Goodbye, Data Lake: Why Continuous Analytics Yield Higher ROI

June 2018

The Data-Driven Business Challenge

From Reactive to Proactive

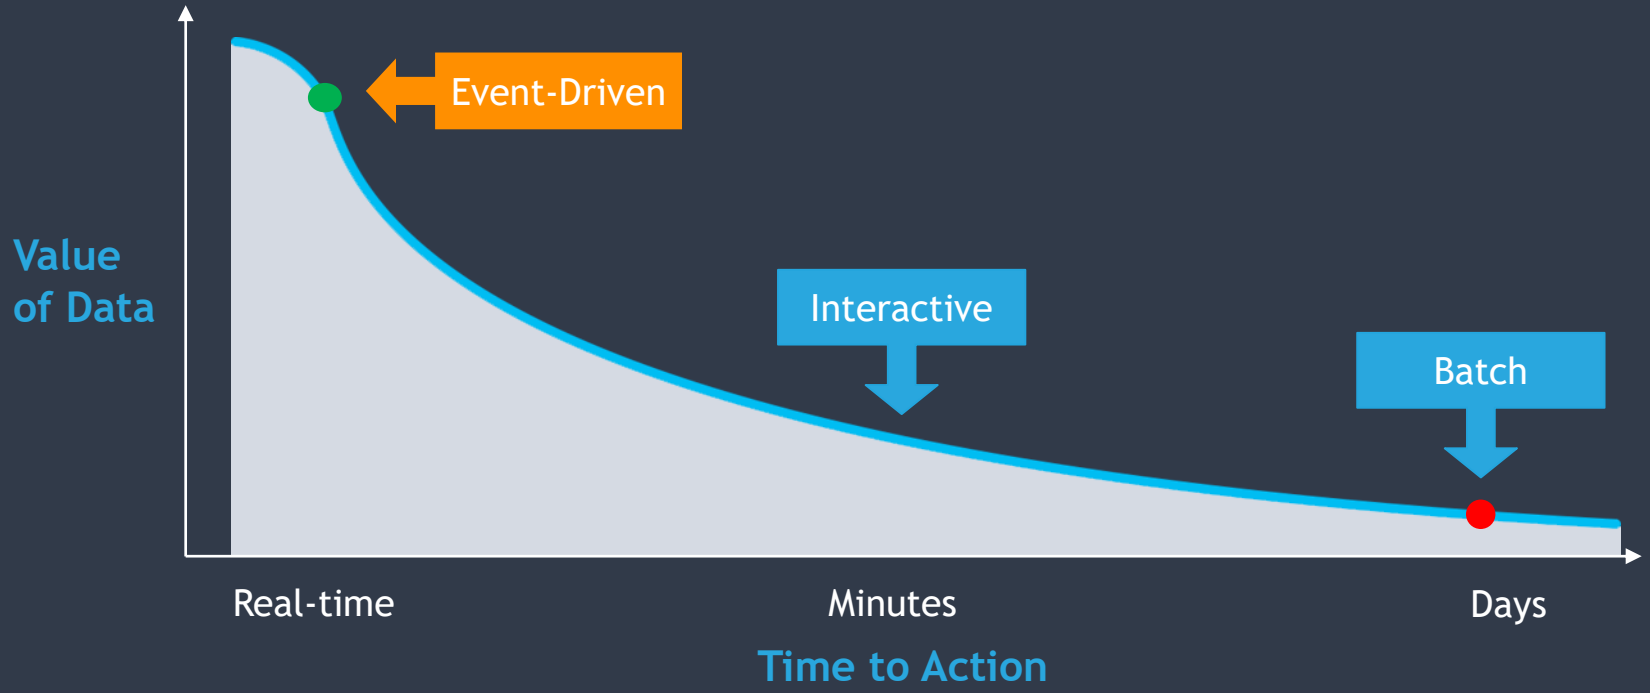

iguazio

Big and Slow or Small and Fast

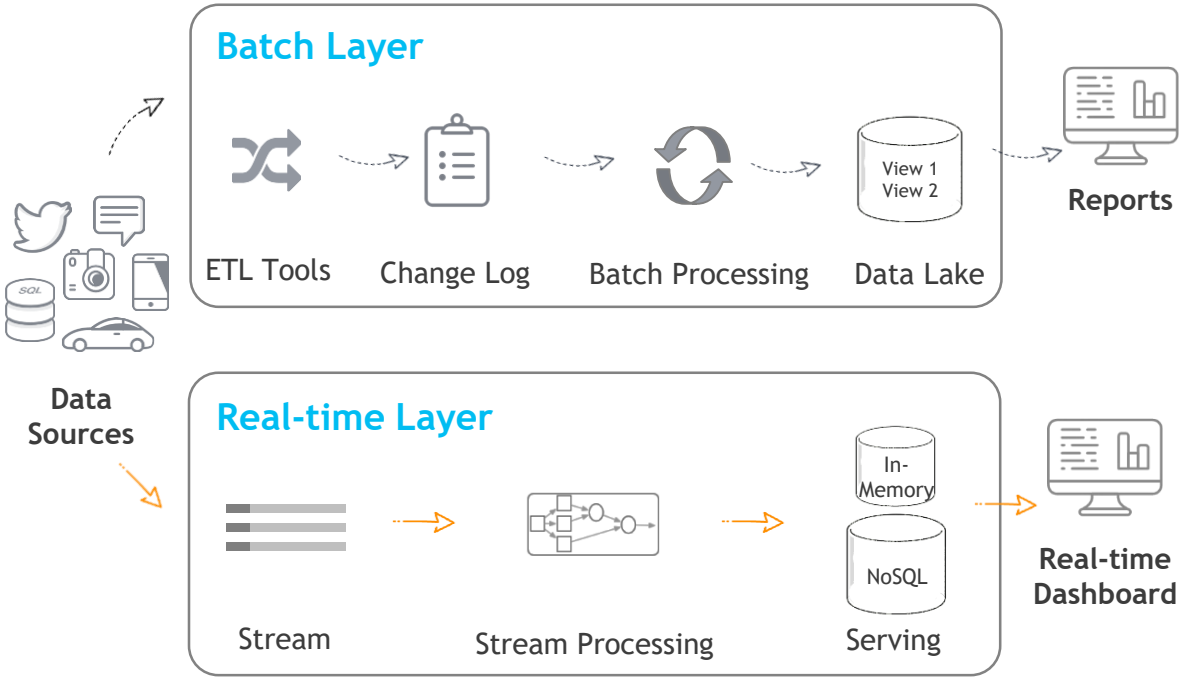

Too slow

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

OR

Limited context

- Small amounts of data
- Expensive
- Lacks context

iguazio

How To Deliver Volume, Velocity and Variety ?

Traditional Layered Approach

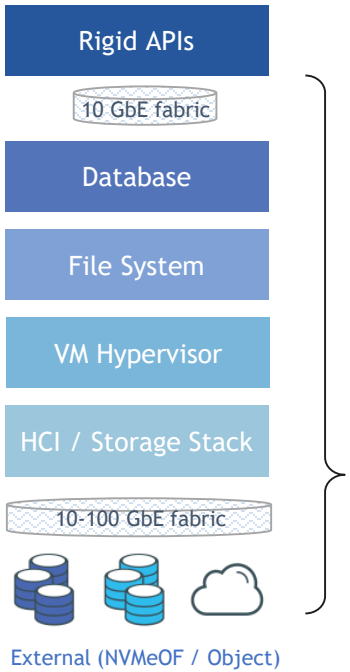

- Slow
- Complex
- Expensive

Optimized Approach

Decouple the APIs from the DB Engine

Use NVMe Flash as an extension of OS Memory

- In-memory speed
- Simple
- 1/3rd the TCO

iguazio

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

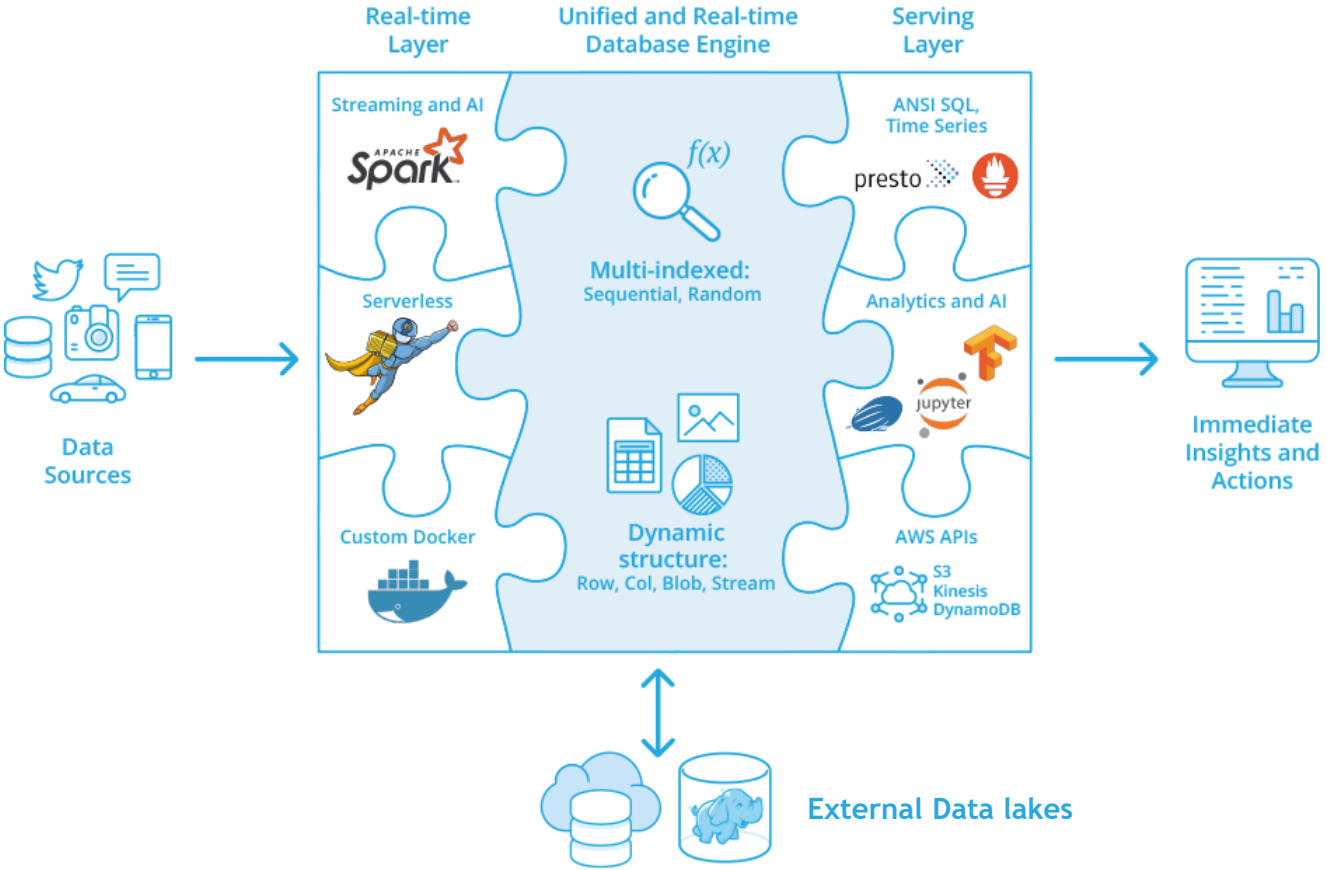

iguazio

Serverless, Eliminating 80% of The Work

Traditional Dev and Ops Model

- **Write code + local testing**
- Build code and Docker image
- CI/CD pipeline
- Add logging and monitoring
- Harden security
- Provision servers + OS
- Handle data/event feed
- Handle failures/auto-scaling
- Handle rolling upgrades
- Configuration management

80%

"Serverless" Development Model

- **Write code + local testing**
- **Provide spec, push deploy**

1. Automated by the serverless platform
2. Pay for what you use

Open-Source Serverless, A Simpler Lock-free Alternative

Laptop

Cloud

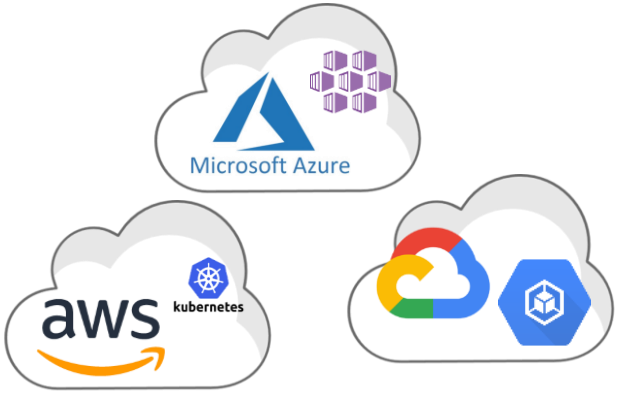

On-Prem or Edge

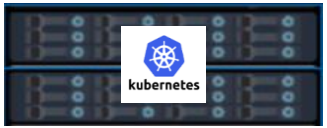

Same APIs, Same User Experience, Anywhere

With native integration into each cloud platform

iguazio

Addressing Serverless Limitations With Nuclio

Performance

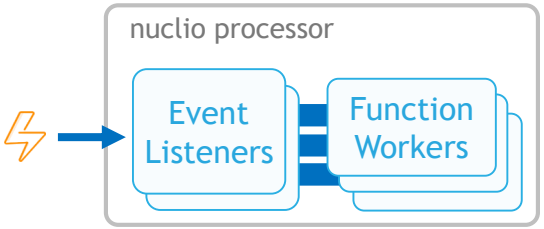

- Non-blocking, parallel
- Zero copy, buffer reuse
- Up to 400K events/sec/proc

Streaming and Batch

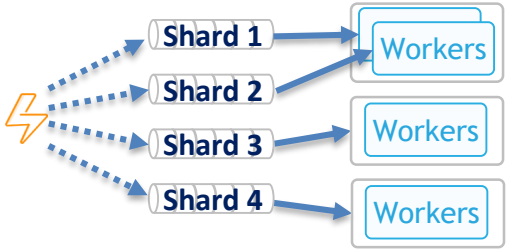

- Auto-rebalance, checkpoints
- Any source: Kafka, NATS, Kinesis, event-hub, iguazio, pub/sub, RabbitMQ, Cron

Statefulness

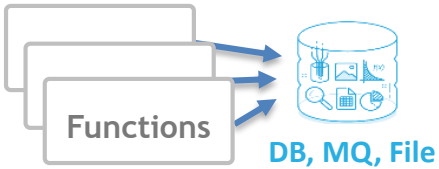

- Data bindings
- Shared volumes
- Context cache

Serverless for compute and data intensive tasks
100x faster than AWS Lambda !

iguazio

Evolve Into a Future Proof Cloud-Native Architecture

Once upon a time there was a beast …

You spent years feeding it with DevOps

And then it became cloudy !

Middleware

Orchestration

Data

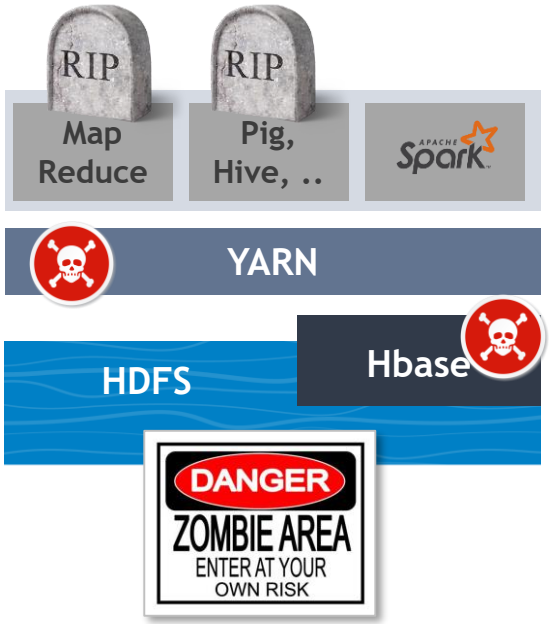

Innovate

Consume

iguazio

Delivering Intelligent Decisions in Real-Time

500TB of Raw Data

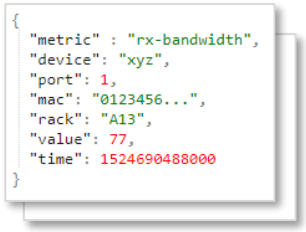

Ingested in real-time (compressed to **10TB**)

External Context

ML Models

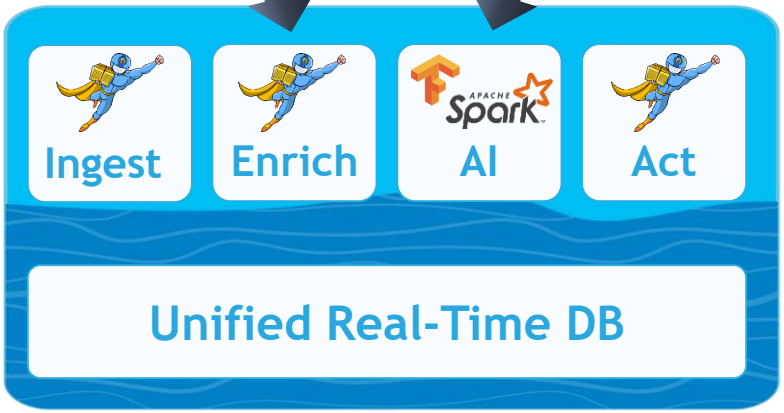

Real-time triggers

Serve

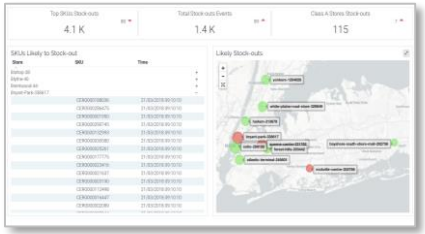

Real-time and historical dashboards

Zero Ops:

Via DBaaS, AIaaS, and Serverless

iguazio

Cyber and Network Ops

A leading telco needs to predict network behavior in real-time:

- Processing high message throughput from multiple streams at the rate of > 50K events/sec
- Cross correlating with historical and external data in real-time
- AI predictions/inferencing conducted on live data
- Small footprint to fit network locations

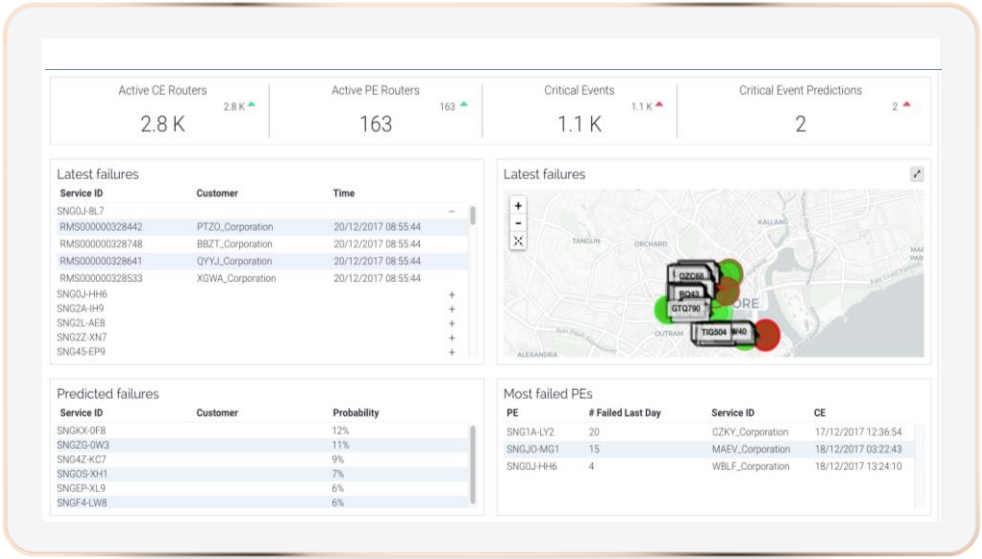

iguazio

Build and Operationalize Proactive Systems Faster

Traditional

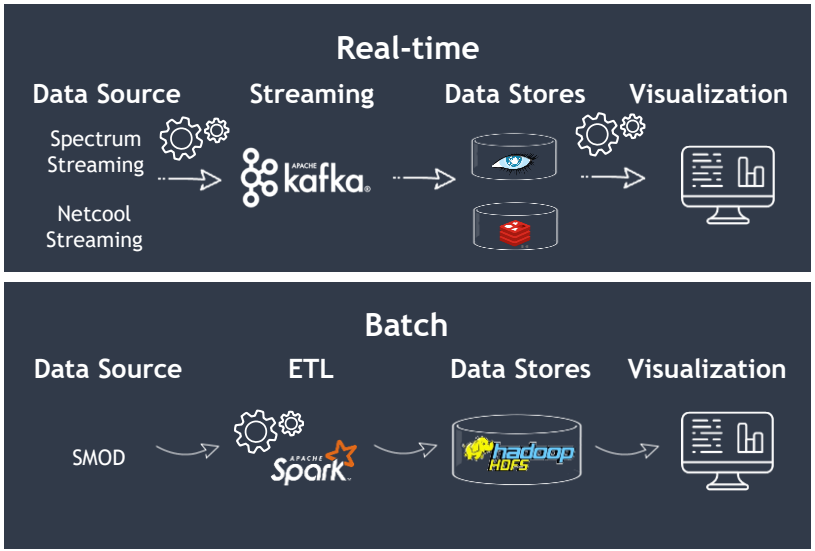

- Complex, skill gaps, slow to productize
- No single view of ops, real-time, history
- Reactive (no actions)

Continuous Analytics

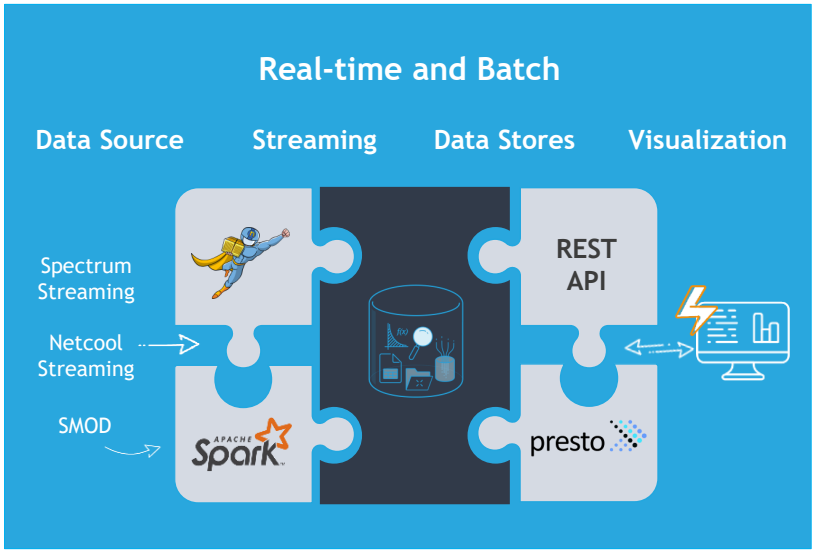

- Simple, just a few weeks to a working app
- Unified view across ALL data
- AI driven, proactive

iguazio

Demo: Voice Driven Real-Time Analytics

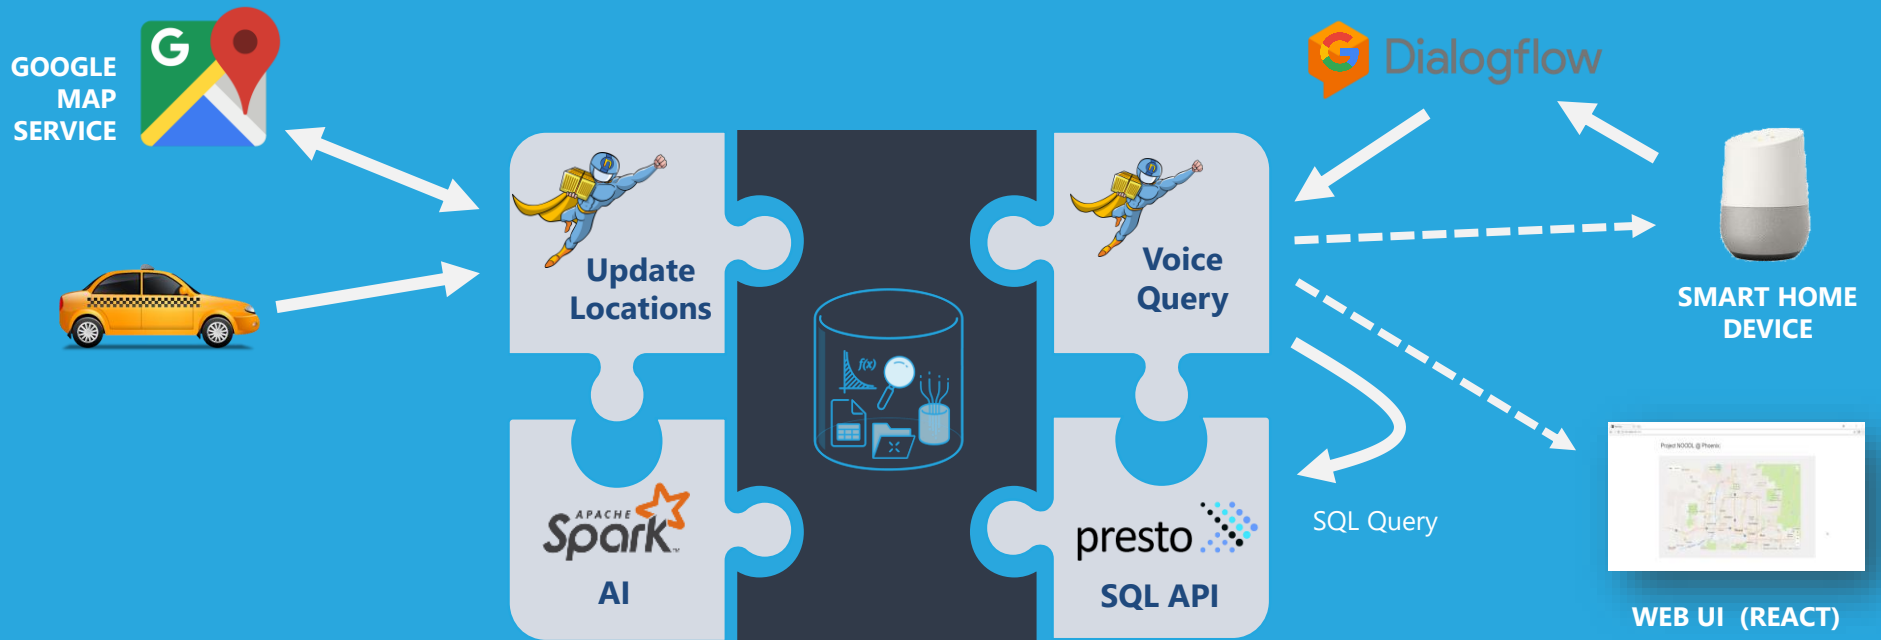

Project NOODL @ Phoenix

Summary

Build continuous, data-driven and proactive apps

- Deliver real-time analytics on fresh, historical and operational data
- Optimize Flash usage 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

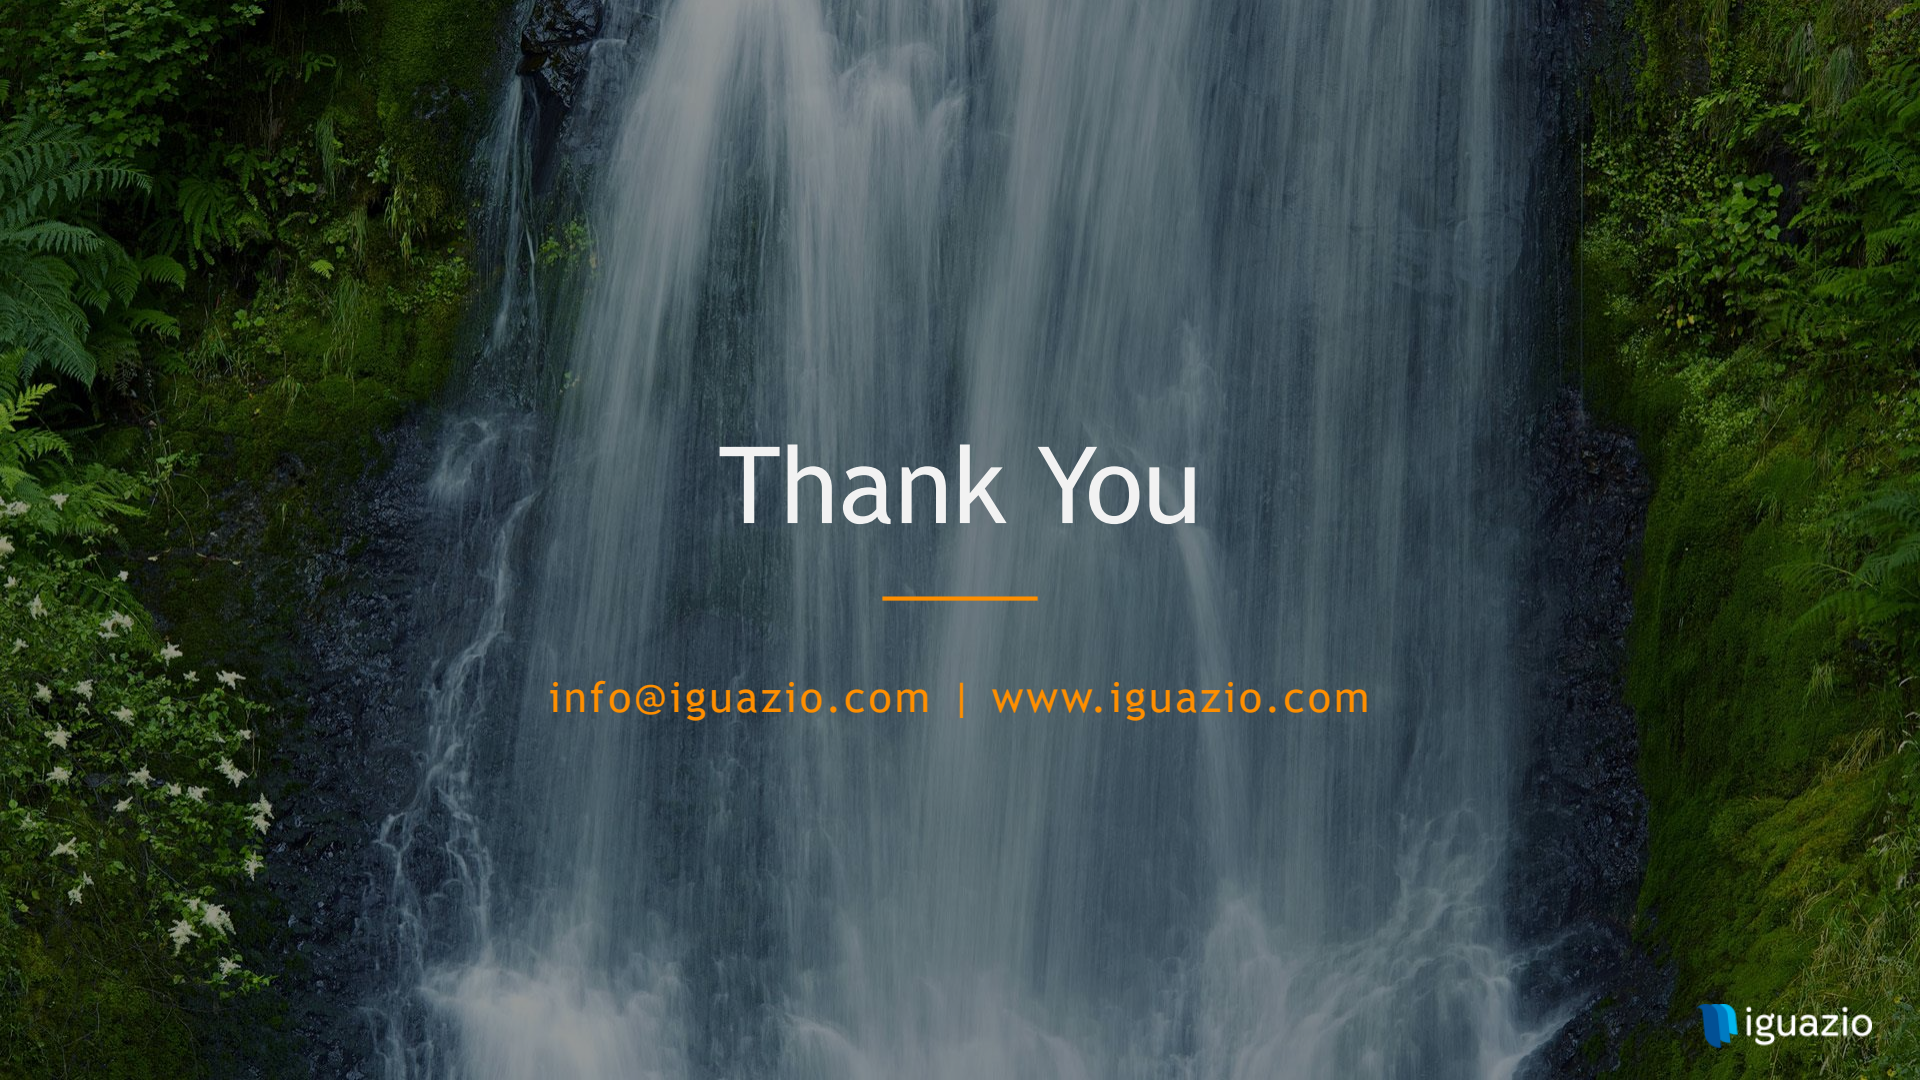
Thank You
info@iguazio.com | www.iguazio.com
iguazio