#ExploreSAS

SAS EXPLORE

Level Up Your Skills in AI and Analytics

Sept. 11–14 // Las Vegas

sas

Learn SAS Event Stream Processing open-source integration with GIT and Grafana on Azure Marketplace

Daniele Cazzari and Steven Allan

SAS EXPLORE

sas

Streaming Data

SAS EXPLORE

Streaming Analytics

Capturing Value Almost Immediately, as the Events Occur

Streaming Analytics is the **application of analytics to data while it's in motion, and before it's stored** – and ranges from data manipulation, normalization all the way to machine learning.

Provides insights into Connected things (IoT devices), sensors and devices, and networked machines

Analyzes data at the edge, as close to the event location as possible, before its value is diminished due to information lag and before the volume of data overwhelms traditional analytics

Identifies and examines patterns as events occur, **so immediate actions can be taken on those events of interest as they happen**






SAS EXPLORE

sas

Streaming vs Batch Processing

Streaming vs Batch

	Streaming Analytics	Batch Processing
Architecture	Streaming analytics platforms are typically based on event-driven architectures.	Batch processing platforms are typically based on collected sets of data, using Hadoop or Spark.

Streaming vs Batch

	Streaming Analytics	Batch Processing
Architecture	Streaming analytics platforms are typically based on event-driven architectures.	Batch processing platforms are typically based on collected sets of data, using Hadoop or Spark.
Time of processing	Data is processed in real time as it is generated or ingested into the system.	Data is processed in batches, which can be scheduled or triggered at regular intervals.

Streaming vs Batch

	Streaming Analytics	Batch Processing
Architecture	Streaming analytics platforms are typically based on event-driven architectures.	Batch processing platforms are typically based on collected sets of data, using Hadoop or Spark.
Time of processing	Data is processed in real time as it is generated or ingested into the system.	Data is processed in batches, which can be scheduled or triggered at regular intervals.
Latency	Latency is very low, as data is processed as soon as it is available.	Latency can be high, as data may not be processed until the end of the batch.

Streaming vs Batch

	Streaming Analytics	Batch Processing
Architecture	Streaming analytics platforms are typically based on event-driven architectures.	Batch processing platforms are typically based on collected sets of data, using Hadoop or Spark.
Time of processing	Data is processed in real time as it is generated or ingested into the system.	Data is processed in batches, which can be scheduled or triggered at regular intervals.
Latency	Latency is very low, as data is processed as soon as it is available.	Latency can be high, as data may not be processed until the end of the batch.
Data volume	Streaming analytics is typically used for processing smaller data packages but at very high speed	Batch processing is typically used for processing larger data volumes.

Streaming vs Batch

	Streaming Analytics	Batch Processing
Architecture	Streaming analytics platforms are typically based on event-driven architectures.	Batch processing platforms are typically based on collected sets of data, using Hadoop or Spark.
Time of processing	Data is processed in real time as it is generated or ingested into the system.	Data is processed in batches, which can be scheduled or triggered at regular intervals.
Latency	Latency is very low, as data is processed as soon as it is available.	Latency can be high, as data may not be processed until the end of the batch.
Data volume	Streaming analytics is typically used for processing smaller data packages but at very high speed	Batch processing is typically used for processing larger data volumes.
Use cases	Streaming analytics is used for applications that require real-time insights, such as fraud detection, anomaly detection, and real-time marketing.	Batch processing is used for applications that do not require real-time insights, such as data mining, reporting, and retrospective analytics.

SAS Event Stream Processing

Overview and Capabilities

SAS EXPLORE

sas

Real time Industry Solutions

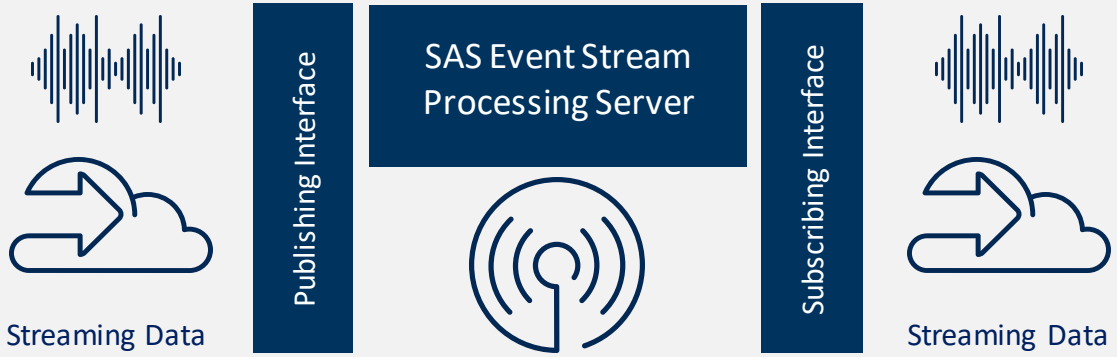

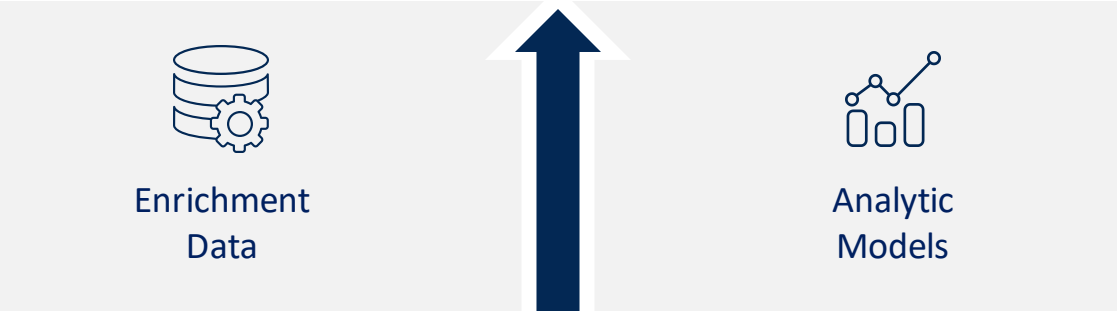

ESP becomes the enabling platform for solutions.

Processes data continuously, on the move, in-memory with very high speed and low latency

Flexible Publish and Subscribe framework

Performs actions such as filtering, aggregation, pattern detection, calculations, correlations, machine learning, geofencing, image analytics and much more…

Add historical context or enrichment data to what is being observed in real time

Score events using externally trained analytical models

Design and test projects in the **low code**, **graphical** design environment

Orchestrate, deploy and monitor ESP projects.

Track and update when new champion models are promoted, provides a fully featured analytical model lifecycle.

SAS Event Stream Processing (ESP)

Engineered for Agility

- Small footprint **Docker Containers** engineered for Edge
- Support for **x86_64** and **ARM** processor
- Support **CUDA**, **TensorRT** and **Openvino** GPU hardware acceleration
- Supports lightweight embedded technology to cloud distributed architecture

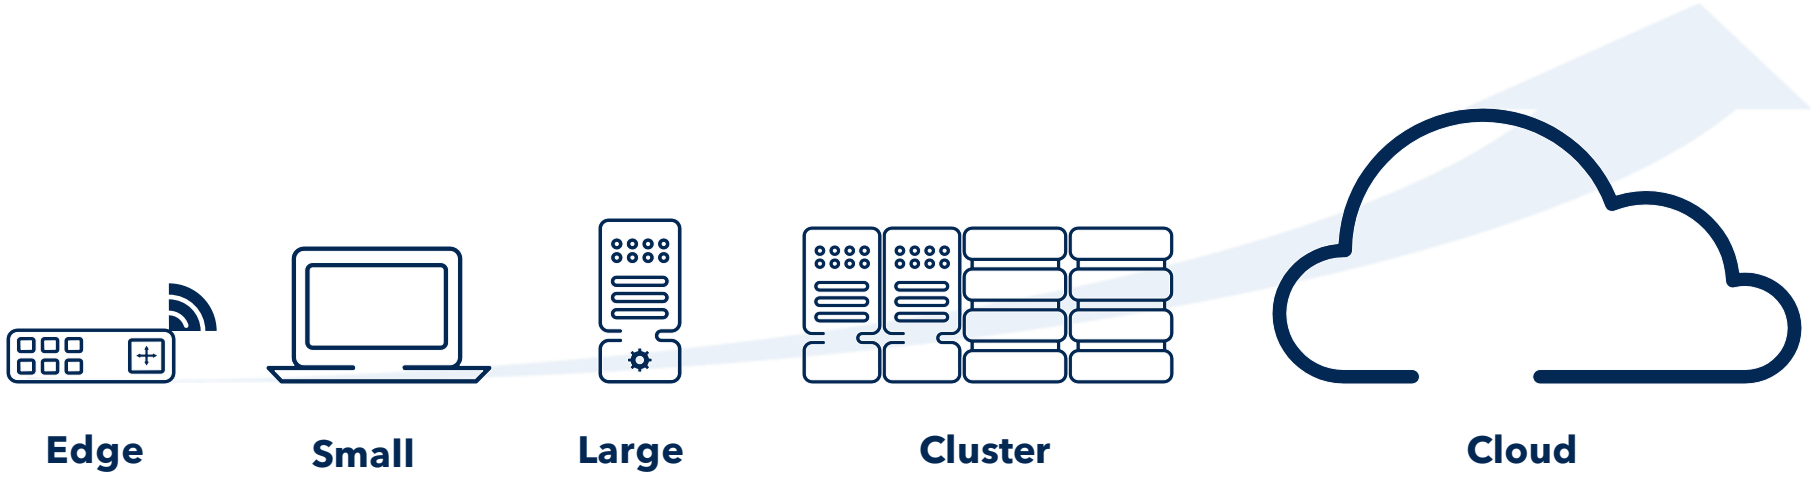

Streaming Analytics Depth and Breadth

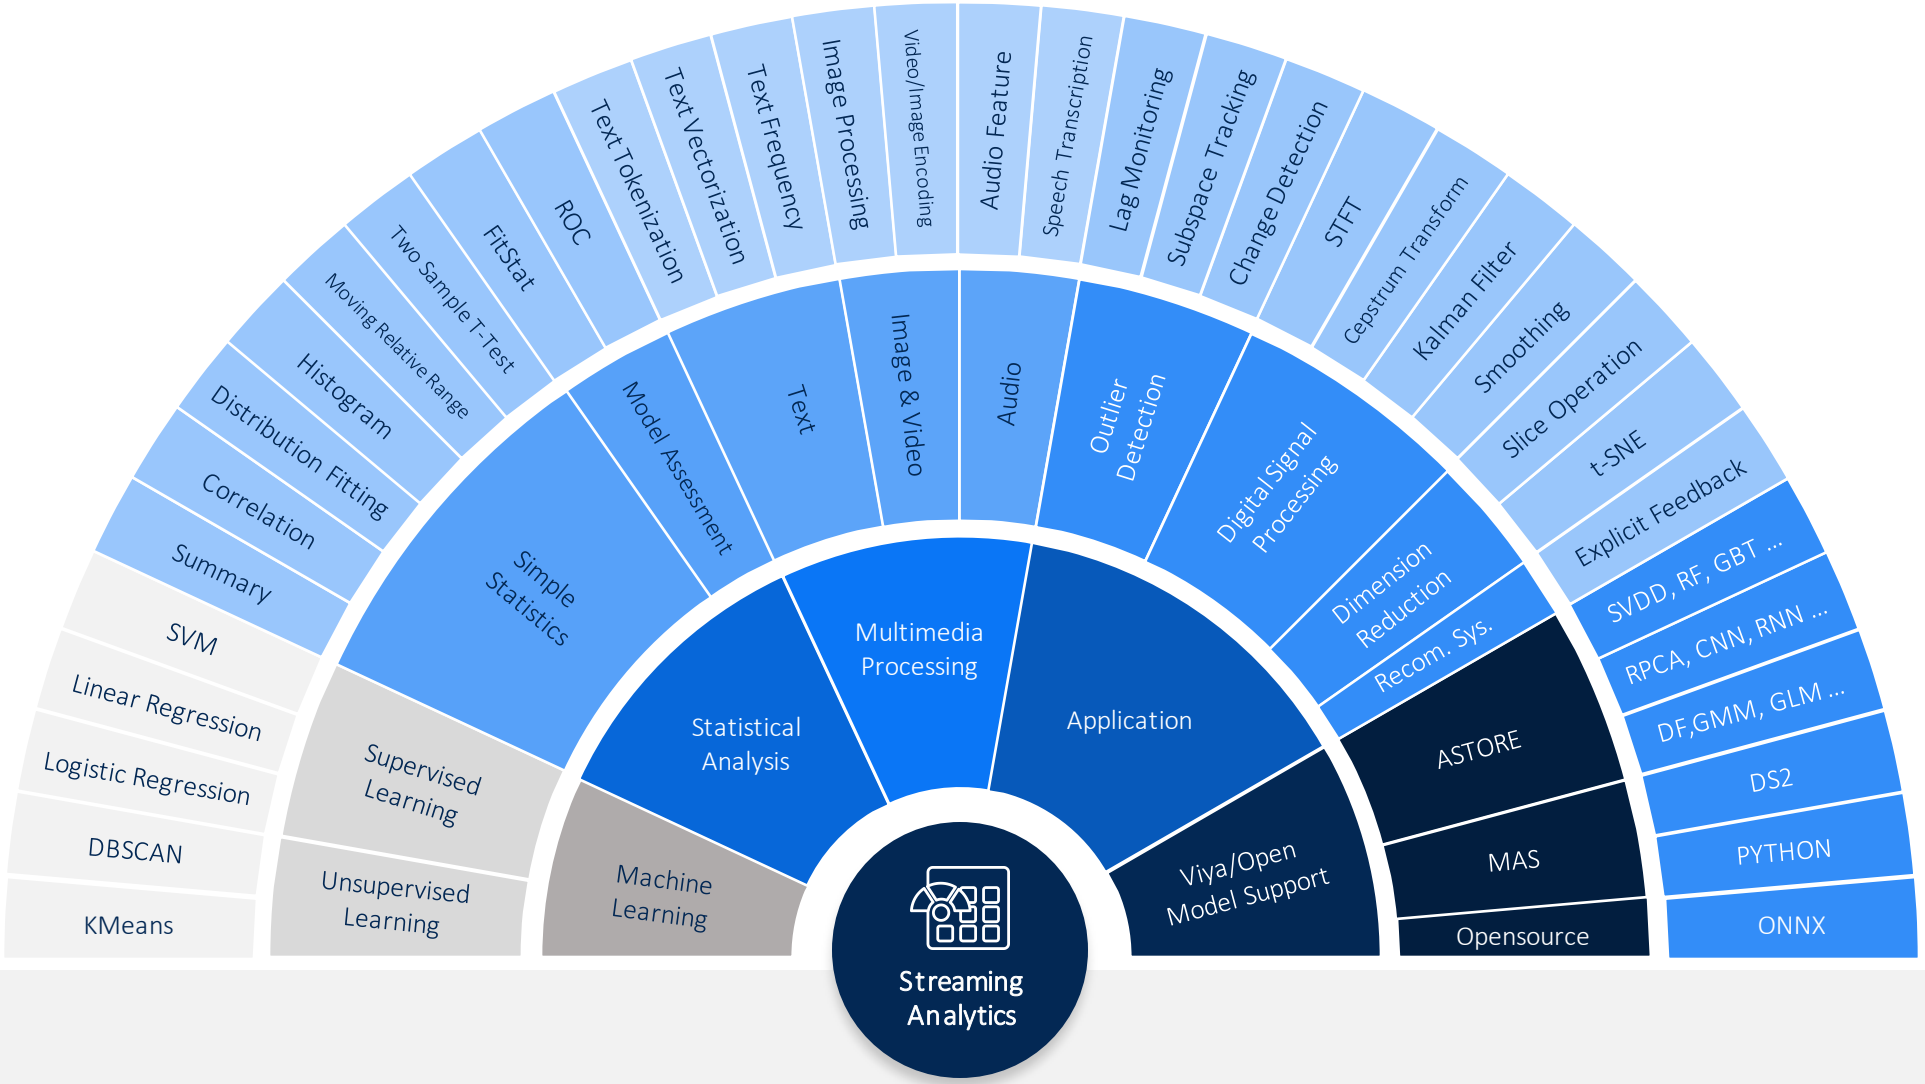

SAS® Internet of Things

Ecosystem Integration - 300+ Endpoints

SYSTEMS & APPLICATIONS

BACnet International
intel
OPC
GREENPLUM DATABASE
Microsoft Azure Event Hubs
amazon KINESIS
PI SYSTEM OSIsoft.
RTSP Real-Time Streaming Protocol
cassandra
protobuf Protocol Buffers
Solace Systems®
QuasarDB
NVIDIA
ORACLE
Modbus
MQTT
kafka
JMS
Asterisk
RabbitMQ
IBM WebSphere MQ
docker
Azure IoT Edge
kubernetes

STANDARDS

FILE/SOCKET
XML / JSON
ODBC
JMS

MQTT
OPC-UA
HTTP RESTFUL
WEB SERVICES

WEBSOCKETS
SMTP
UVC
RTSP

PUBLISH & SUBSCRIBE API

CONNECT TO ANY SYSTEM WITH JAVA, C++, PYTHON
FULLY DOCUMENTED AND EASY TO USE

SAS EXPLORE

sas

ESP and Opensource Integration

Opensource Integration

Kubernetes Operator

Container Orchestration

Grafana Plugin

Visualization

Git Support

Versioning

ONNX Support

Model Portability

ESP and Kubernetes

ESP Kubernetes Operator:

Desired State based process to create, deploy, monitor, and scale ESP Projects

Why provide this capability?

- Automated and configurable Docker container orchestration
- Standard, open-source Kubernetes APIs
- Provide demand based autoscaling
- Interoperates with popular K8s tools

Benefits: **Resiliency, visibility, failover and elasticity**

ESP and Git Versioning

Promotion and Continuous Integration

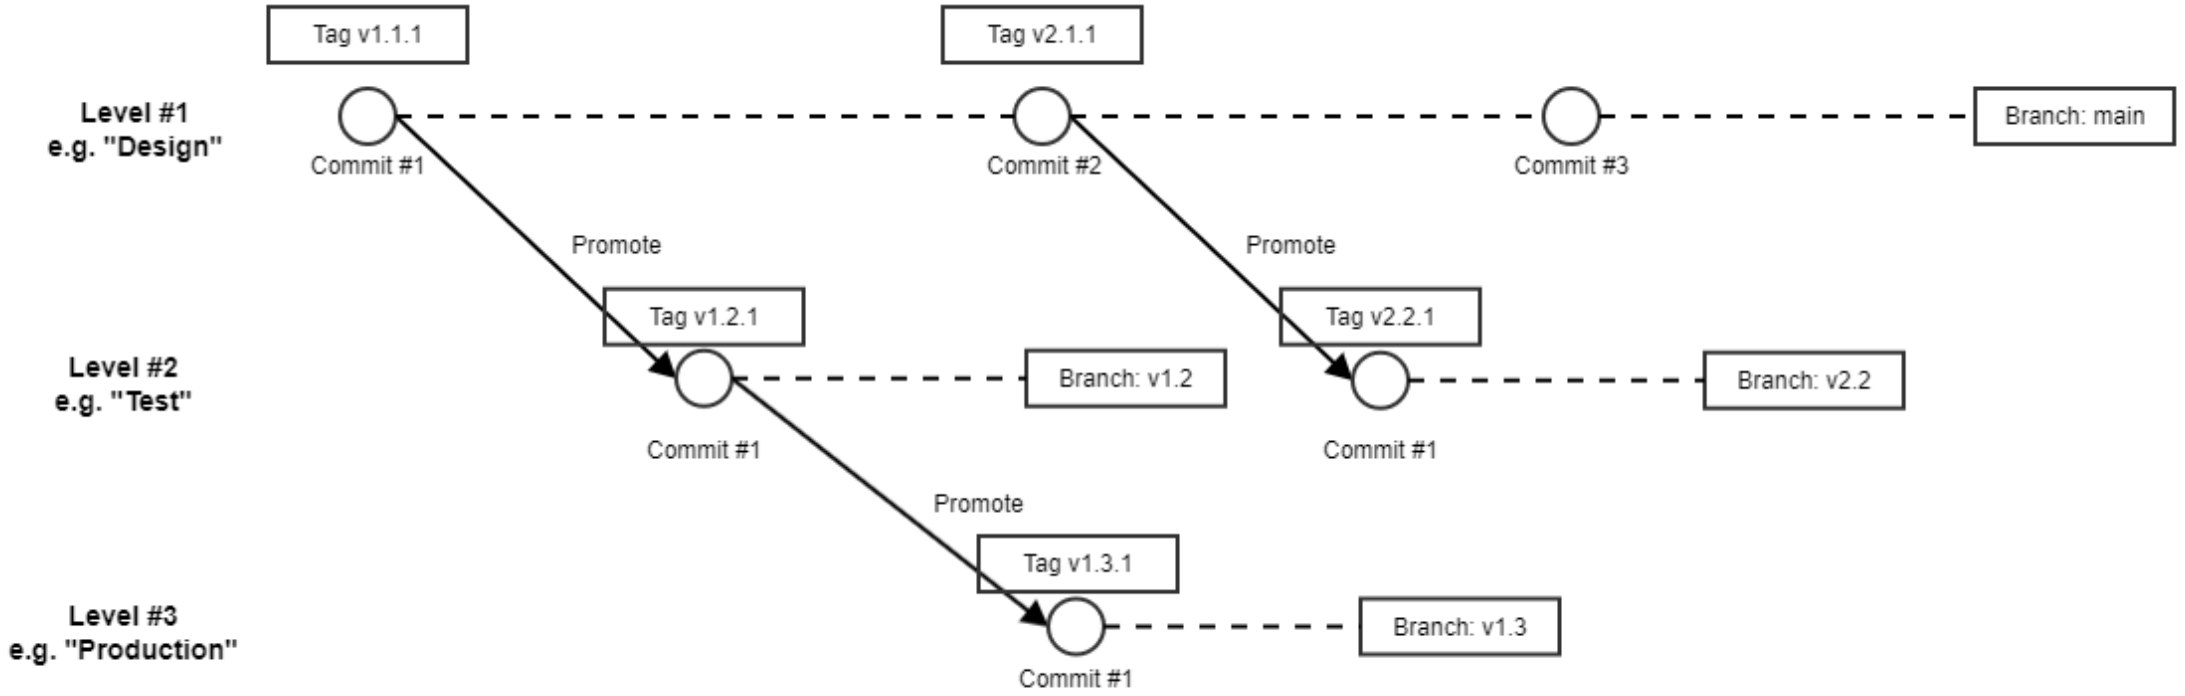

ESP and Git Versioning

Promotion and Continuous Integration

Ensures a robust, streaming analytics lifecycle:

- Design and test ESP projects in ESP Studio
- Promote to Test for deployment in Event Stream Manager (ESM)
- Execute long running tests in ESM to validate logic
- Promote to Production and manage the project when “Live”

ESP and Grafana

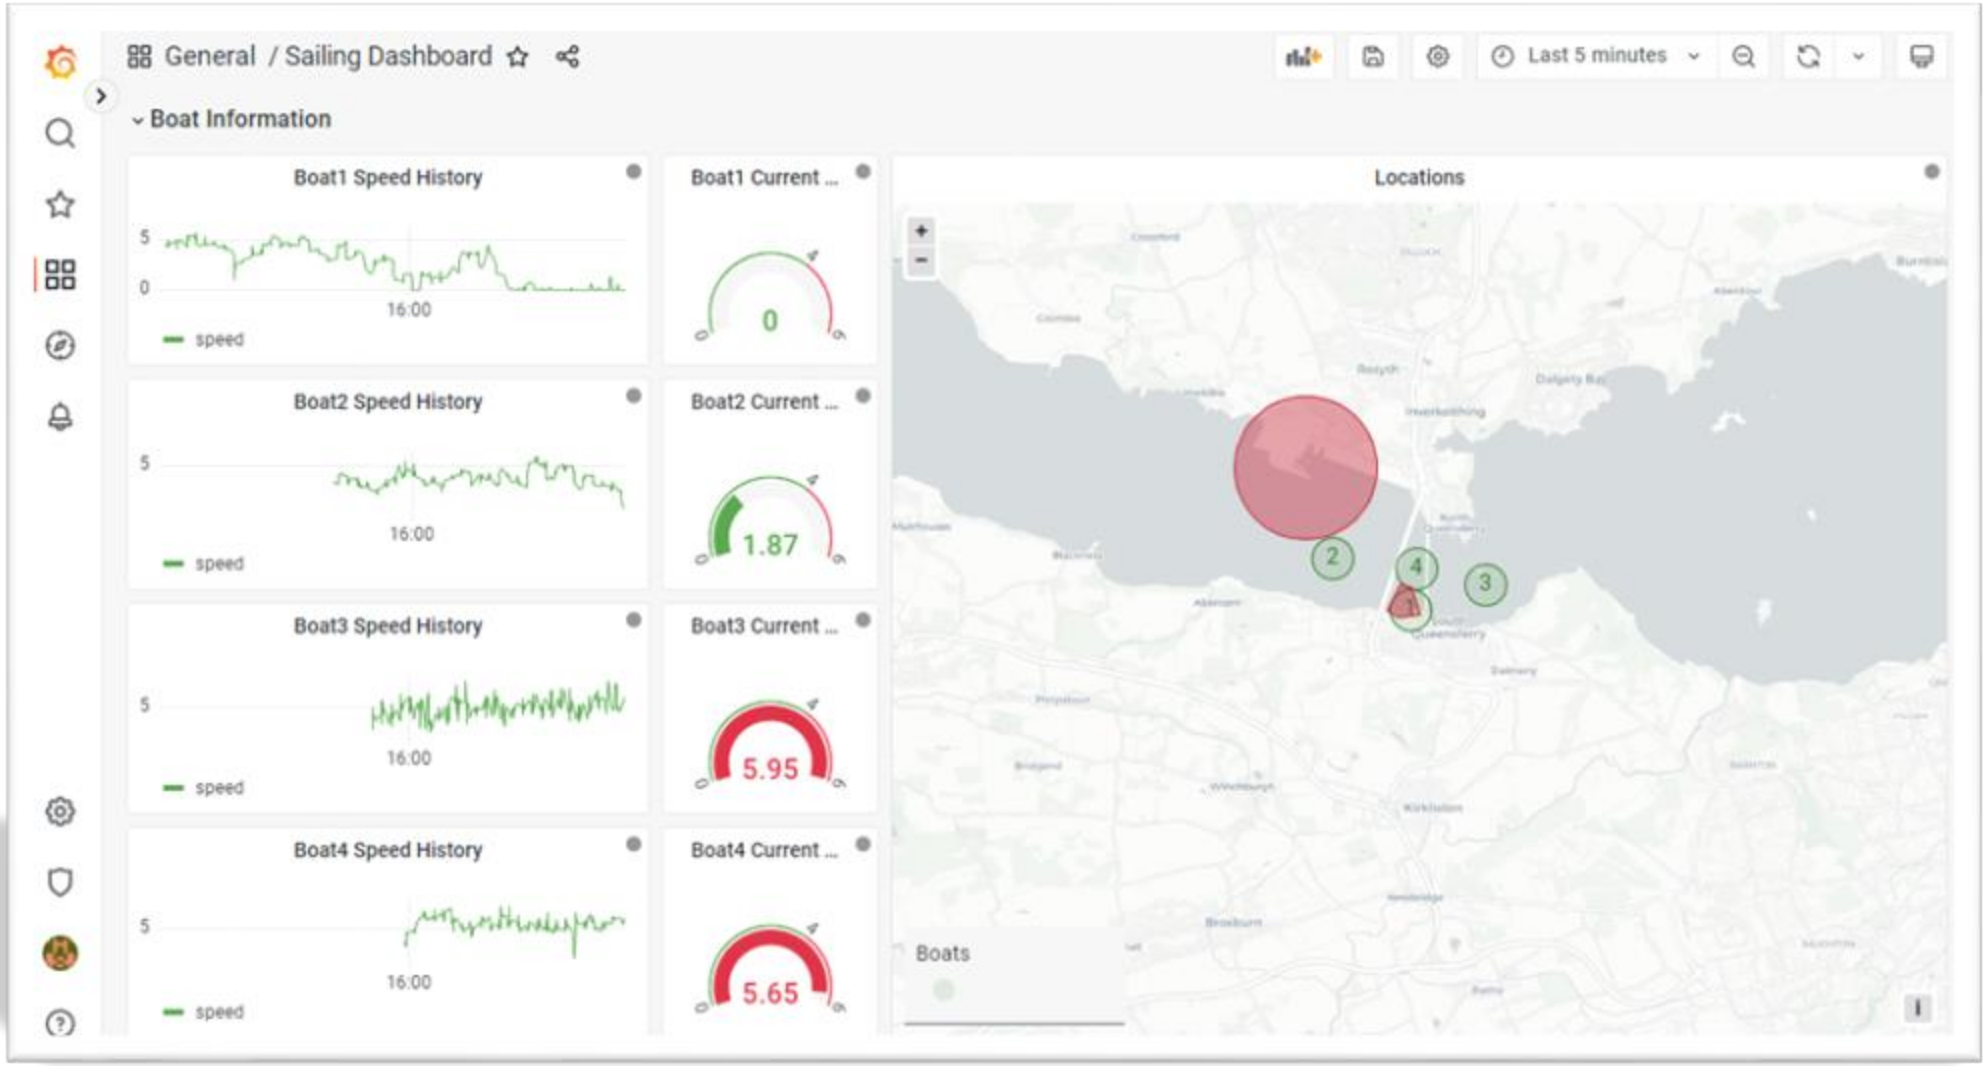

ESP and Grafana

Visualization and Reporting

Visualize your running ESP projects in Grafana.

The Grafana ESP plugin…

- Authenticates to the ESP client midtier
- Discovers ESP server pods in the cluster
- Makes direct web socket connections to the running projects
- Manages connections to be as efficient as possible!

ESP and ONNX

Open
Neural
Network
e**X**change

ESP and ONNX

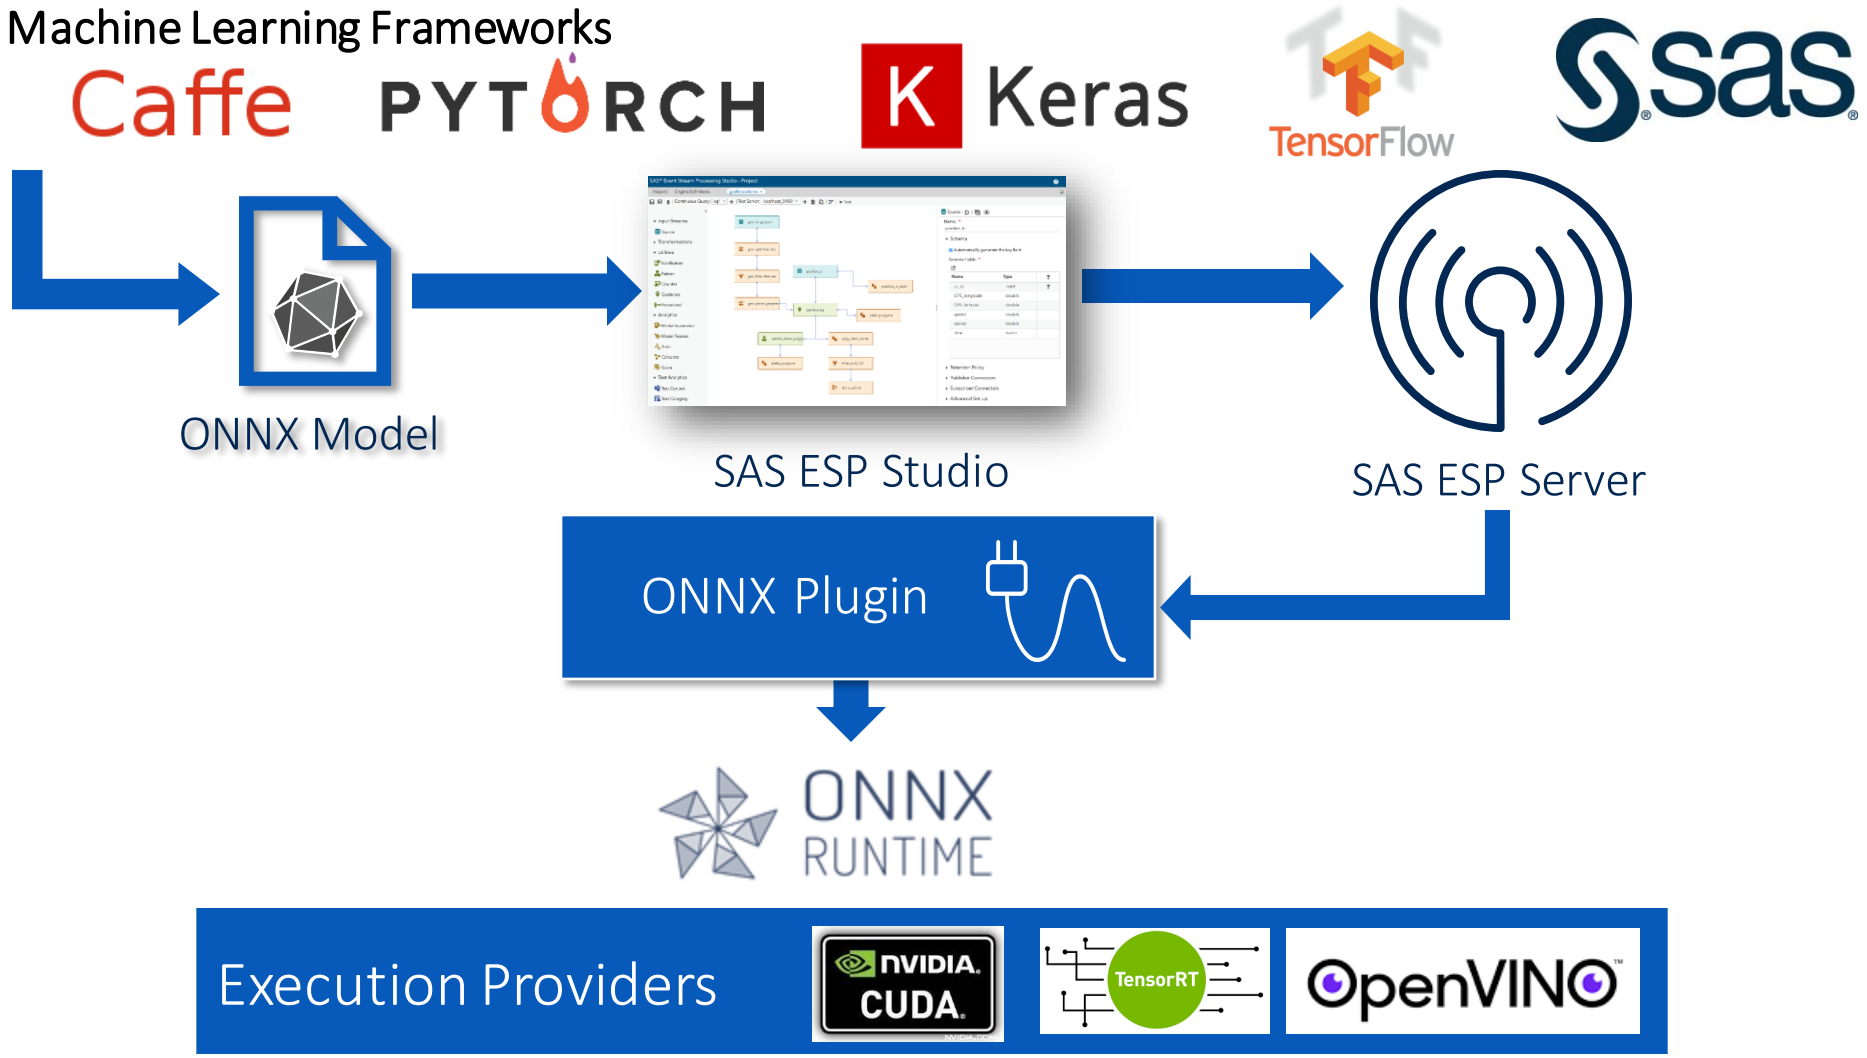

SAS Event Stream Processing

Available on the Azure Marketplace

SAS EXPLORE

sas

Event Stream Processing

Azure Marketplace Deployment

- Simple and intuitive provisioning of a cost effective Kubernetes infrastructure that fully support SAS Event Stream Processing enabled with Git and Grafana
- Support for scalable CPU and GPU workload tailored for you application needs
- Automatic configuration of managed Postgres DB that could be also used to store streaming analysis outcome
- Optional integration with EventHub for data ingestion and external blob storage

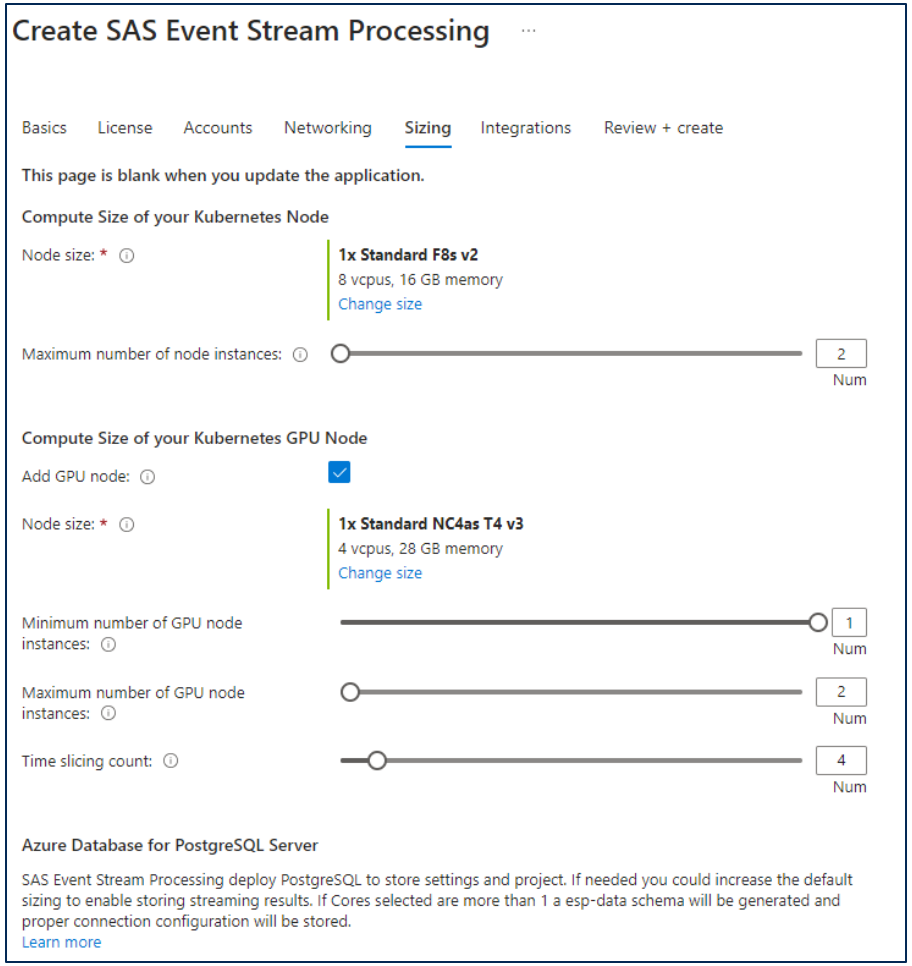

SAS EXPLORE

sas

Event Stream Processing

Azure Marketplace Architecture

Sensor Data
(via IoT Hub or
custom streaming)

Stream

Streaming Input

Preconfigured Event Hub Input Topic

Custom Streams

SAS Event Stream Processing Server

Streaming Output (actionable insight)

Preconfigured Event Hub Output Topic

Preconfigured PostgreSQL

Enterprise Systems

Tools

SAS Event Stream Manager

SAS ESP Studio

Grafana

Infrastructure

Azure Kubernetes Services
Automating deployment, scaling, and management of containerized applications

Keycloak
Identity and Access Management

GIT
Distributed version control system

PostgreSQL
Tools Configuration and Identity storage

Azure Storage
ESP Project persistence and data sharing

SAS Event Stream Processing

Key Takeaways

SAS offers the most built-in analytics
Irrespective of where you want to execute your analytics…

1

SAS is striving to demystify streaming analytics
User friendly UIs, low code development and deployment tools

2

SAS is an enabler
Strategic partnerships, opensource integration and complimentary technologies

3

Learn more at: www.sas.com/esp

SAS EXPLORE

sas

Read more about ESP

SAS Packs the Most Analytics Punch

"SAS Event Stream Processing (ESP) stands out as the platform with the most built-in analytics for machine learning and other advanced analytics, as well as a mature edge analytics capability for IoT applications."

The Forrester Wave™
Streaming Analytics, Q2 2021

Event Stream Processing

is available on Azure Marketplace

- You could deploy it with your existing license now!
- For a trial license or for any question please write to iotcontact@sas.com

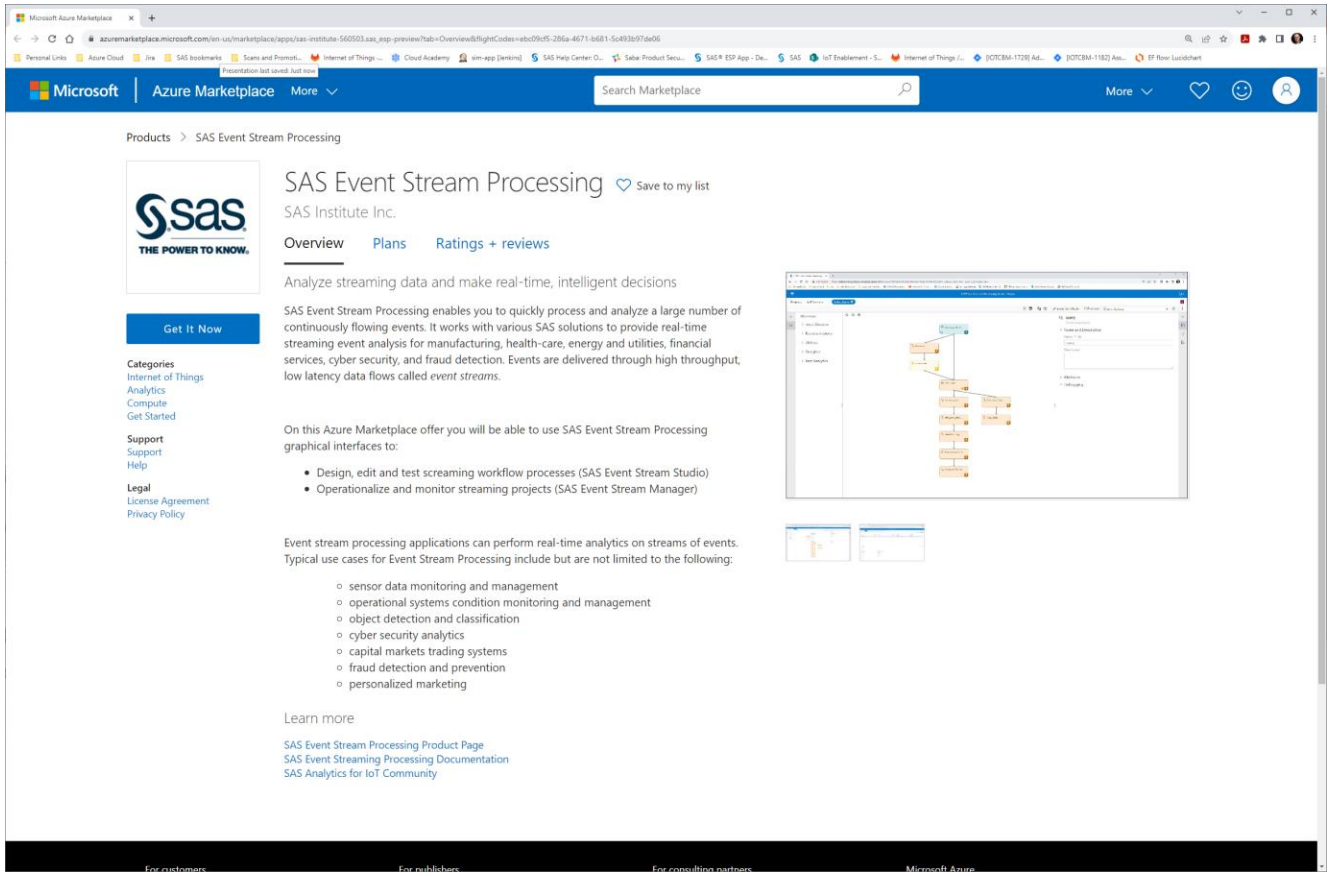

SAS EXPLORE

sas

Thank you!

You can contact us via email:

Daniele.Cazzari@sas.com & Steven.Allan@sas.com

For more information, visit: www.sas.com/esp

SAS EXPLORE

sas