



Managing Edge Computing with Serverless

Lev Radomislensky, Iguazio

Data is Everywhere !
Not limited to Cloud



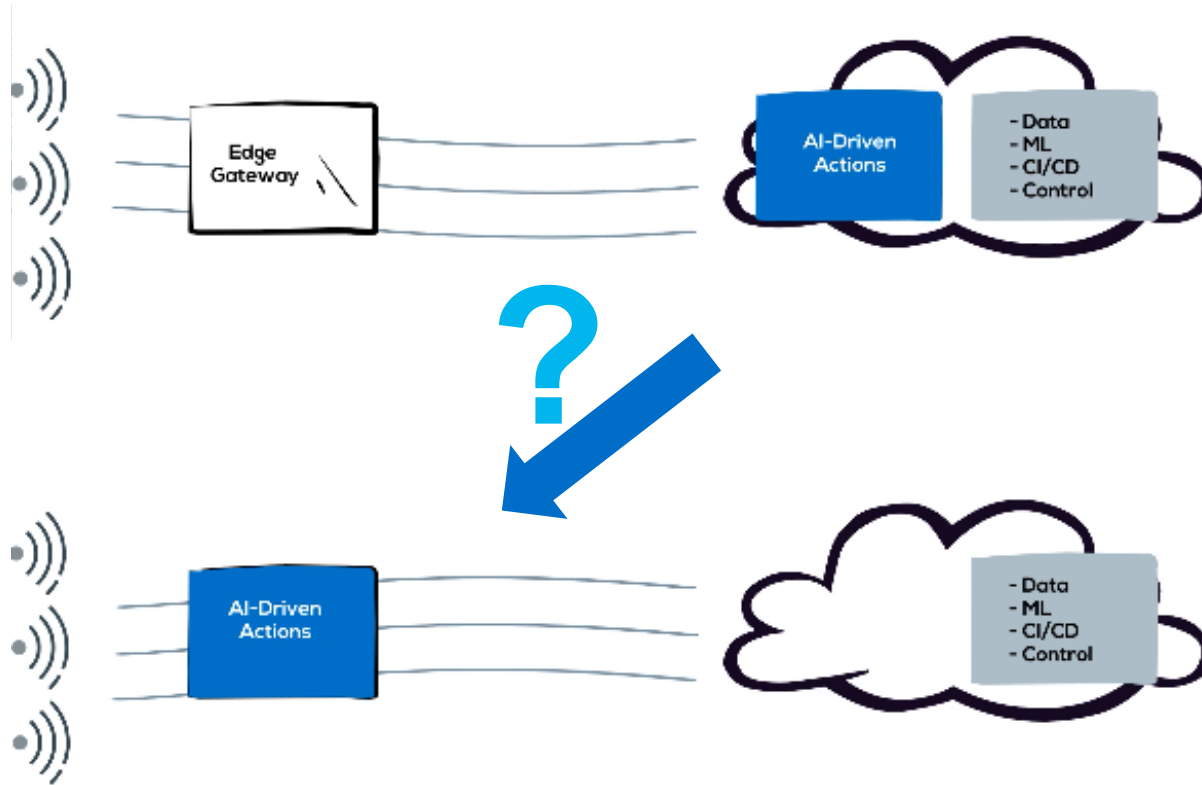
Current Approaches are Problematic



- Run apps locally in a traditional IT approach
- Requires local IT staff
- Hard to manage and update across branches
- Combining data between apps in complex
- Send all the data to the cloud and run all apps there
- Bandwidth limitations
- Intermittent connectivity
- Latency problems


A true hybrid approach is needed with an intelligent edge

What Does it Take to Move Intelligence to the Edge?

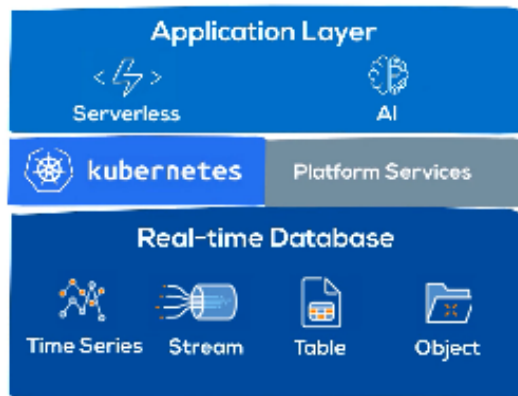


Serverless at the Intelligent Edge

Ingest data from cameras and sensors



Make real-time decisions at the branch



Upload important data to the cloud and manage it all from there



- High-performance pipeline from events to responses, predictions and dashboards
- Simpler deployment and upgrades through integrated apps, AI and data access
- “Set-top box” federated architecture allowing to place apps close to data sources


But, Serverless Comes with Challenges

- Slow performance, lack of concurrency
- Stateless, limited application patterns
- Limited number of cloud specific event sources
- Hard to debug, diagnose and build dependencies
- Cloud vendor API lock-in

Nuclio: Taking Serverless to the Next Level




Extreme Performance




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

Advanced Data & AI Features



- Auto-rebalance, checkpoints
- Any source: Kafka, NATS, Kinesis, event-hub, iguazio, pub/sub, RabbitMQ, Cron, ..
- Jupyter, Spark, Rapids integration

Statefulness



- Data bindings
- Shared volumes
- Context cache

Open-source Serverless for compute and data intensive tasks, **100x faster** than AWS Lambda !

Best Match: Serverless + Kubernetes



- Abstraction: auto generated code, Docker files and YAMLs
- Automated dev & ops flow, observability
- Maximizing performance and resource efficiency
- Everything is a Kubernetes resource

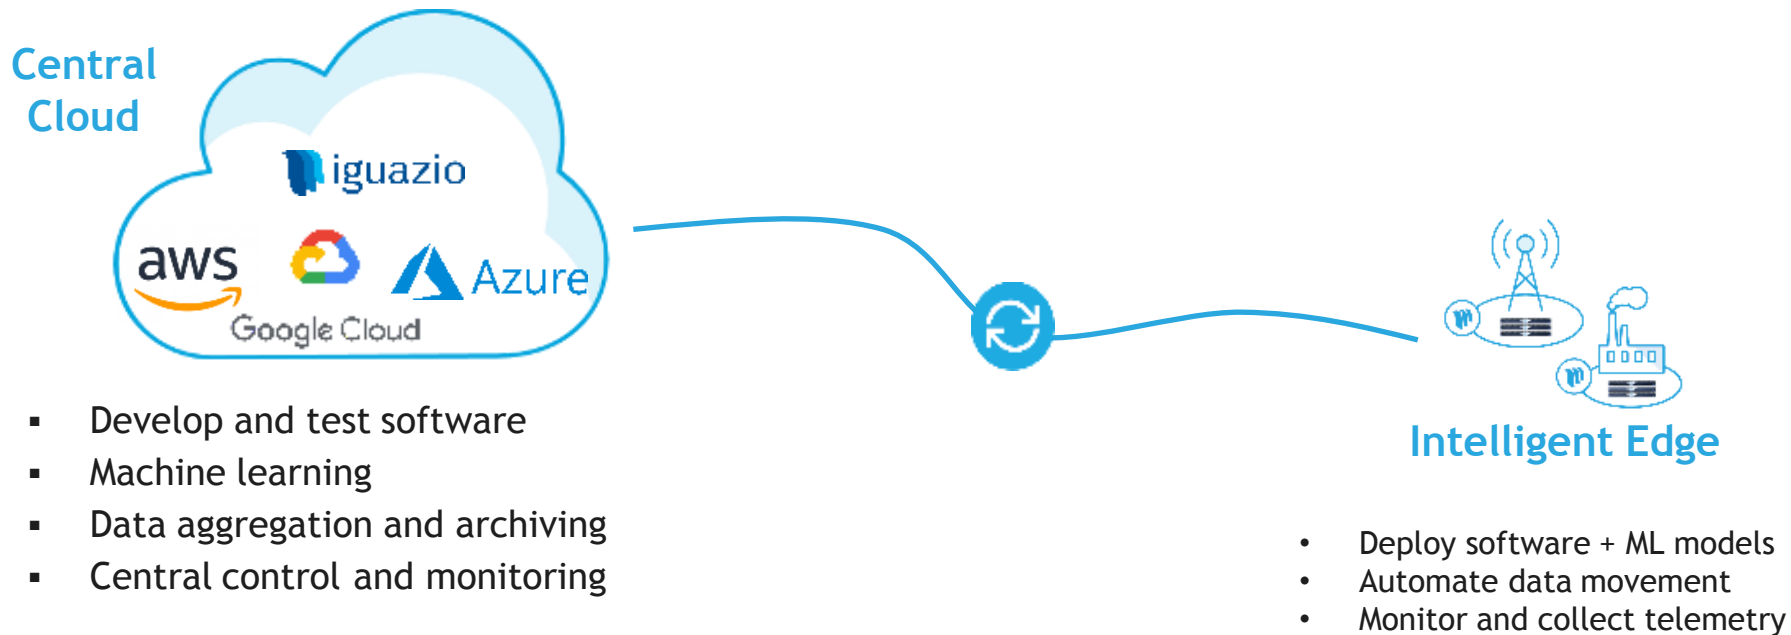


- Cloud independent APIs, on-prem, edge
- Auto-scaling and abstract infrastructure
- One platform for serverless and containers



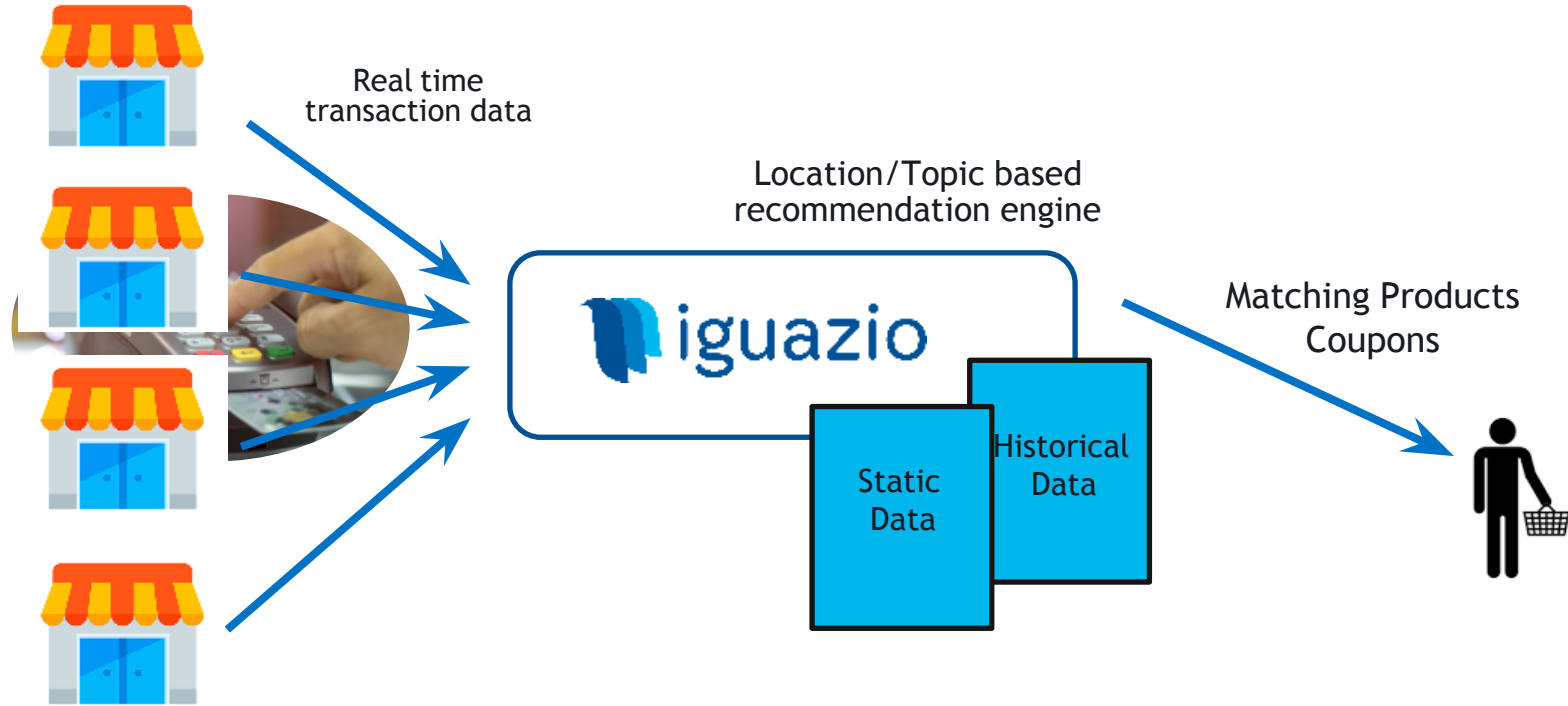
- “Serverless” data services
- Endless scalability

Run Anywhere, Automate Data and Workload Movement



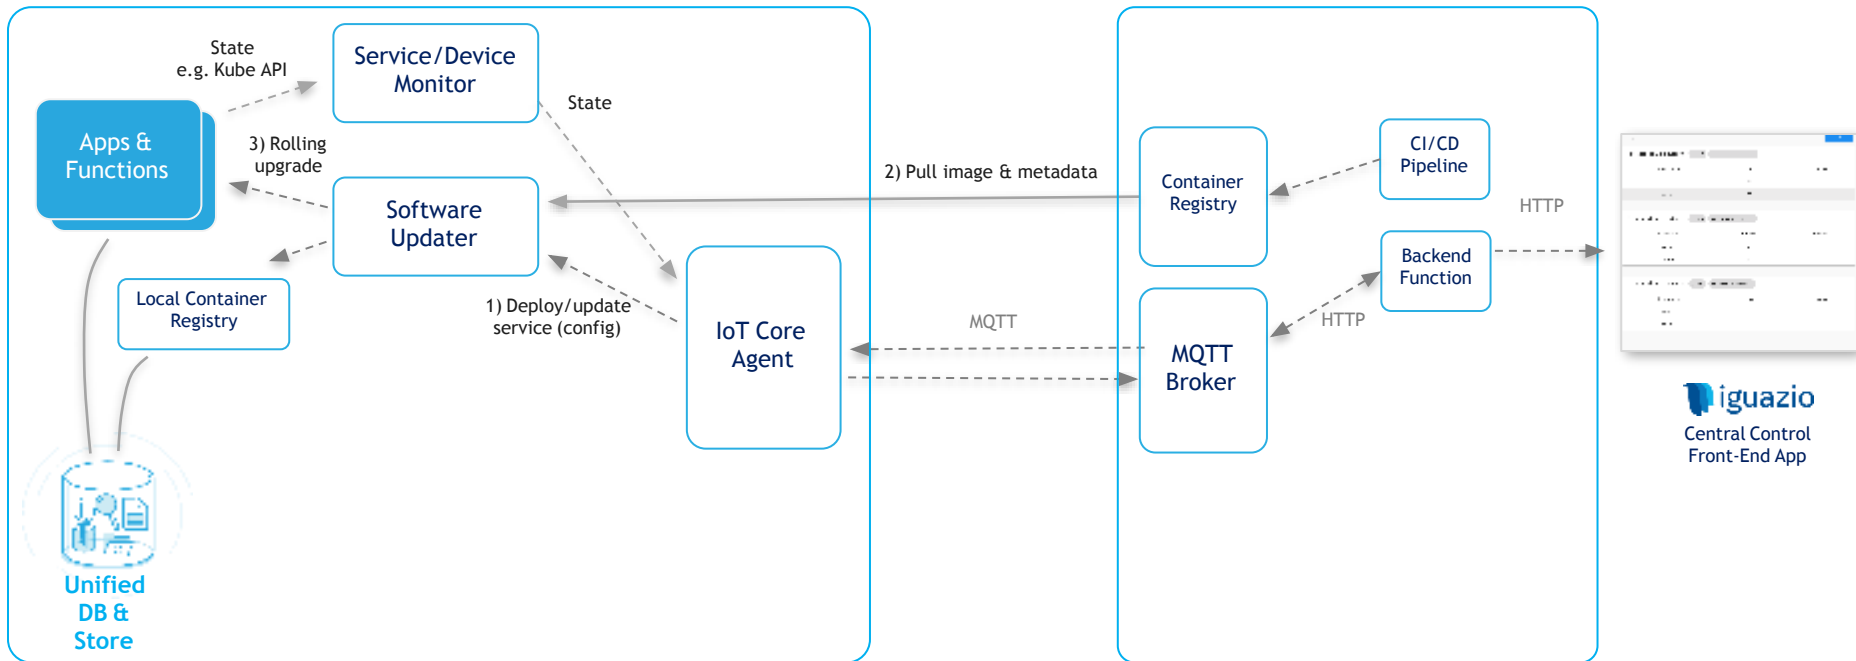
Watch Iguazio-Google Intelligent edge video at:
<https://www.youtube.com/watch?v=ZMnZNh5XB0s>

Real-time location based recommendation engine



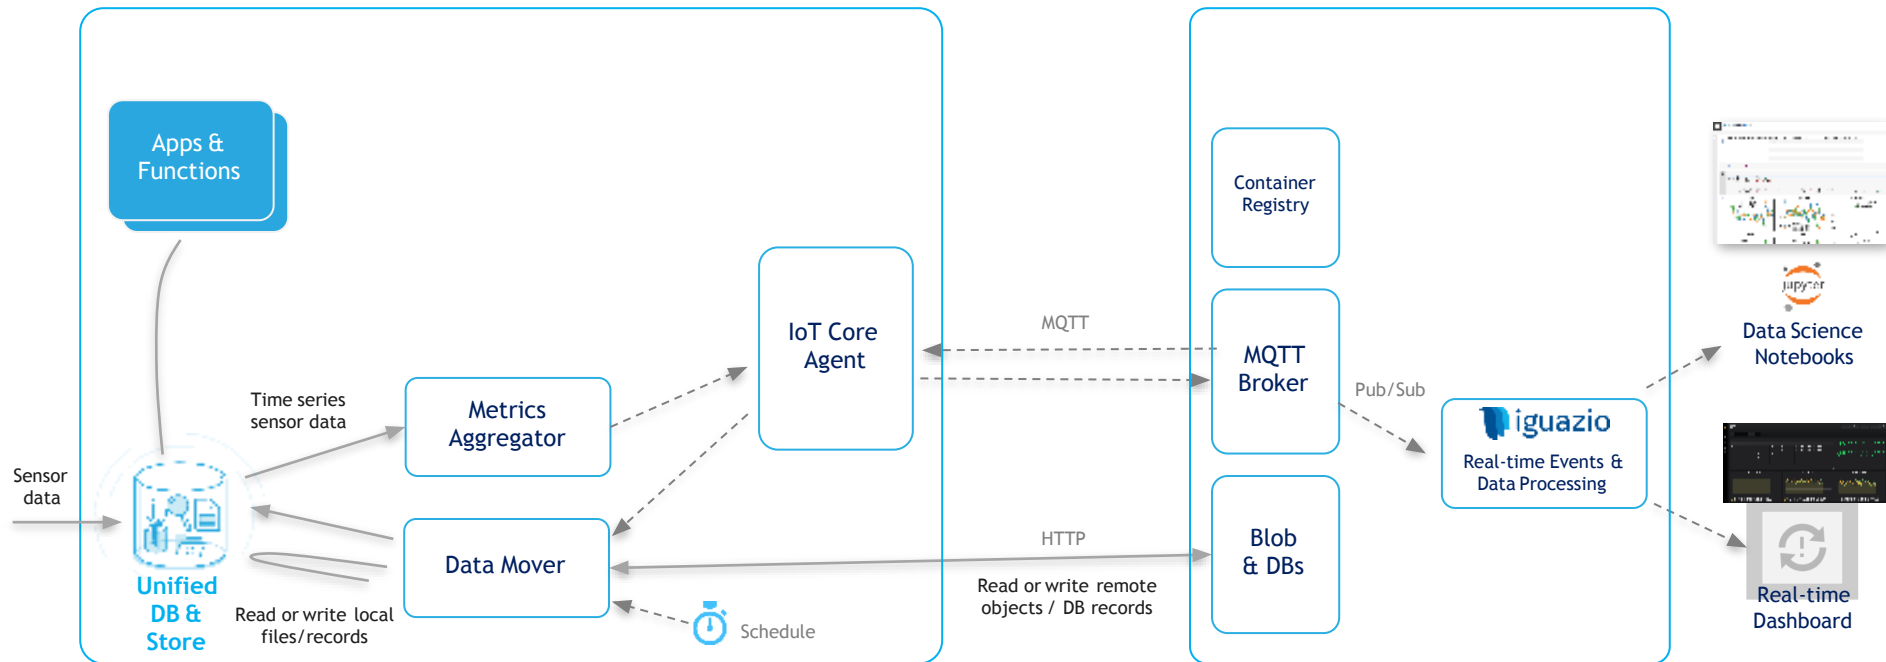
Demo: Control Flow Details

iguazio Intelligent Edge - Control Functions

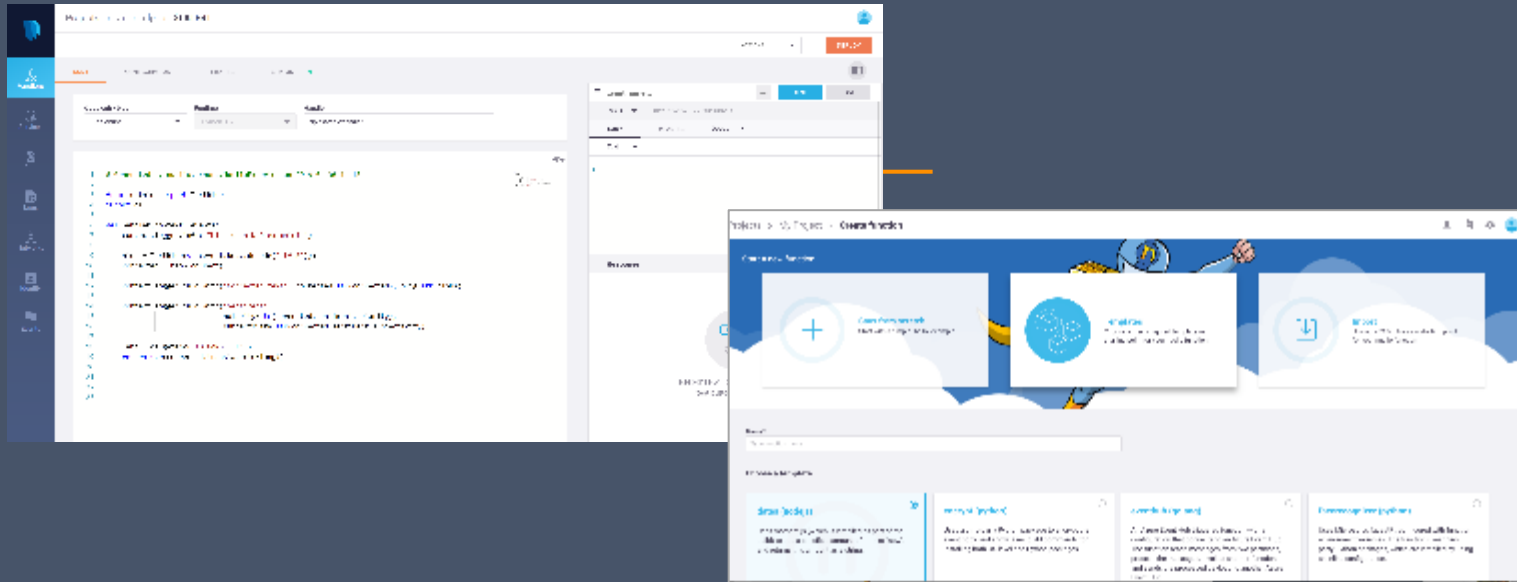


Demo: Data Movement & Analysis Flow Details

 Intelligent Edge - Control functions



Demo: Build & Test an App in 2 min



Simple Configuration UI (Over MQTT Broker)

Services

CONFIGURE

iot-core-demo-device-0 go to location iot-core-demo-device-0

Service name	Version	Replicas
opensearch	0.1	1
elasticsearch	0.0.1	1

iot-core-demo-device-1 go to location iot-core-demo-device-1

Service name	Version	Replicas
opensearch	0.0.1	1
elasticsearch	0.1	1

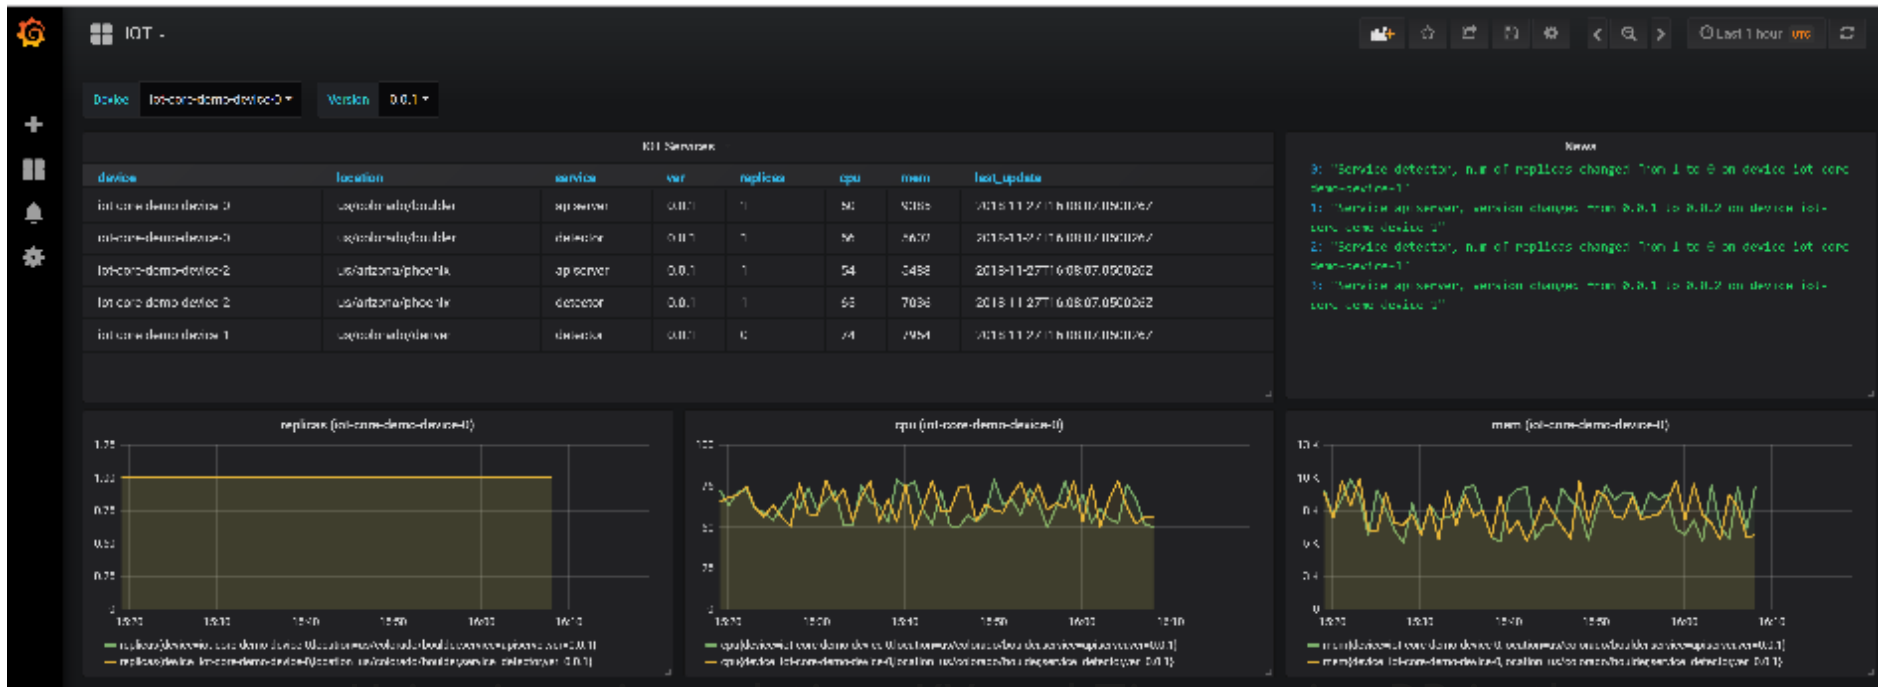
iot-core-demo-device-2 go to location iot-core-demo-device-2

Service name	Version	Replicas
opensearch	0.1	1
elasticsearch	0.0.1	1

Push software, config or commands to a set of devices

View data and metadata per device and service

Real-time Aggregation Dashboard (Grafana)

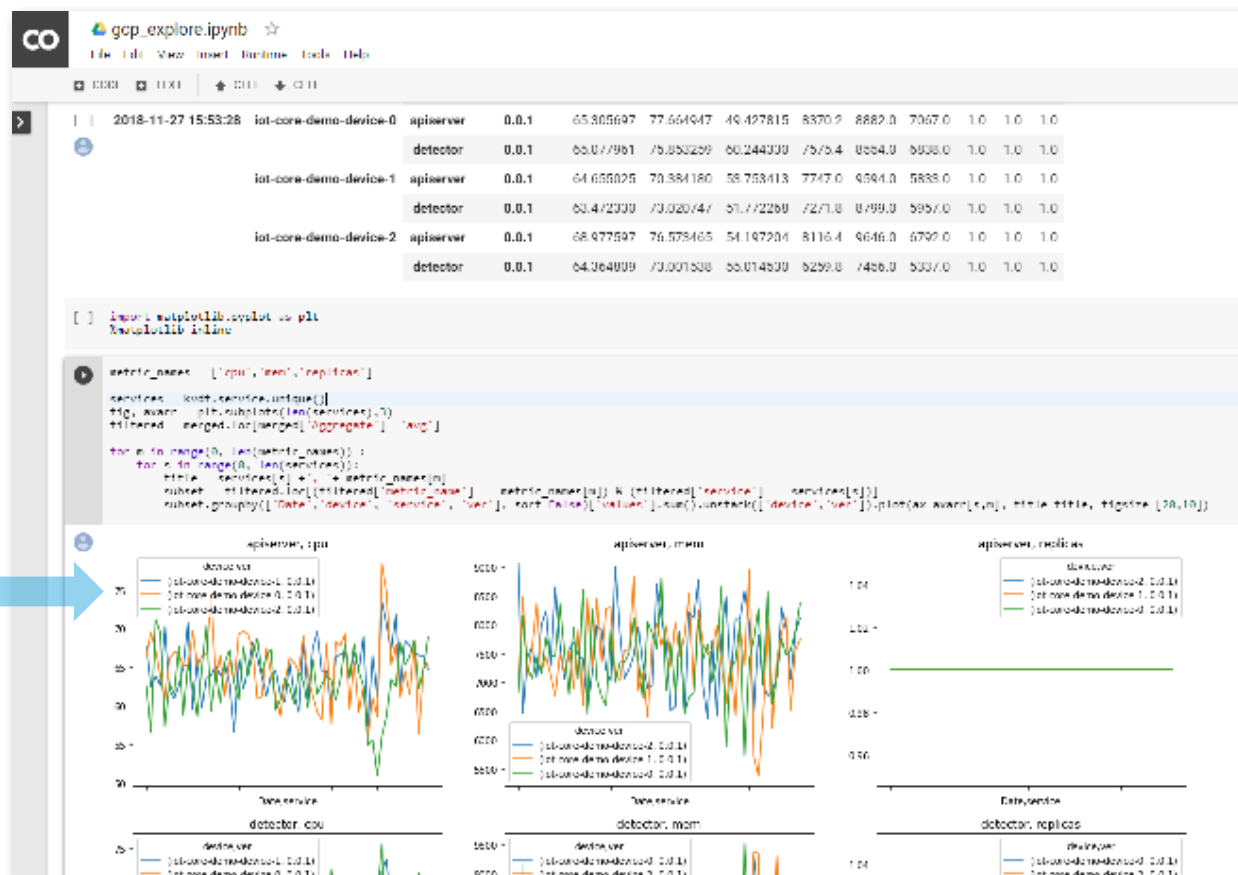


Using iguazio real-time KV and Time-series DB in the cloud

Interactive Data-Science Workbench on Real-time Data



Using Google
Collaboratory Service
(Jupyter)



Data served by iguazio real-time analytics & TSDB solution

Summary

Serverless is the Optimal Solution for the Edge

- Enrich data, take actions and persist the information within the database
- Accelerate application performance (400k events/sec/proc, 0.1ms latency)
- Cut dev and ops overhead, shorten time to production and reduce costs