In-mem DB Performance, Flash Cost Enabling Real-time AI

June 2018

The Data-Driven Business Challenge

From Reactive to Proactive

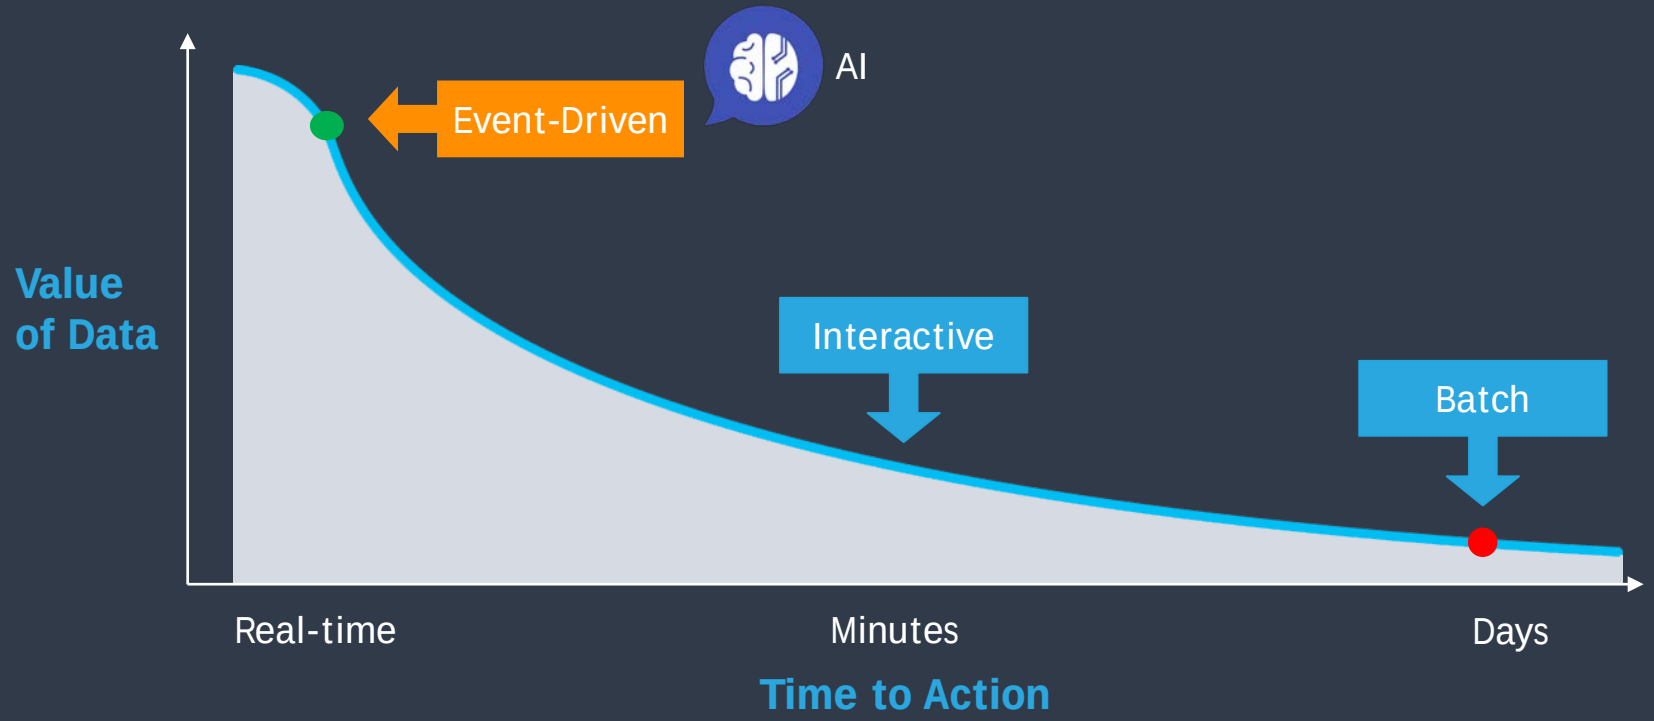

iguazio

Big and Slow or Small and Fast

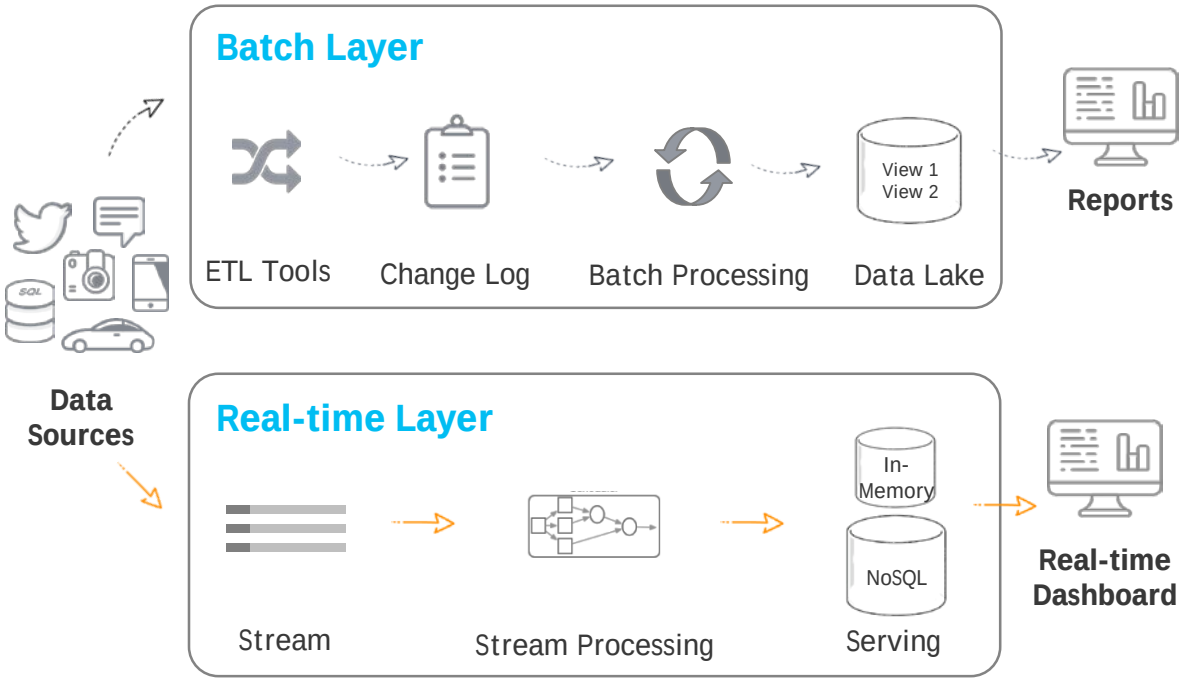

Too slow

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

OR

Limited context

- Small amounts of data
- Expensive
- Lacks context

iguazio

Traditional Approach, DB over File over Flash

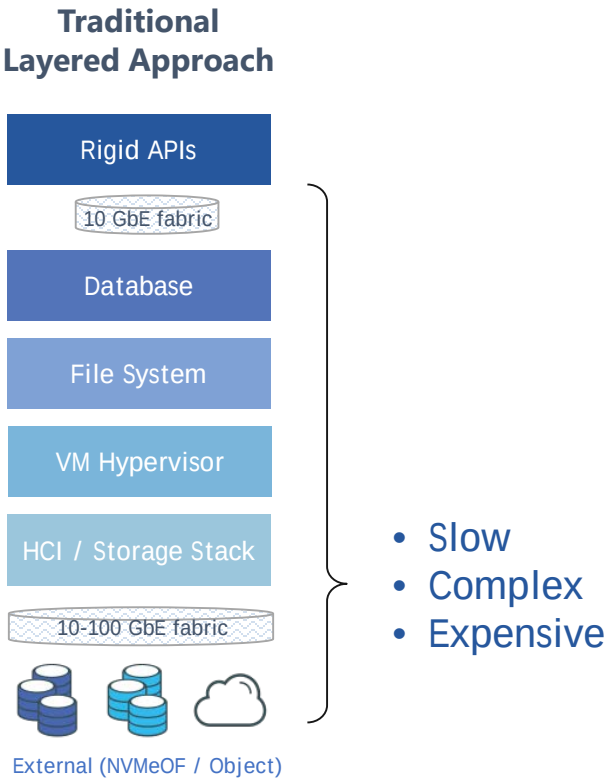

- Slow
- Complex
- Expensive

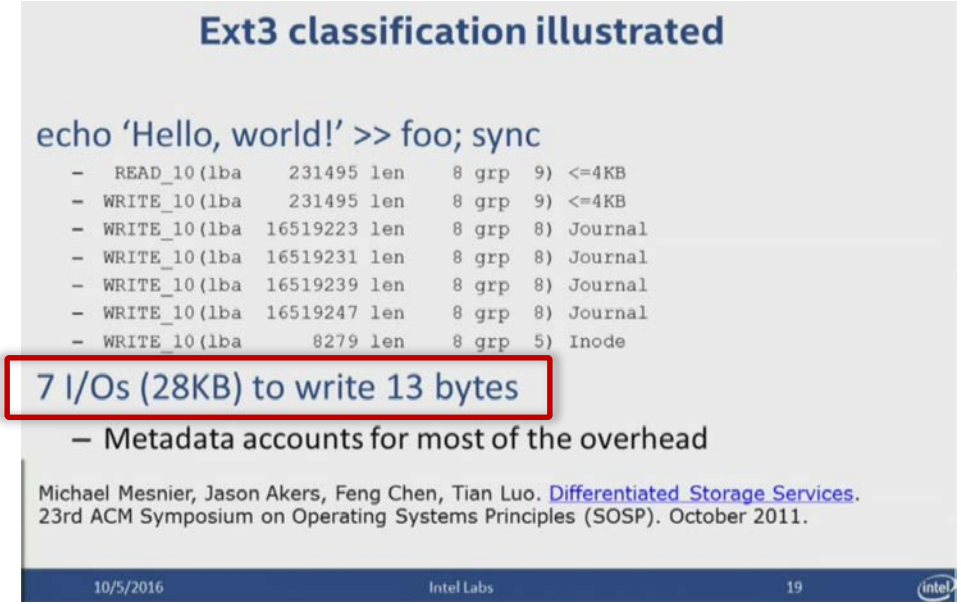

For every file IOs conducted by the DB
(Record, Redo/Undo, Metadata, ..)

iguazio

New Cloud Databases Are Built to Scale Ops & Capacity

Current DB architectures are monolithic

Shared Disk

Coupled at the caching & storage layer

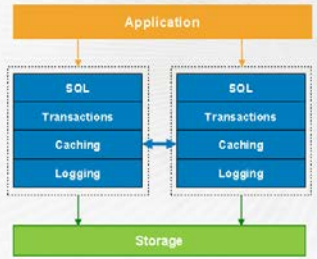

Even when you scale it out, you're still replicating the same stack

Amazon Aurora

- Service-oriented architecture applied to the database
- Moved the logging and storage layer into a multi-tenant, scale-out database-optimized storage service
- Integrated with other AWS services like Amazon EC2, Amazon VPC, Amazon DynamoDB, Amazon SWF, Amazon Route 53 for control plane operations
- Integrated with Amazon S3 for continuous backup

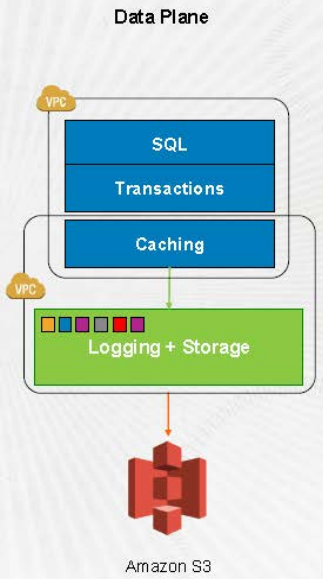

API & Transaction

Distributed Processing & Cache

Capacity (Object)

Decouple access, processing, and capacity and eliminate storage serialization

iguazio

Breaking The Volume and Velocity Barrier

Support many standard APIs on a common DB Engine

Unique architecture which use NVMe Flash as an extension of OS Memory

Re-engineer the stack to deliver memory speed with Flash density

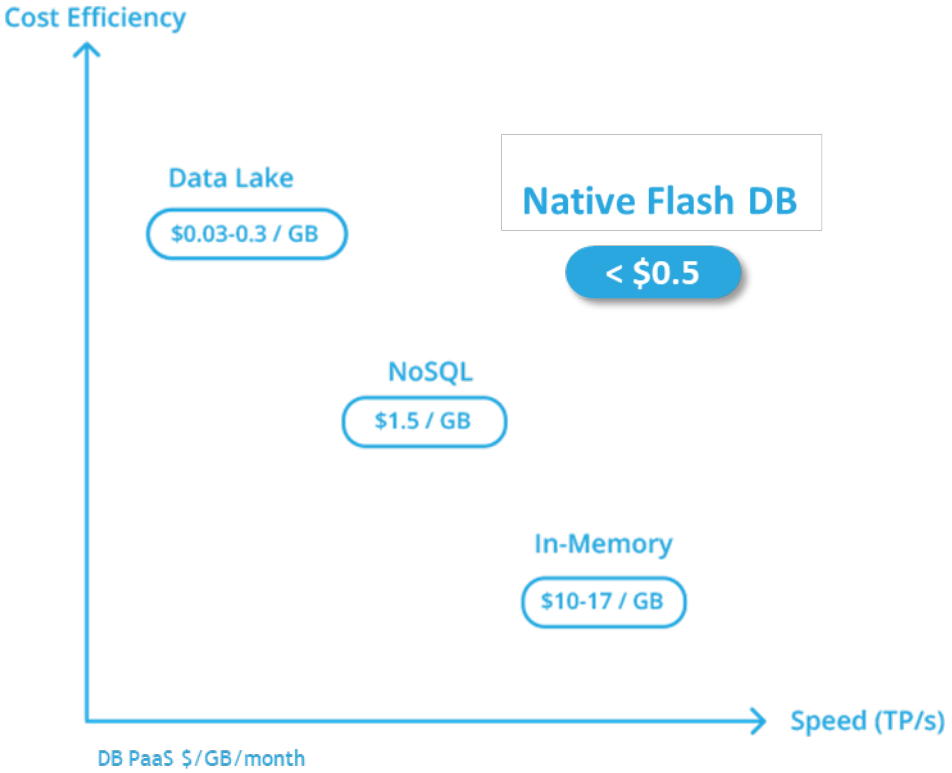

iguazio

Breaking Performance Barriers – Design Principles

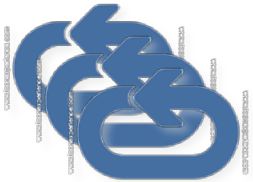

Never blocking, never locking, 100% parallelism
Latency optimized, QoS aware, data scheduler
Lockless, preempt less memory management
True scale out through parallelism

Zero processing wastes
CPU cache optimization and prediction
E2E zero buffer data flow (NIC to Disk, accelio)
Complete OS bypass

HW awareness
RDMA, NVMe (3DXP)
Vector processing operations
IRQ balancing and throttling

Ok, any other challenges on the way to real-time AI ?

90% of AI Today

Build feature vectors using batch and CSVs

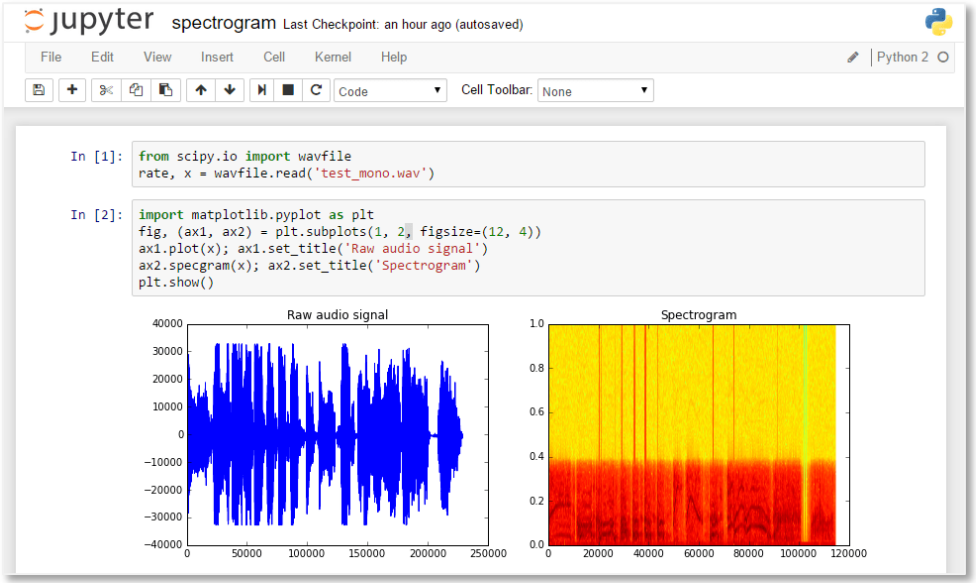

Inspect, Improve

How do we form complex feature vectors in real-time?

How do we visualize or act on the results in real-time?

iguazio

Moving to Continuous Ingest + AI + Serve Flow

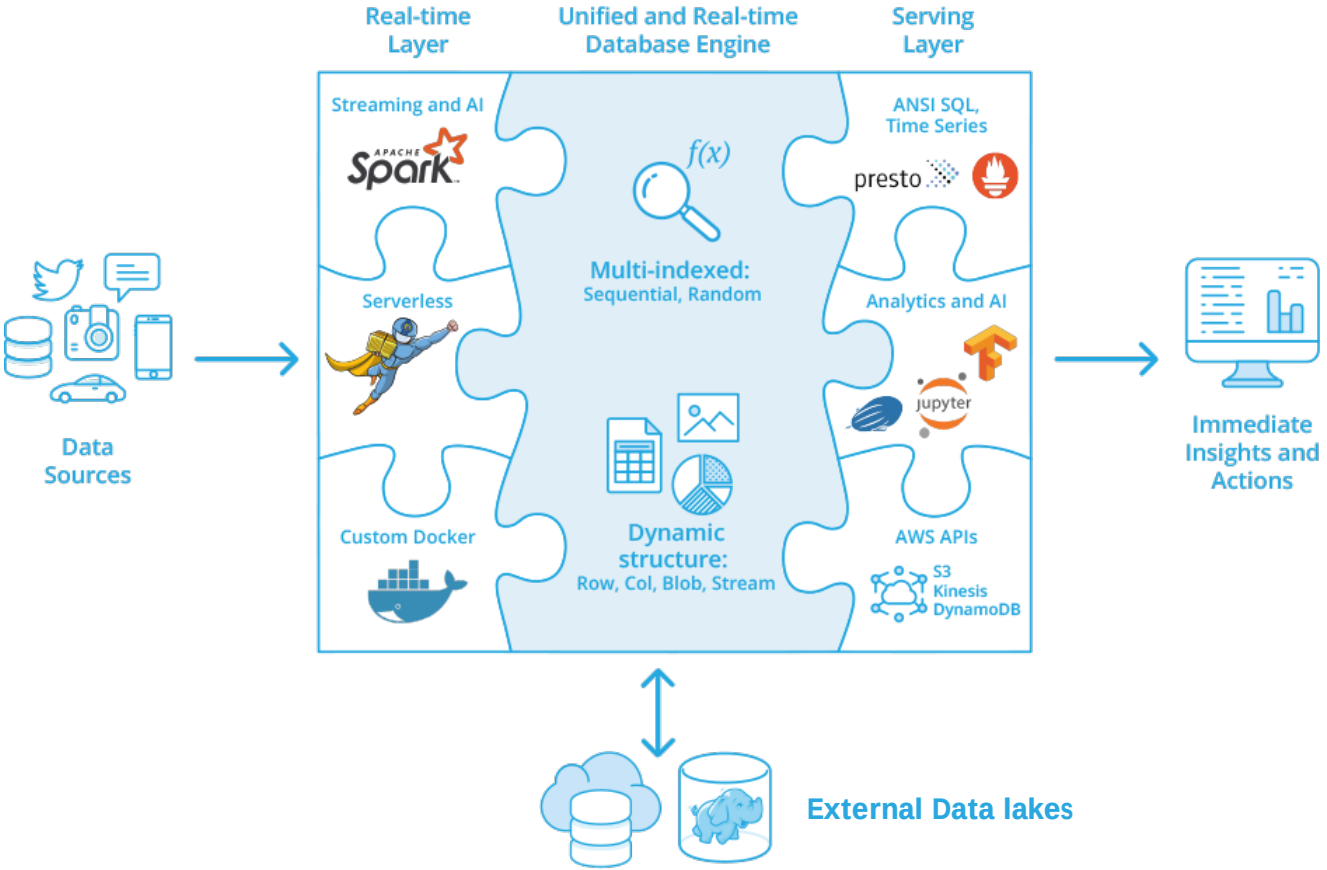

iguazio

From Silos and ETLs to All-in-one DBs

Traditional: Unique Model Per Store

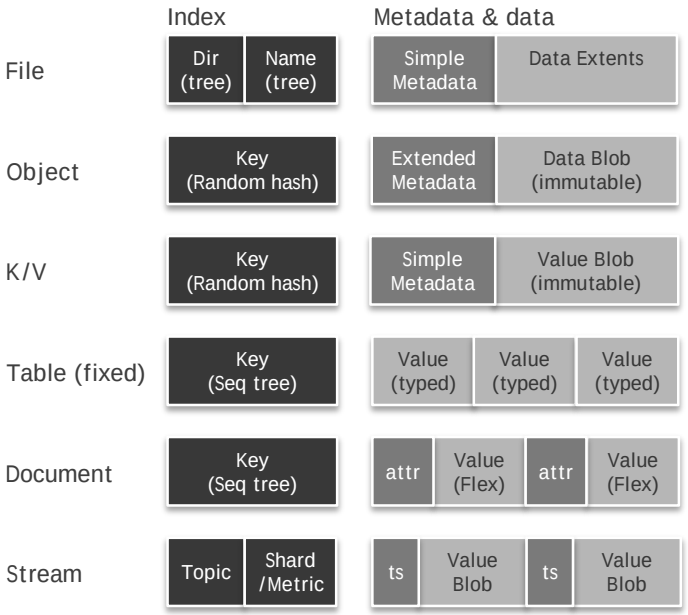

Multi-Model Store

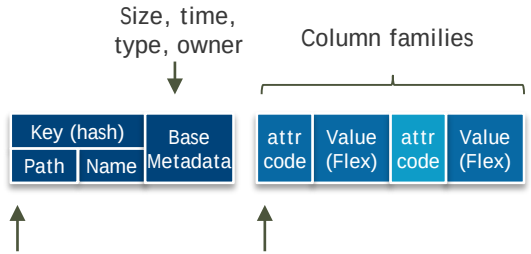

Multiple Indexes

Random, sequential and hierarchical

Any Data Type

- Nested attributes (encoded)
- Flexible value types
- Can be organized and viewed as extents, rows, cols or logs

Independent tiering logic for indexes, metadata and data

iguazio

Time Series Data Example

Raw time series sample data

Thousands of samples

Labels

Data

```
{
  "metric" : "rx-bandwidth",
  "device": "xyz",
  "port": 1,
  "mac": "0123456...",
  "rack": "A13",
  "value": 77,
  "time": 1524690488000
}
```

Ingest / compress
In real-time

Optimized TSDB Layout (per unique metric)

```
{
  "__name__" : "rx-bandwidth",
  "device": "xyz",
  "port": 1,
  "mac": "0123456...",
  "rack": "A13",

  "_v_count": [...],
  "_v_sum": [...],
  ...

  "_v0": <compressed blob>,
  "_v1": <compressed blob>,
  ...
}
```

Filter based on labels

Pre-aggregation arrays:
(to accelerate queries)

T/V chunks with 10:1
Gorilla compression

Real-time Consistency	50 : 1 Compression	10–100x Faster Queries

Serverless, The New Stored Procedure

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

iguazio

Addressing Serverless Limitations With Nuclio

Performance

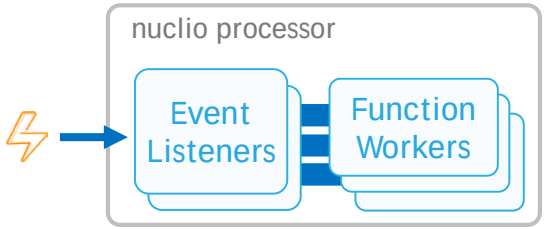

- Non-blocking, parallel
- Zero copy, buffer reuse
- Up to 400K events/sec/proc

Streaming and Batch

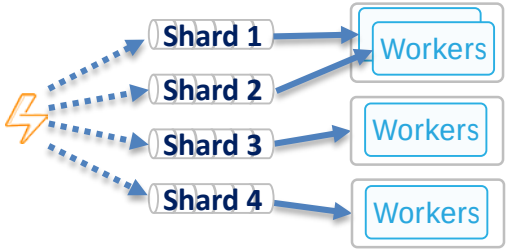

- Auto-rebalance, checkpoints
- Any source: Kafka, NATS, Kinesis, event-hub, iguazio, pub/sub, RabbitMQ, Cron

Statefulness

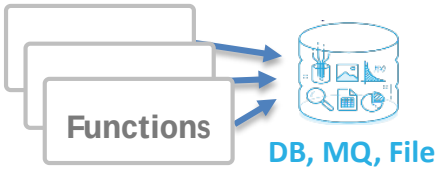

- Data bindings
- Shared volumes
- Context cache

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

iguazio

Delivering Intelligent Decisions in Real-Time

500TB of Raw Data

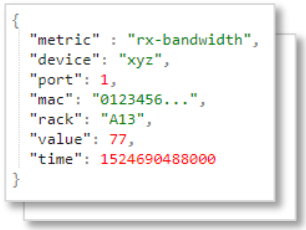

Ingested in real-time
(compressed to **10TB)**

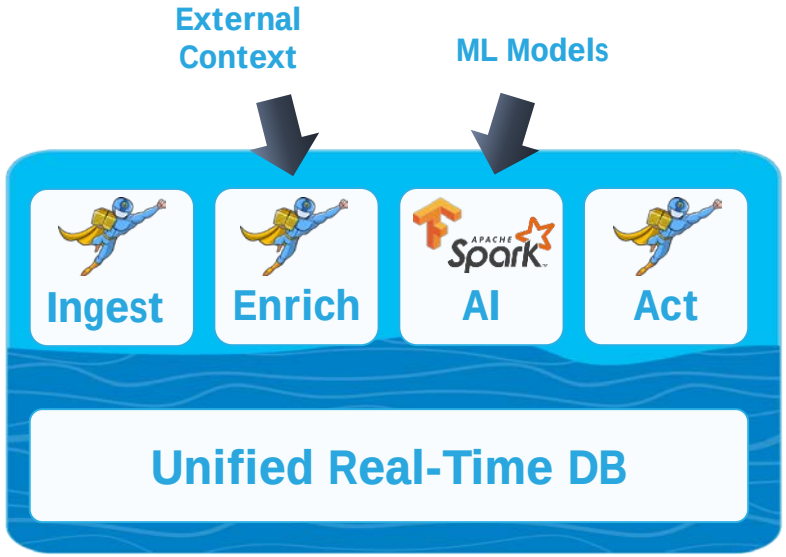

Real-time triggers

Serve

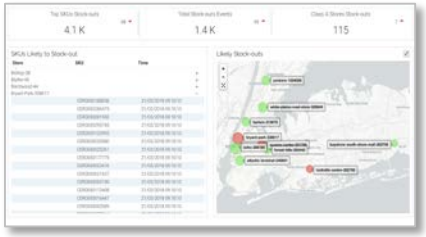

Real-time and historical dashboards

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

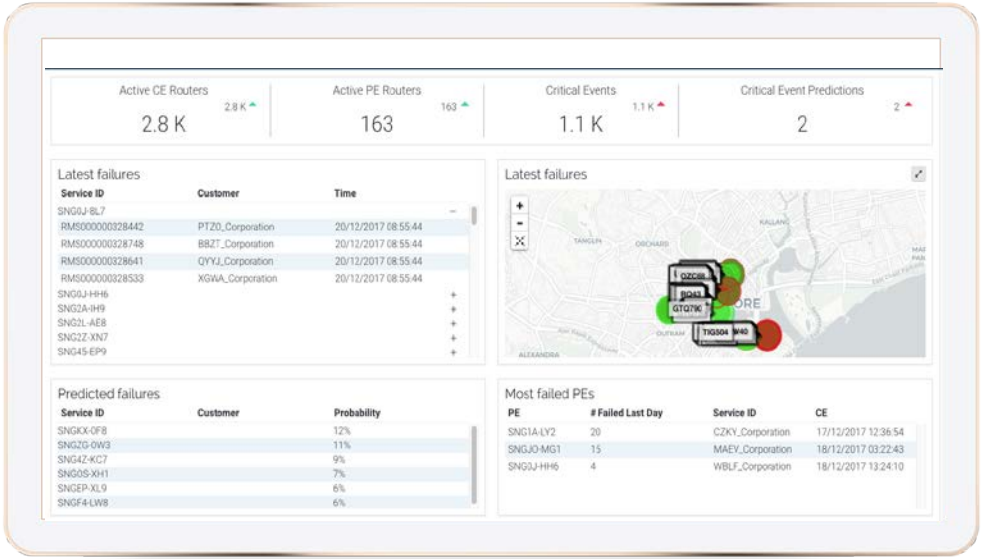

iguazio

Build and Operationalize Proactive Systems Faster

Traditional

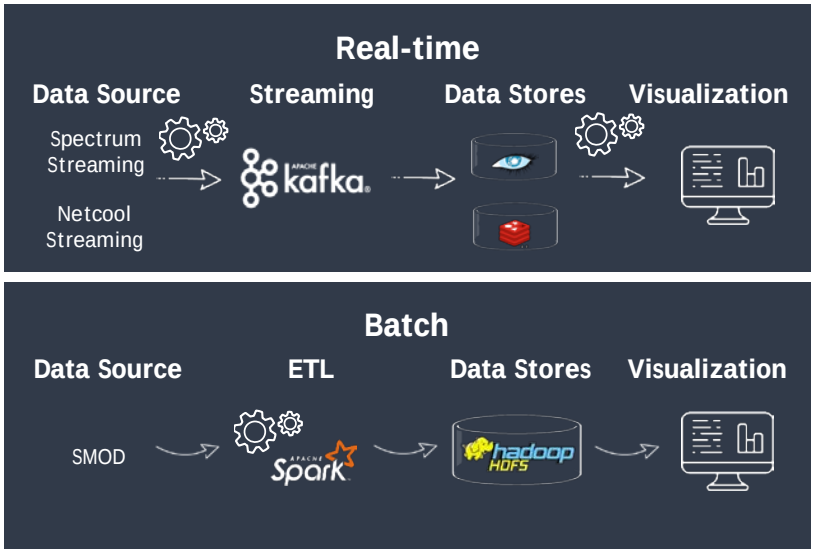

- Complex, skill gaps, slow to productize
- No single view of ops, real-time, history
- Reactive (no actions)

Continuous Analytics

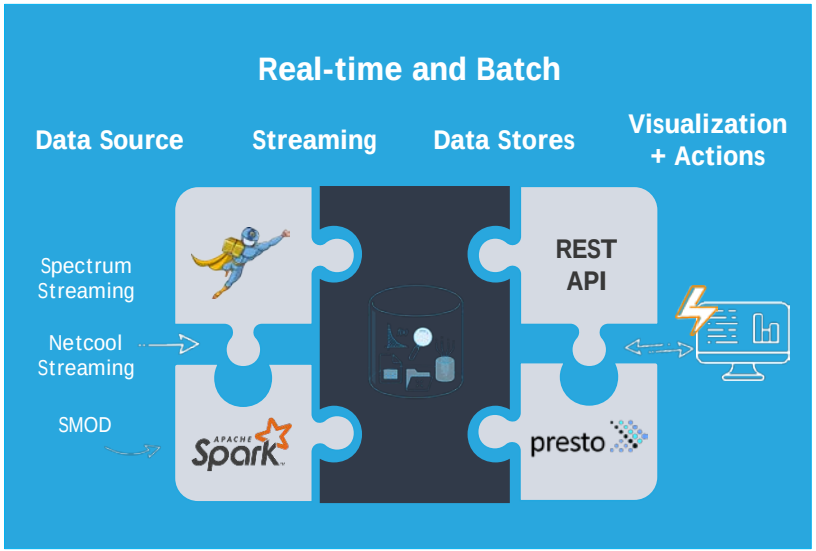

- Simple, just a few weeks to a working app
- Unified view across ALL data
- AI driven, proactive

iguazio

Predictive Maintenance Based on Real-time + Historical + Ops Data

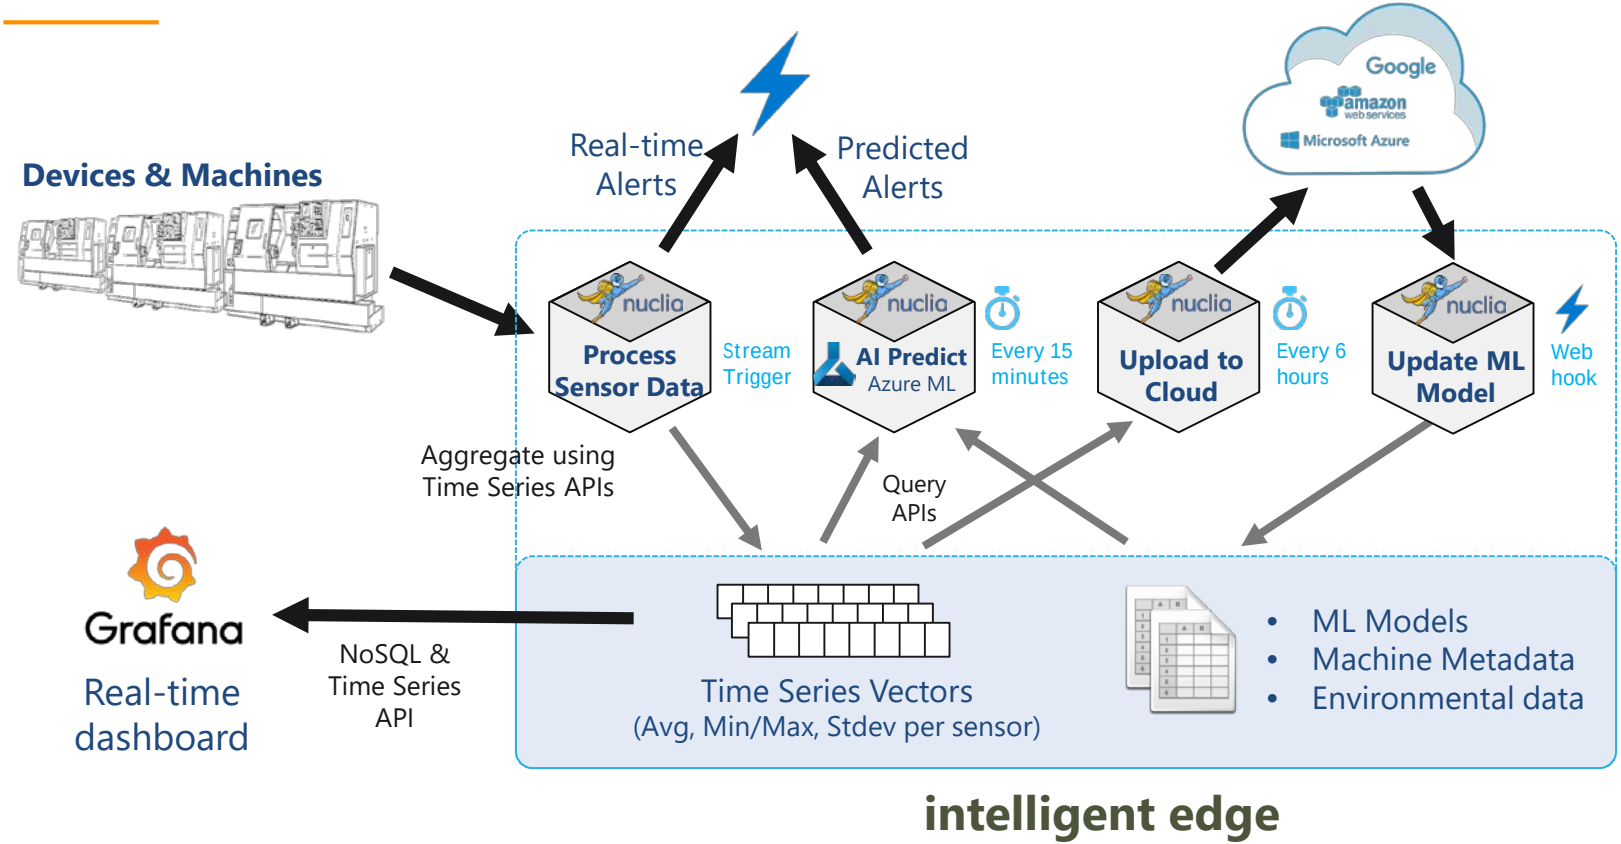

iguazio

Demo: Voice Driven Real-Time Analytics

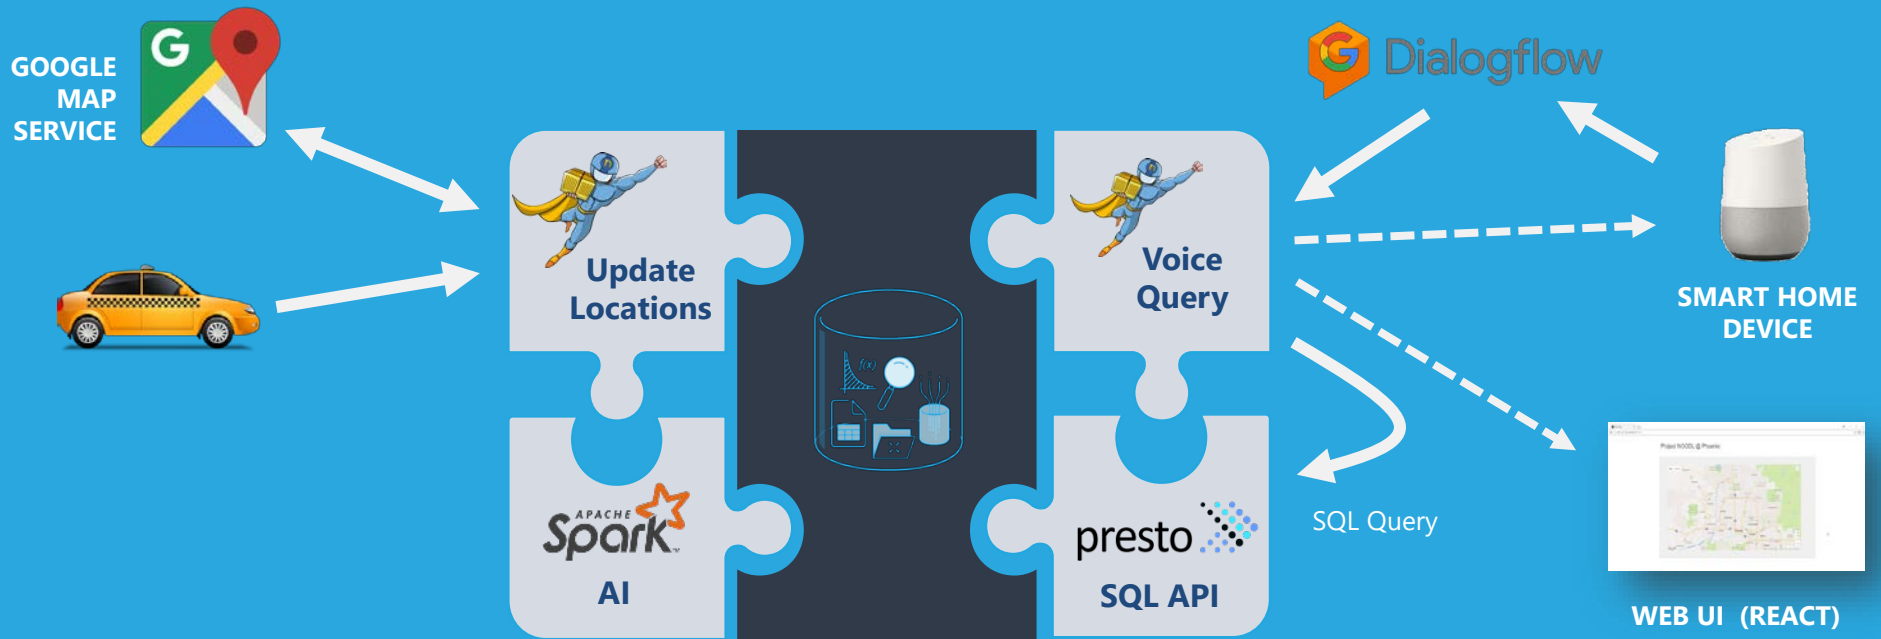

Summary

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