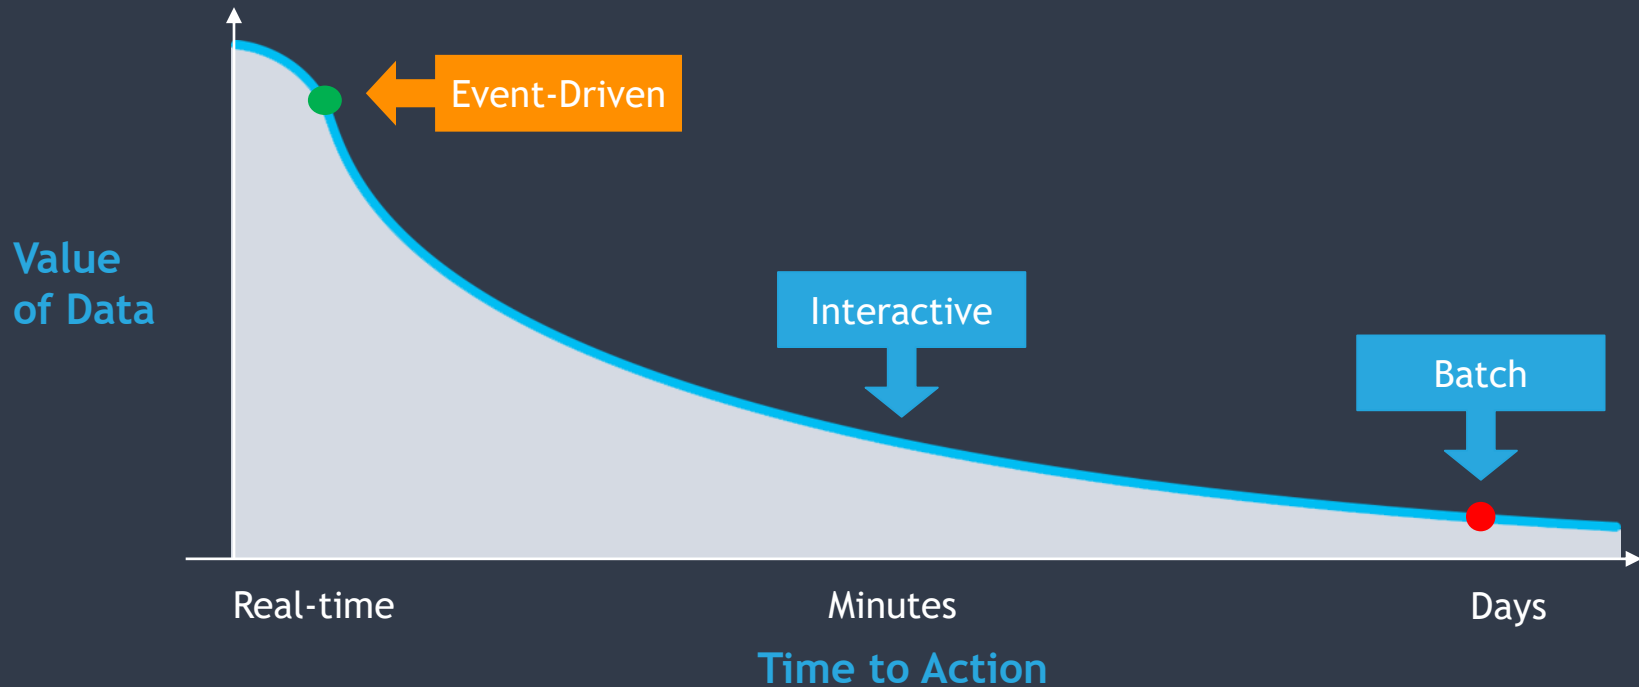


- Fast and intelligent data-driven actions
- Rapid and continuous innovation
- Building for scale and distribution



# The Data-Driven Business Challenge

## From Reactive to Proactive






# Evolve Into a Future Proof Cloud-Native Architecture

## Once upon a time there was a data lake

(You spent years feeding it with DevOps)



## Meanwhile in the cloud



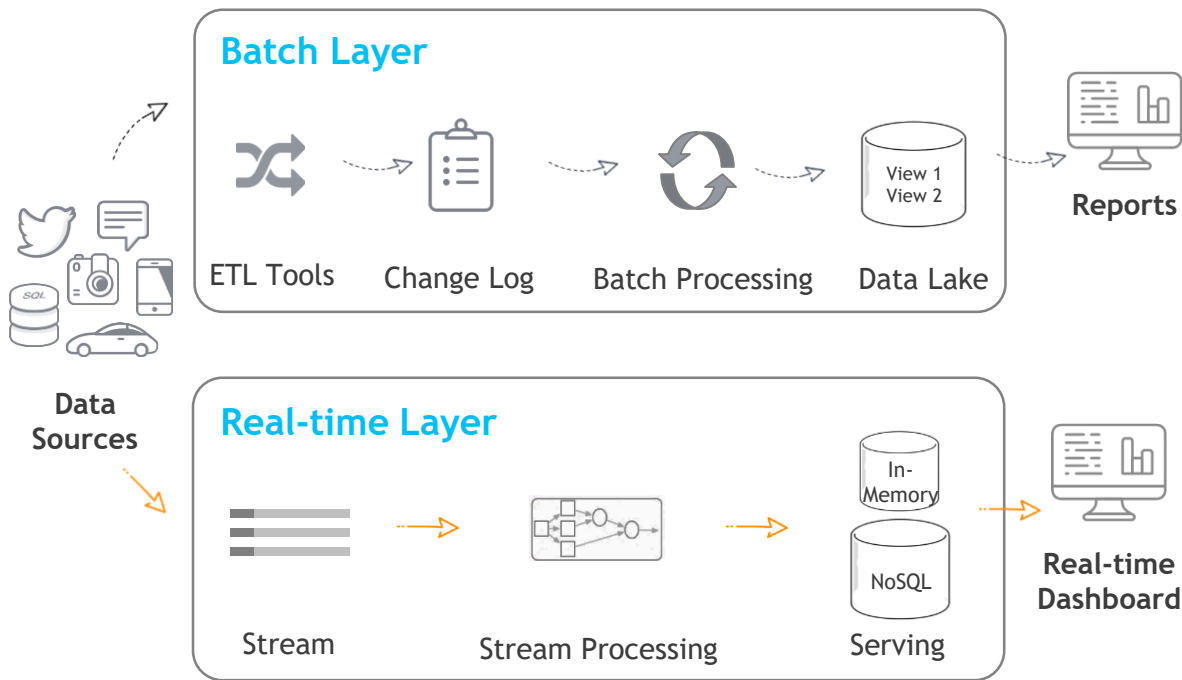
Innovate



Consume



# Big and Slow or Small and Fast



## Too slow

- Big data but slow
- Not up to date
- Complex

OR

## Limited context

- Small amounts of data
- Expensive
- Lacks context



# Addressing The Digital Transformation Challenges

---



## Making AI Real-time

*“Real-time data integration and access to data remain core challenges for data and analytics leaders looking to modernize data management ecosystems.” - Gartner*



## One Solution Across Cloud and Edge

*“Real-time decision-making and interaction requirements; growth in data being produced at the edge; and requirements for autonomy, security and privacy will increase the percentage of compute and storage capability that's closer to people, things and the edge.” - Gartner*

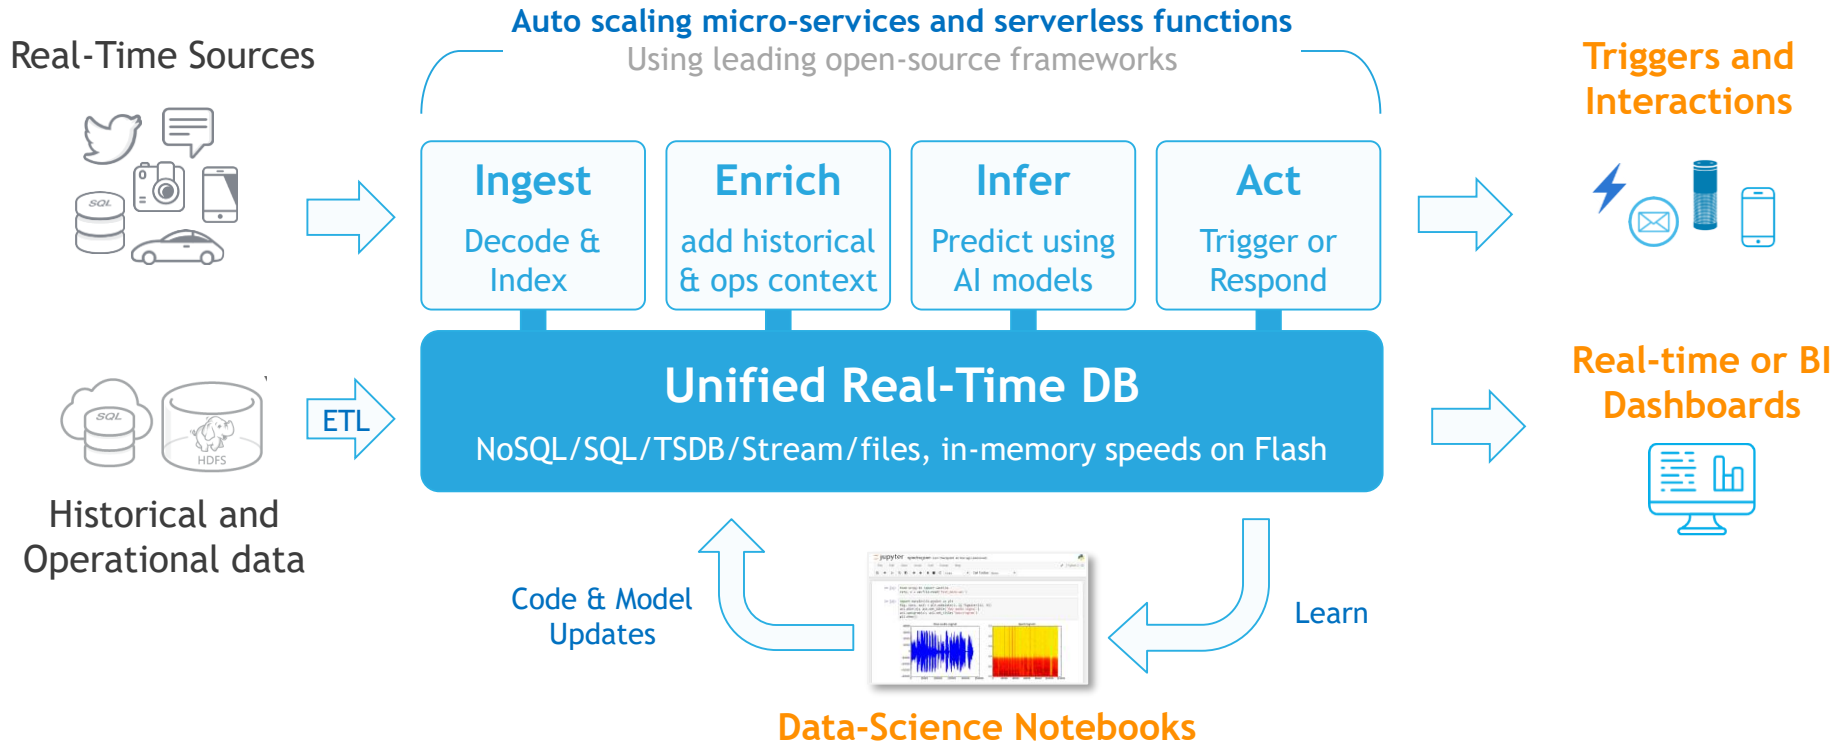


## Simplified Dev&Ops with Serverless

*“No need to spend time and resources on server provisioning, maintenance, updates, scaling, and capacity planning. Instead, all of these tasks and capabilities are handled by a serverless” - CNCF Serverless Paper*

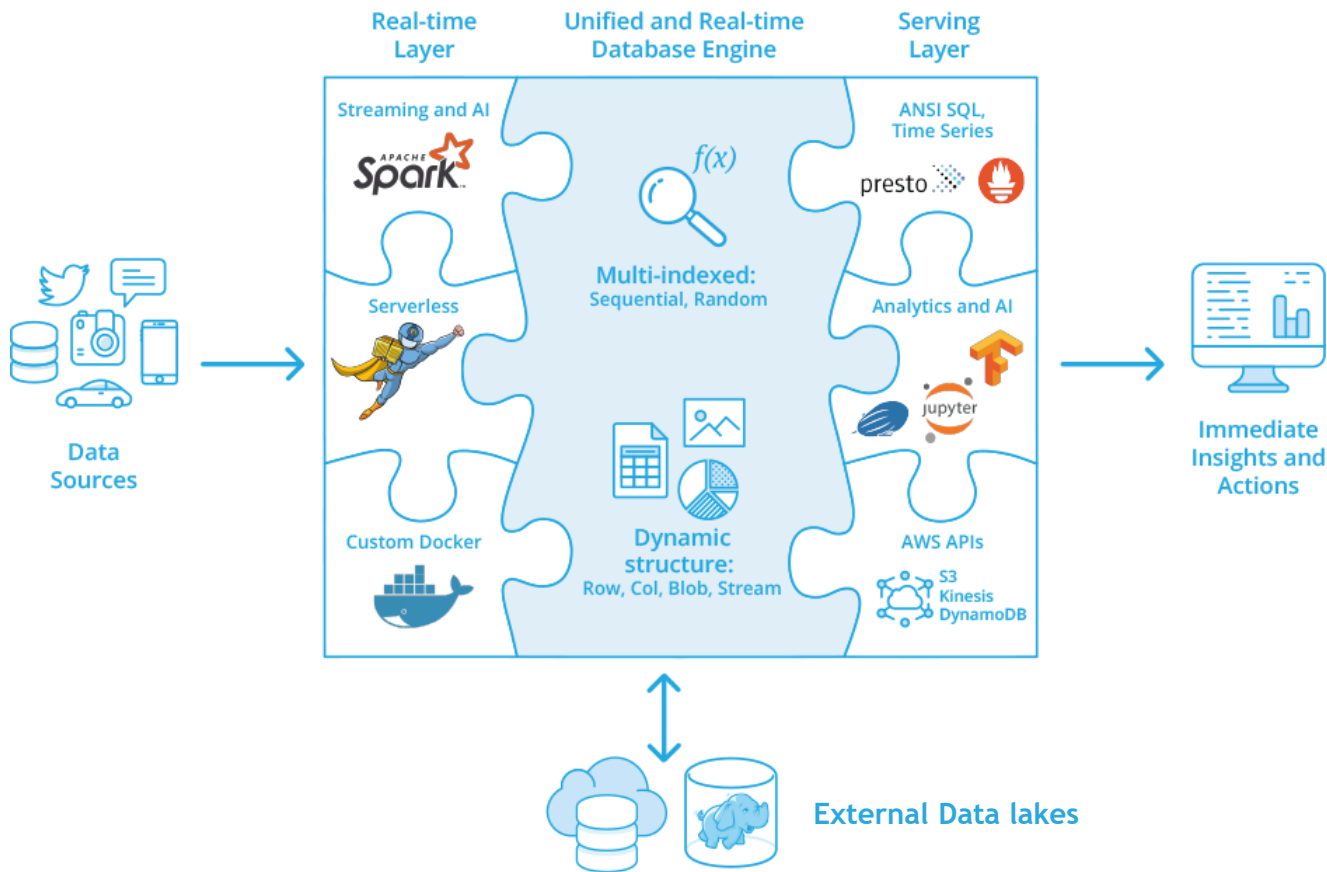


# The Real-Time AI Pipeline






# Ingest, Enrich, AI, and Serve on One DB Engine





# How To Deliver Volume, Velocity and Variety ?

## Traditional Layered Approach



External (NVMeOF / Object)

- Slow
- Complex
- Expensive

## Optimized Approach



Decouple the APIs from the DB Engine

Use NVMe Flash as an extension of OS Memory

- In-memory speed
- Simple
- 1/3<sup>rd</sup> the TCO



# REAL-TIME DATA PIPELINES

---

CUSTOMER USE-CASES



# Latency Prevention for Electronic Trading Platforms

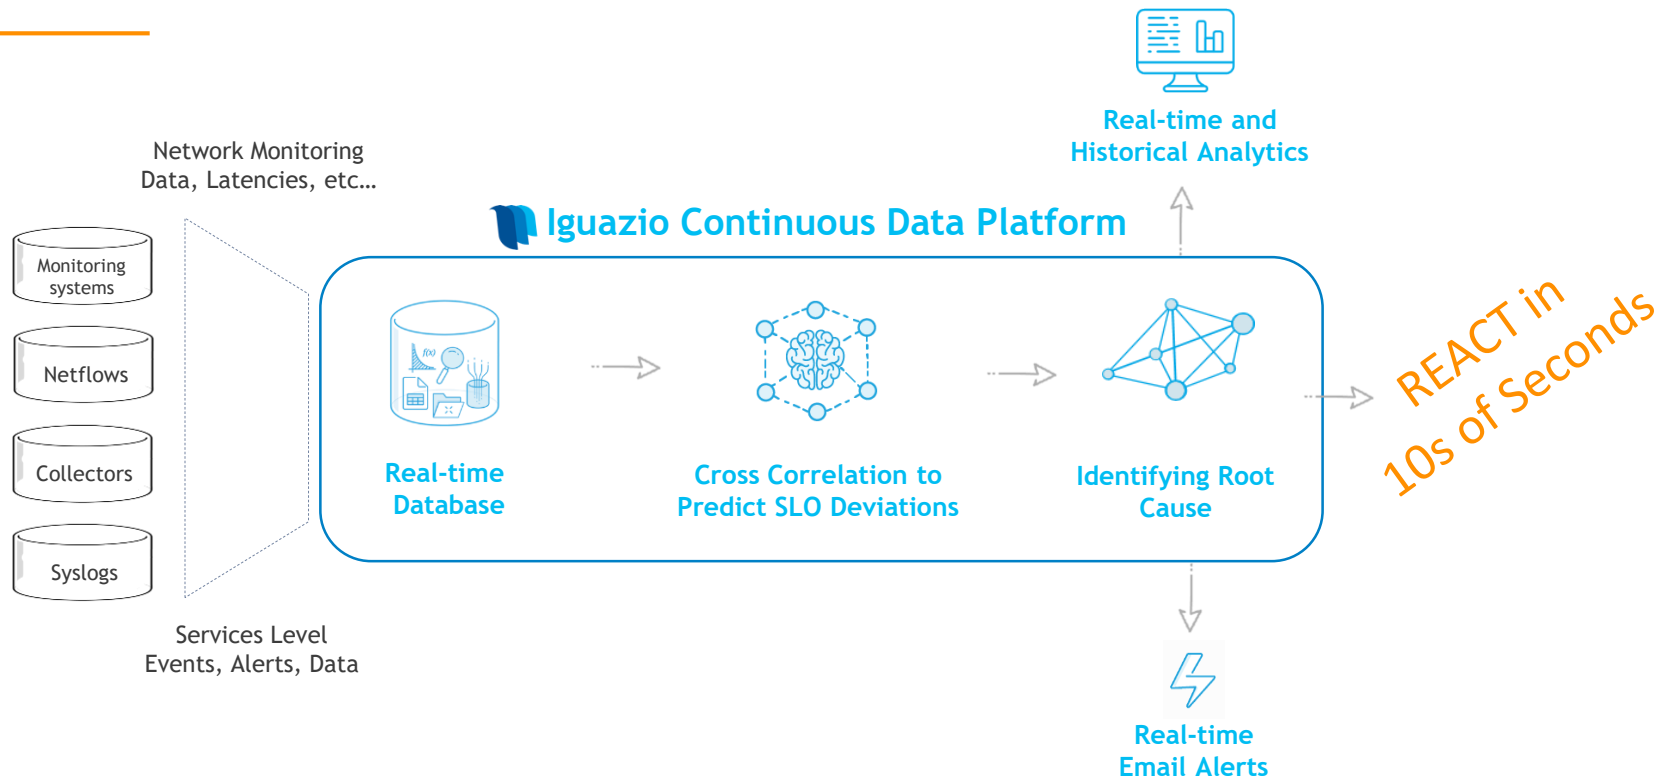
Global banks deployed Iguazio to predict and avoid latencies in their electronic trading platforms

- Replaced traditional latency monitoring solutions - enforcing trading latencies SLO
- Iguazio has changed it to:
  - Real-time predictions latencies trend
  - Root cause analysis via multivariate analysis
  - Actions triggered via serverless functions
  - Trading stays continuous and uninterrupted
- Outcome - Improved SLOs, reduced trading down times





# Data Flow

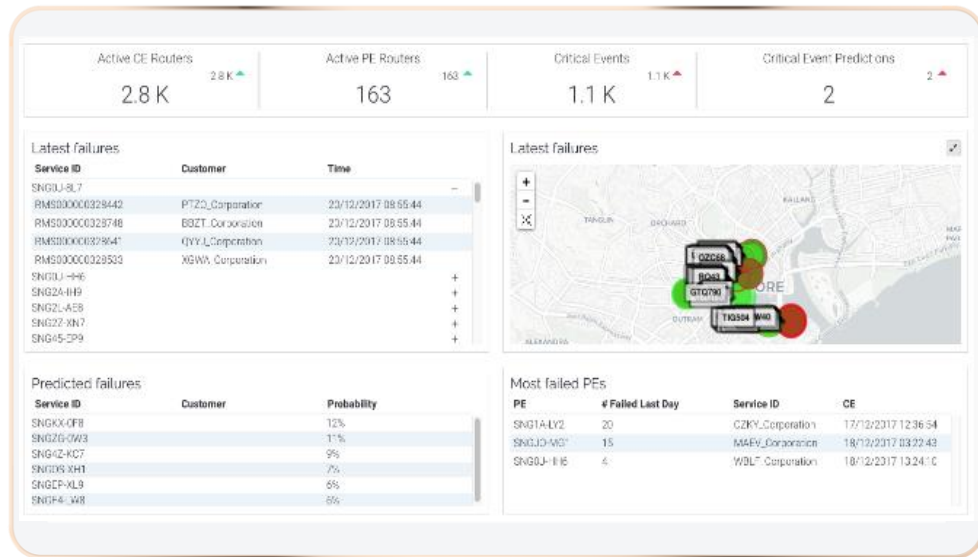




# Auto Healing Network Operations

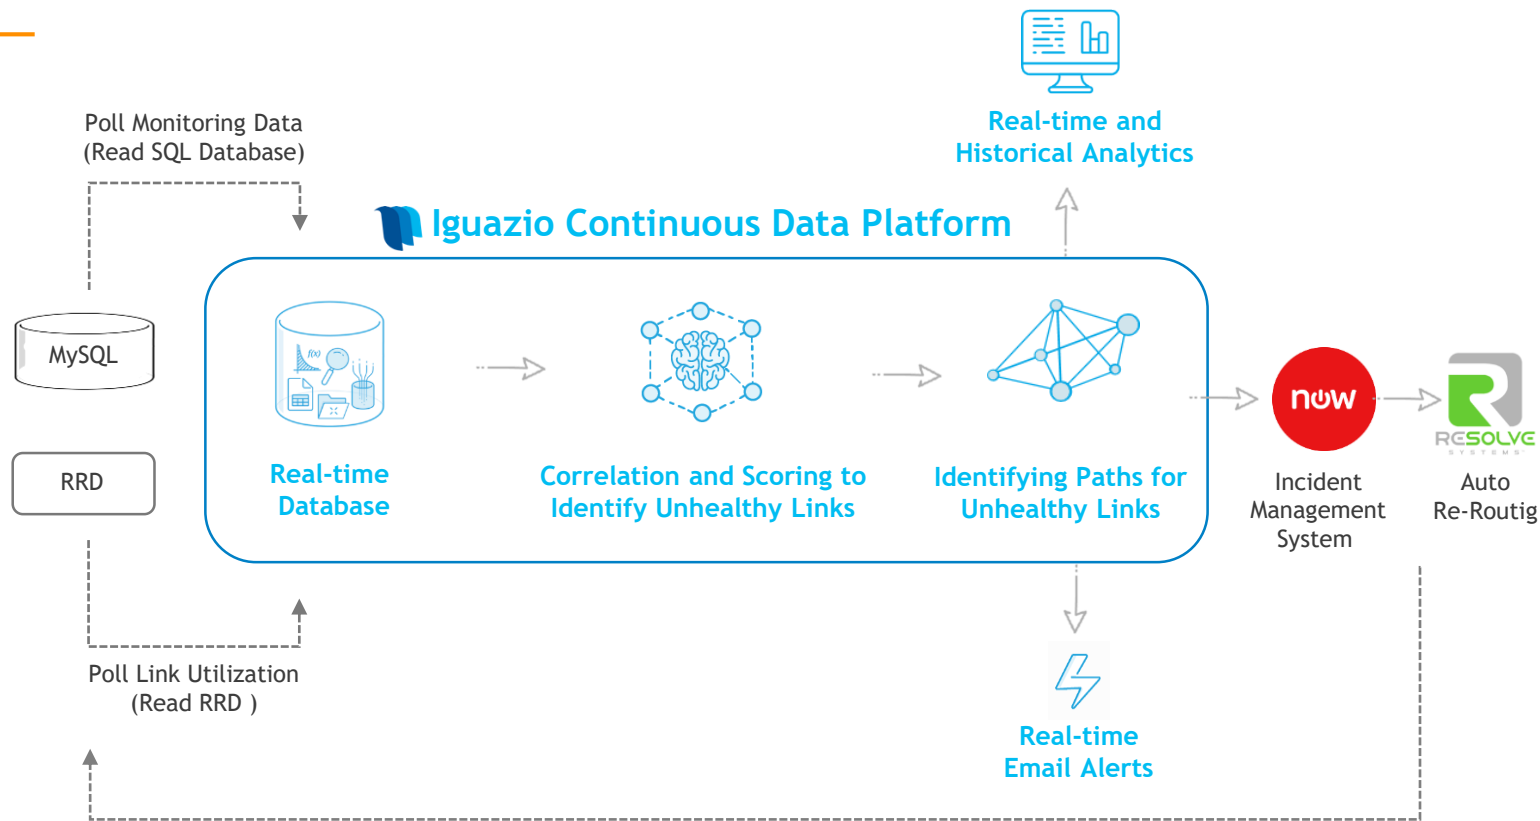
## Global network operator uses Iguazio to predict and avoid network outages in real-time

- Replaced a Hadoop data pipeline (that was never productized)
- Real-time cross correlation of data from multiple sources
- AI based predictions triggering pre-programmed network optimization algorithms
- Time to production < 4 weeks





# Data Flow



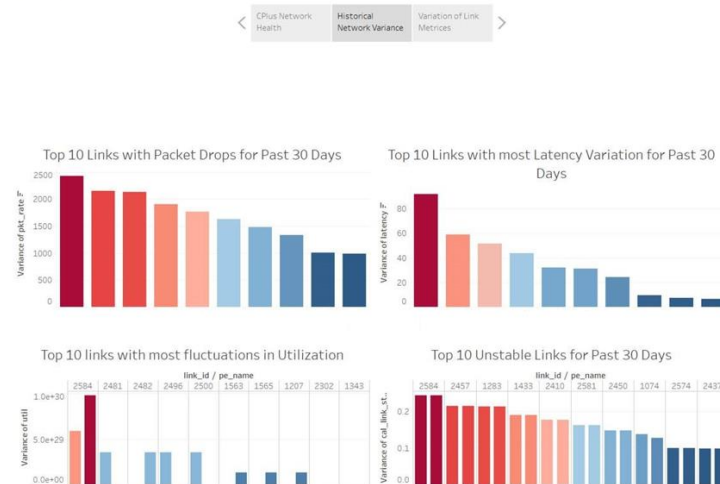


# Real time dashboards

Real-time Link Quality Dashboard



Real-time Link Quality Dashboard

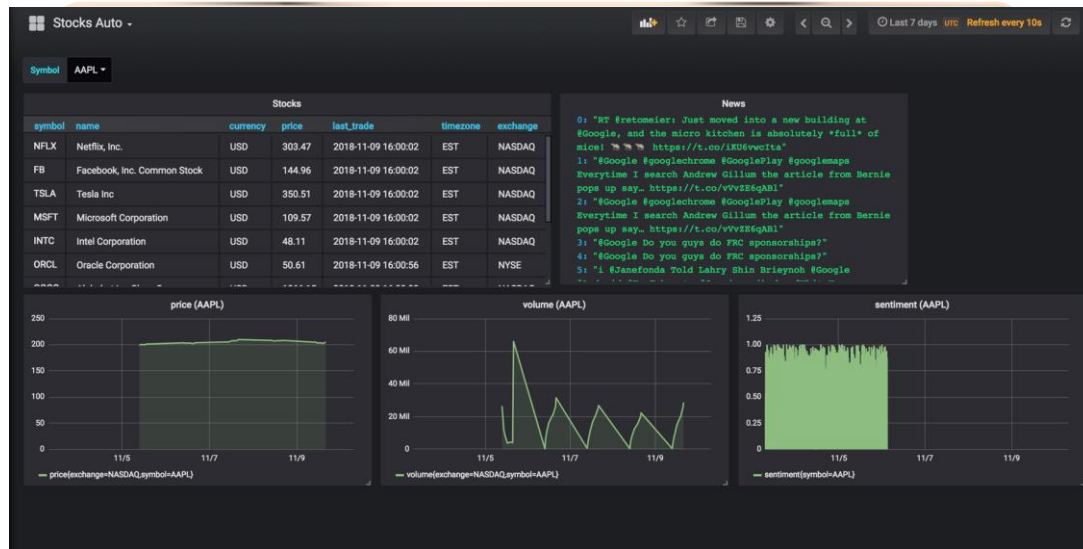




# Multivariate Real-time Analysis of Trade Data


## Major financial institute instruments its trade data ML research with Iguazio

- Orchestrating multiple data feeds and data modalities
- Multi dimensional feature vectors aggregation
- Enabling ML research activities on operational recent data
- Advanced serving and dash-boarding from the same platform
- Time to POC ~ 2 weeks





# Real-Time Analysis of Trade Data





# Real-Time Analysis of Trade Data



Real-time functions for  
Ingestion, decoding and tagging



Tweet  
Sentiment  
Analysis

**N**  
News API



News  
Analysis  
& Tagging



Tick feed  
Analysis  
& Tagging

World Trading Data

News stream  
viewer



  
Real-time Dashboard



  
Real-time Analysis

```
Explore Real-Time Data on Iguazio via Integrated Notebooks
Analyze Real-Time Data Using Spark Streaming, SQL, and ML
Data access to Spark on Iguazio DataFrames

1. Run queries on Iguazio DataFrames
2. Run queries on Iguazio DataFrames
3. Run queries on Iguazio DataFrames
4. Run queries on Iguazio DataFrames
5. Run queries on Iguazio DataFrames
6. Run queries on Iguazio DataFrames
7. Run queries on Iguazio DataFrames
8. Run queries on Iguazio DataFrames
9. Run queries on Iguazio DataFrames
10. Run queries on Iguazio DataFrames
```

**Iguazio Unified Real-Time DB Service**  
SQL | NoSQL | TSDB | Streams | Obj/files

Functions/Model  
Dev & Test



deploy



# Summary

---

## Build continuous, data-driven and proactive apps

- Deliver real-time analytics based on fresh and historical data
- Utilize Flash to deliver in-memory speed at much lower costs
- Create a unified data layer for stream processing, AI and serving
- Adopt cloud-native and serverless approaches to gain agility
- Enable your data scientists with access to fresh, large scale data





# Thank You

---

[orim@iguazio.com](mailto:orim@iguazio.com) | [www.iguazio.com](http://www.iguazio.com)