Scalable Edge Cyberinfrastructure for Science-driven AI Workflows

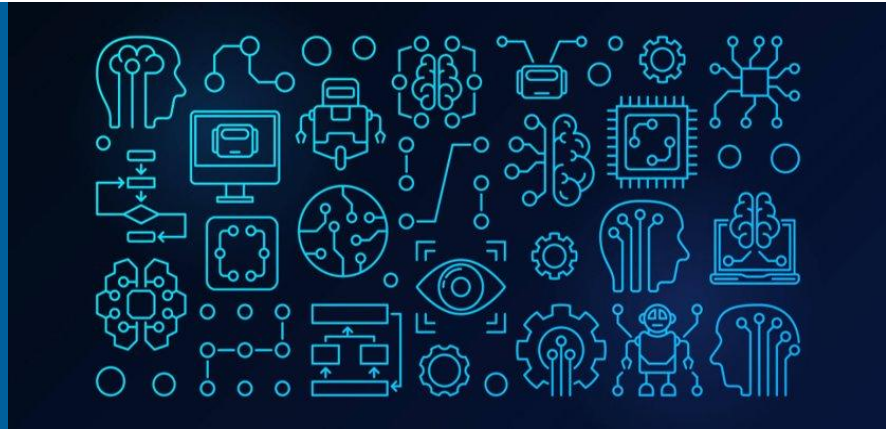

Yongho Kim
yongho.kim@anl.gov
Assistant Computer Scientist
Mathematics and Computer Science division
Argonne National Laboratory

Invited Talk
May 19th, 2023

Wildfire Detection and Prediction

AI@Edge for wildfire detection linked to HPC simulations

Ilkay Altintas, UCSD, Co-PI for SAGE

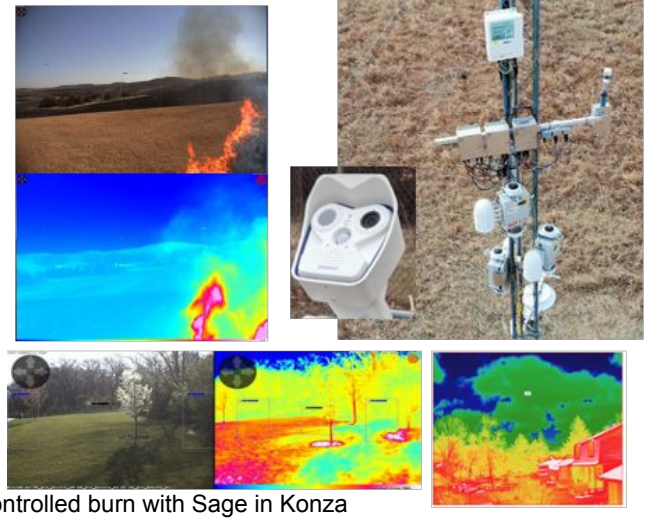

source: controlled burn with Sage in Konza (https://sagecontinuum.org/news/sage-neon-deploy-konza)

Sage project will move Pan-Tilt-Zoom cameras toward suspected outbreaks, and use infrared cameras to build self-supervised AI training

source: WIFIRE: wildfire detection and monitoring (https://wifire.ucsd.edu/)

Edge (Field cameras)

1. Smoke detection at the edge of the network using Camera Imagery

Cloud (SDSC)

2. Wildfire simulations to determine the severity and direction of fires

HPC (Utah)

3. Pollution Concentration maps to support decision-making

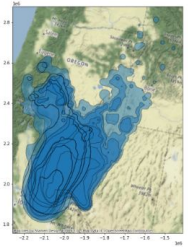

source: a scientific workflow on understanding air pollution caused by wildfires, Daniel Balouek-Thomert, Scientific Computing and Imaging (SCI) Institute, University of Utah

DOE VTO: Understanding traffic type, flow, and density

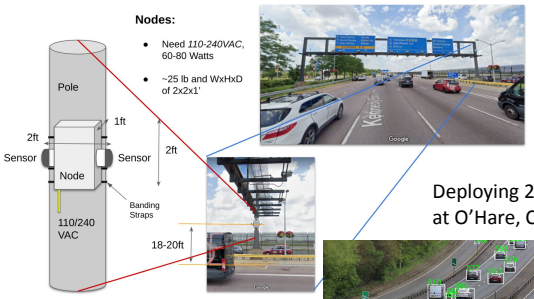

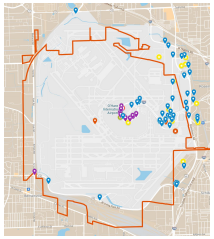

Deploying 25+ at O'Hare, Chicago

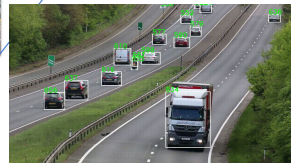

sources: Street view in Google Maps (https://www.google.com/maps)

Dual cameras to capture approaching and leaving traffic for in-situ analysis.

NNSA: Advanced radiation sensors for safer cities

Physical sensors include - Temp, Pressure & Humidity, Rain Gauge, Microphone, Camera, LiDAR, and Rad/Nuc Detector

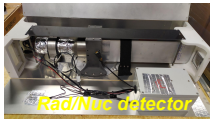

PANDA-DAWN Software Architecture:

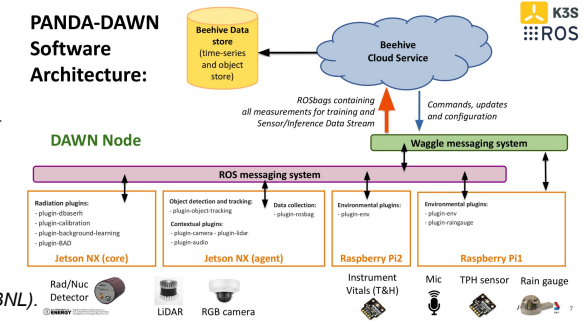

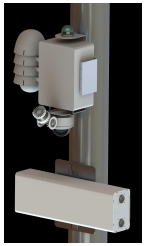

Deploying 20+ Sensors in Chicago

A deployment illustration by Victor Negut (LBNL).

DOE CRADA with Exelon: Advanced Sensors for Grid Stress and Load Forecasting

Simulations predicting locations for sensor deployments.

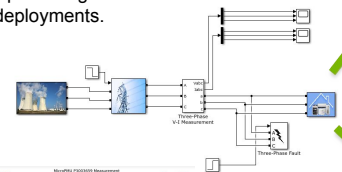

Three-Line-to-Ground Fault

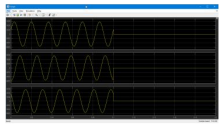

V_a
V_b
V_c

Single-Line-to-Ground Fault

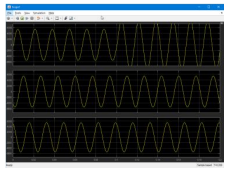

I_a
I_b
I_c

uPMU

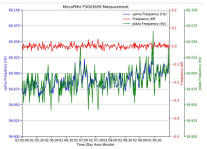

Analysis to quantify stress and forecast load (weather data, traffic data) at the edge.

DOE ASCR - Migratory Computation for the Wireless 5G Digital Continuum

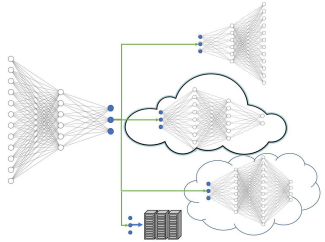

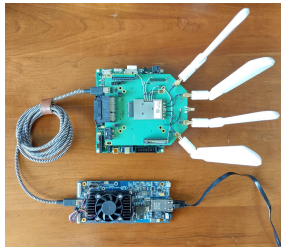

mmWave enabled Waggle nodes will connect with Nokia's NDAC solution with a 5G non-standalone configuration to offload computation on the MEC.

Extend anytime networks to be anywhere/anytime. Explore approaches where ensemble is executed on multiple parts of the continuum, and may use a variety of inputs.

The nodes are about 3 ft tall, 2 ft wide, and 1 ft deep, and wight about 30 lb.

Optical Rain Sensor (HYDREON RG-15)

POE Sky facing camera (HANWHA XNF-8010RV or XNV-8081Z)

POE Sensor Stevenson Shield

Additional Sensors Mounting Location

Relative humidity, barometric pressure, ambient temperature sensor (BOSCH BME680)

NVIDIA Xavier NX GPU, power, wireless communication (4G/WiFi) and node management.

Microphone (ETS ML1-WS)

POE and USB ports for Sensors (additional expansion ports)

110/230V AC Power and Ethernet

POE Ground Facing Camera (HANWHA TEC XNV-8081Z)

Attachable sensors and actuators

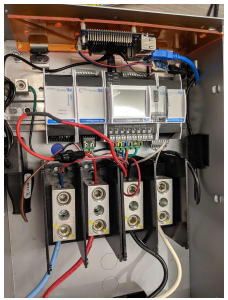

micro synchrophaser

thermal imaging camera with a pan-tilt mount

< image credit: Scott Collis, CROCUS project, Argonne National Laboratory

air-quality sensor

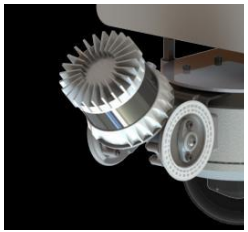

LiDAR

AI/ML Status

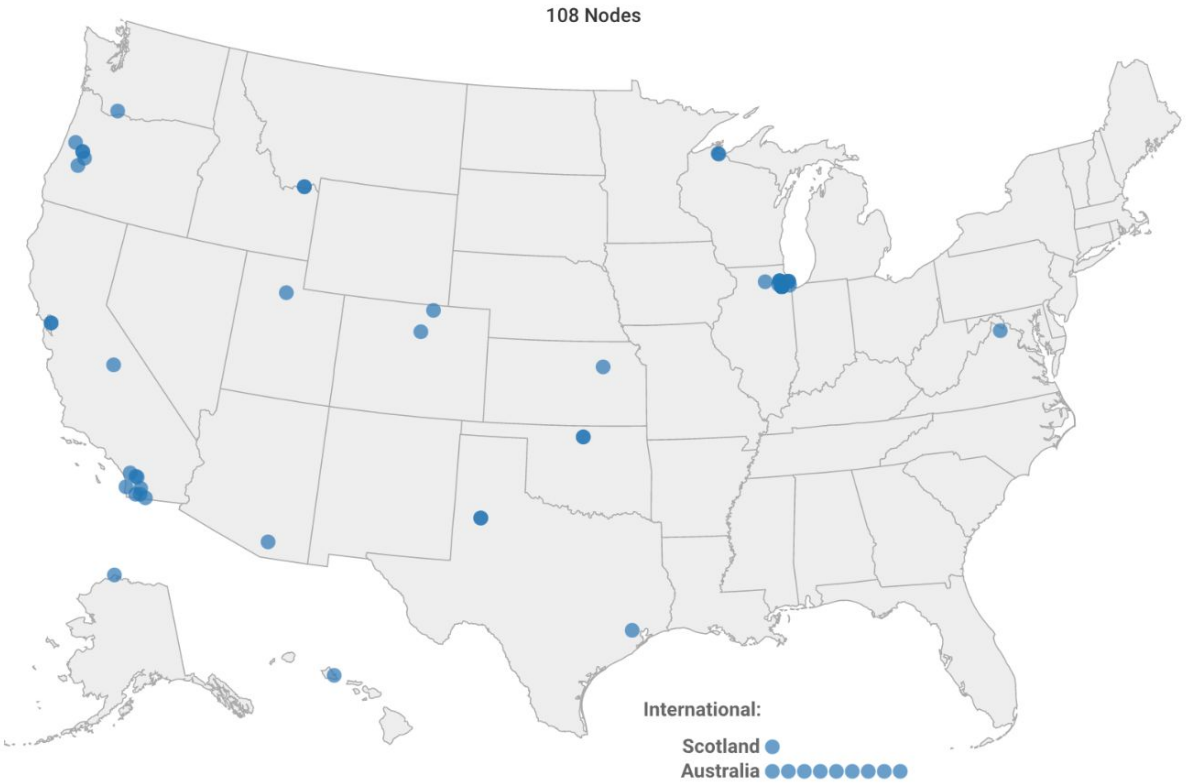

Nodes
108

Active Jobs
62

Recent Apps
16

App Data
1,683,730
Records in the last 24 hours*

Source: https://sagecontinuum.org/

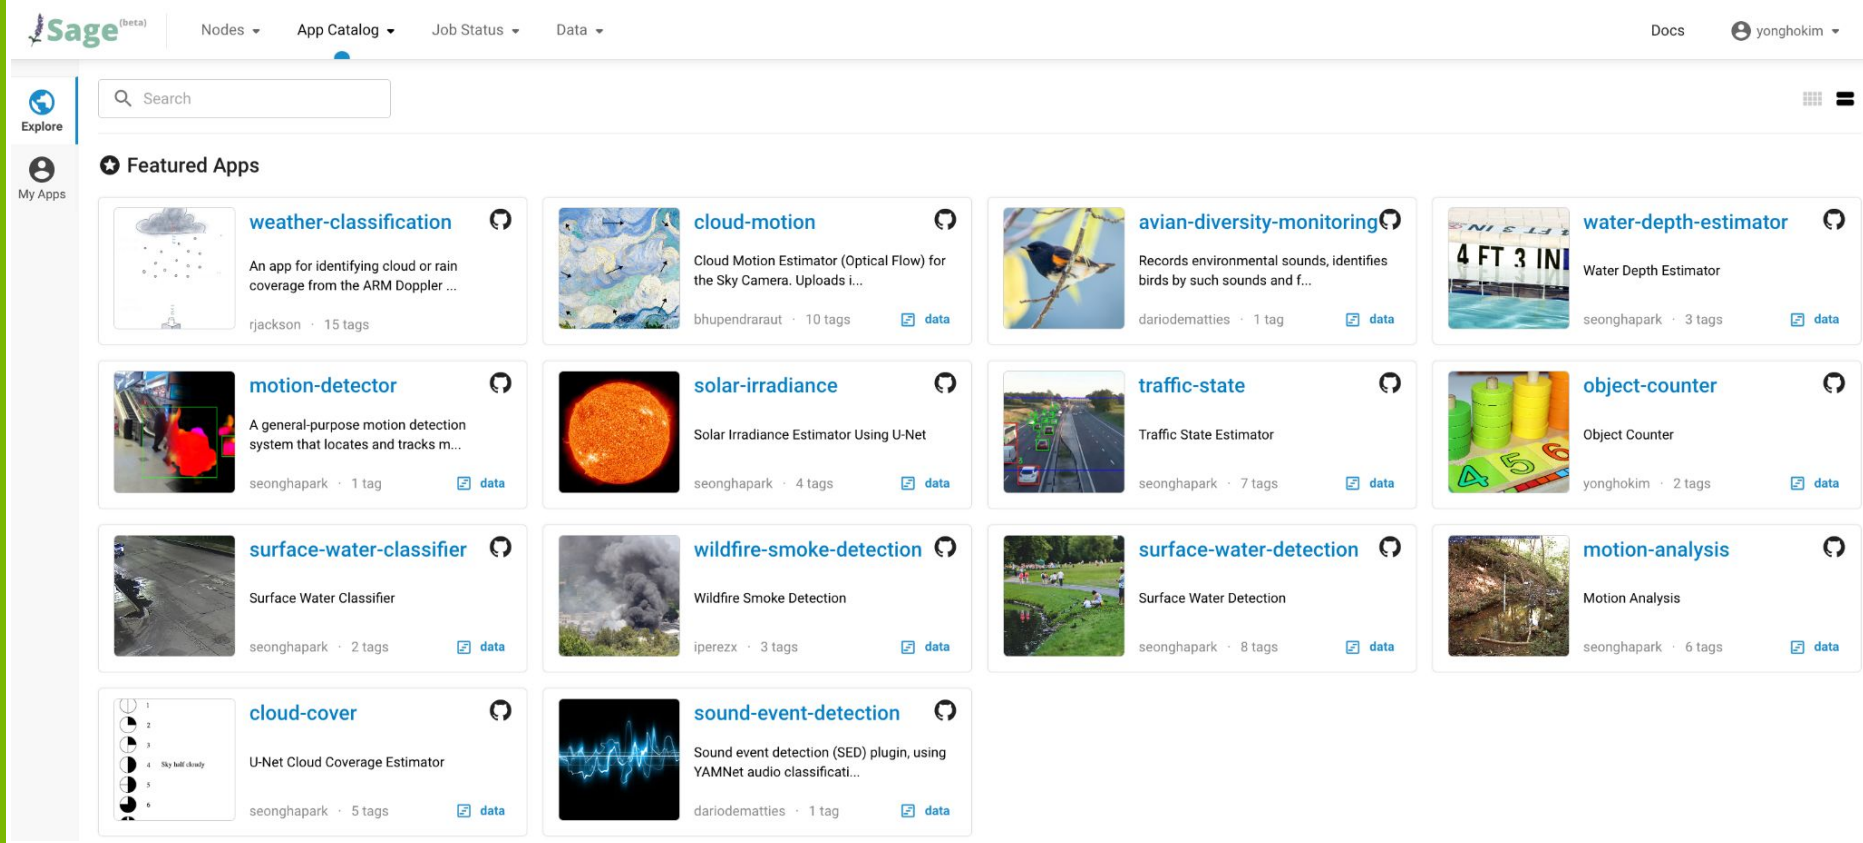

Source: https://portal.sagecontinuum.org/apps/explore

The Waggle Programming Model

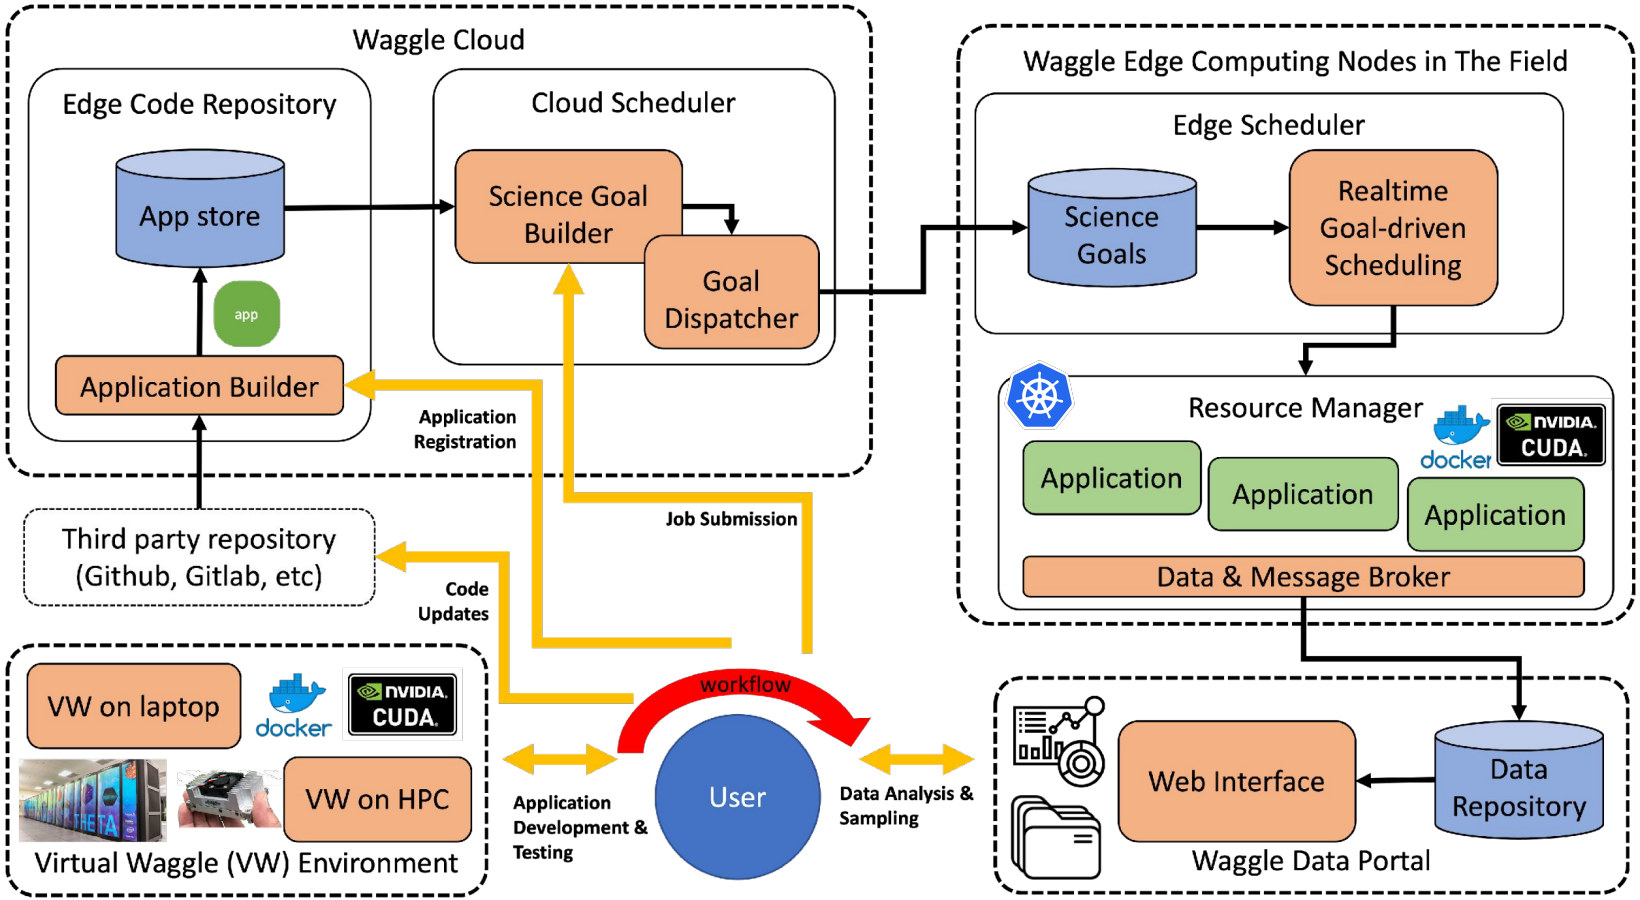

Waggle Software Stack and AI@Edge Applications

- Packaging applications including popular machine learning tools
- The Waggle layer ensures that packaged applications access hardware resources including CUDA on Waggle nodes
- Enabling a multi-tenant environment by isolating computing environment and allocating requested resources to individual application

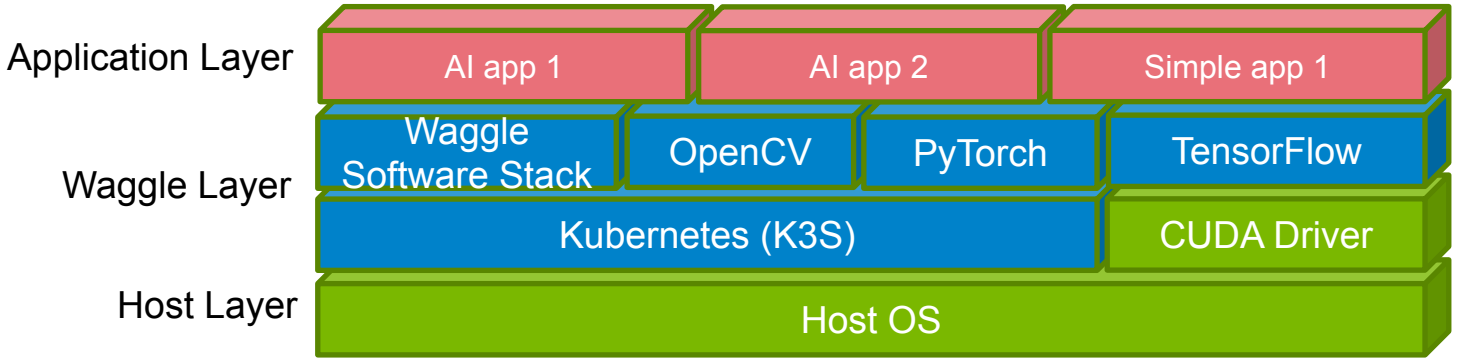

Job Dashboard

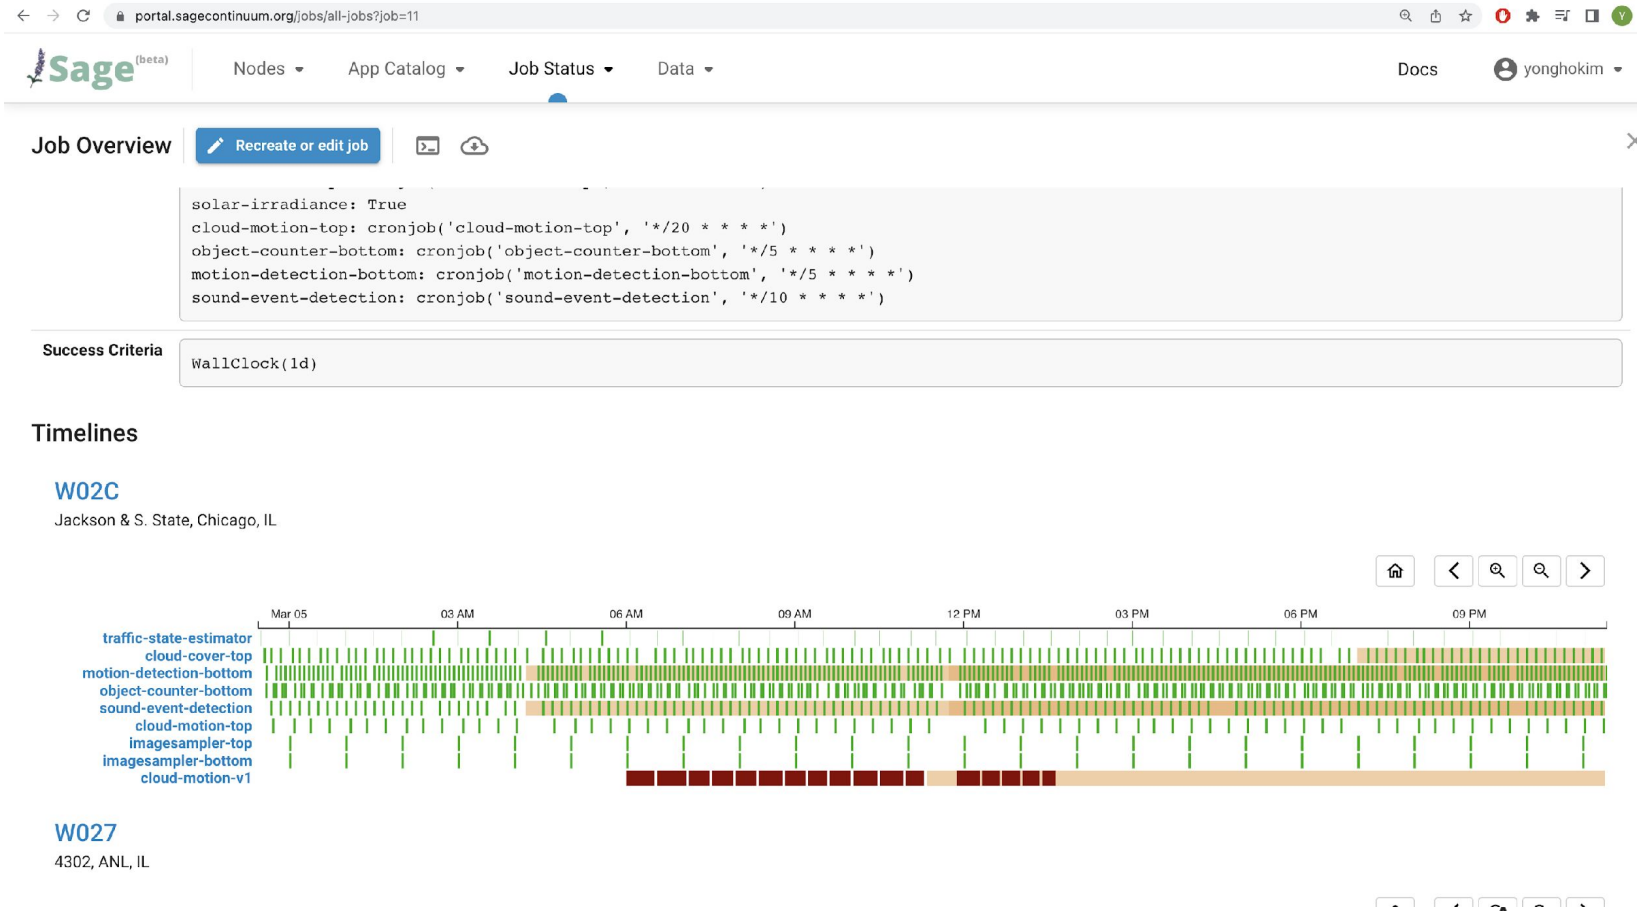

Source: https://portal.sagecontinuum.org/jobs/all-jobs

Viewing Data from Cloud (Waggle Beehive)

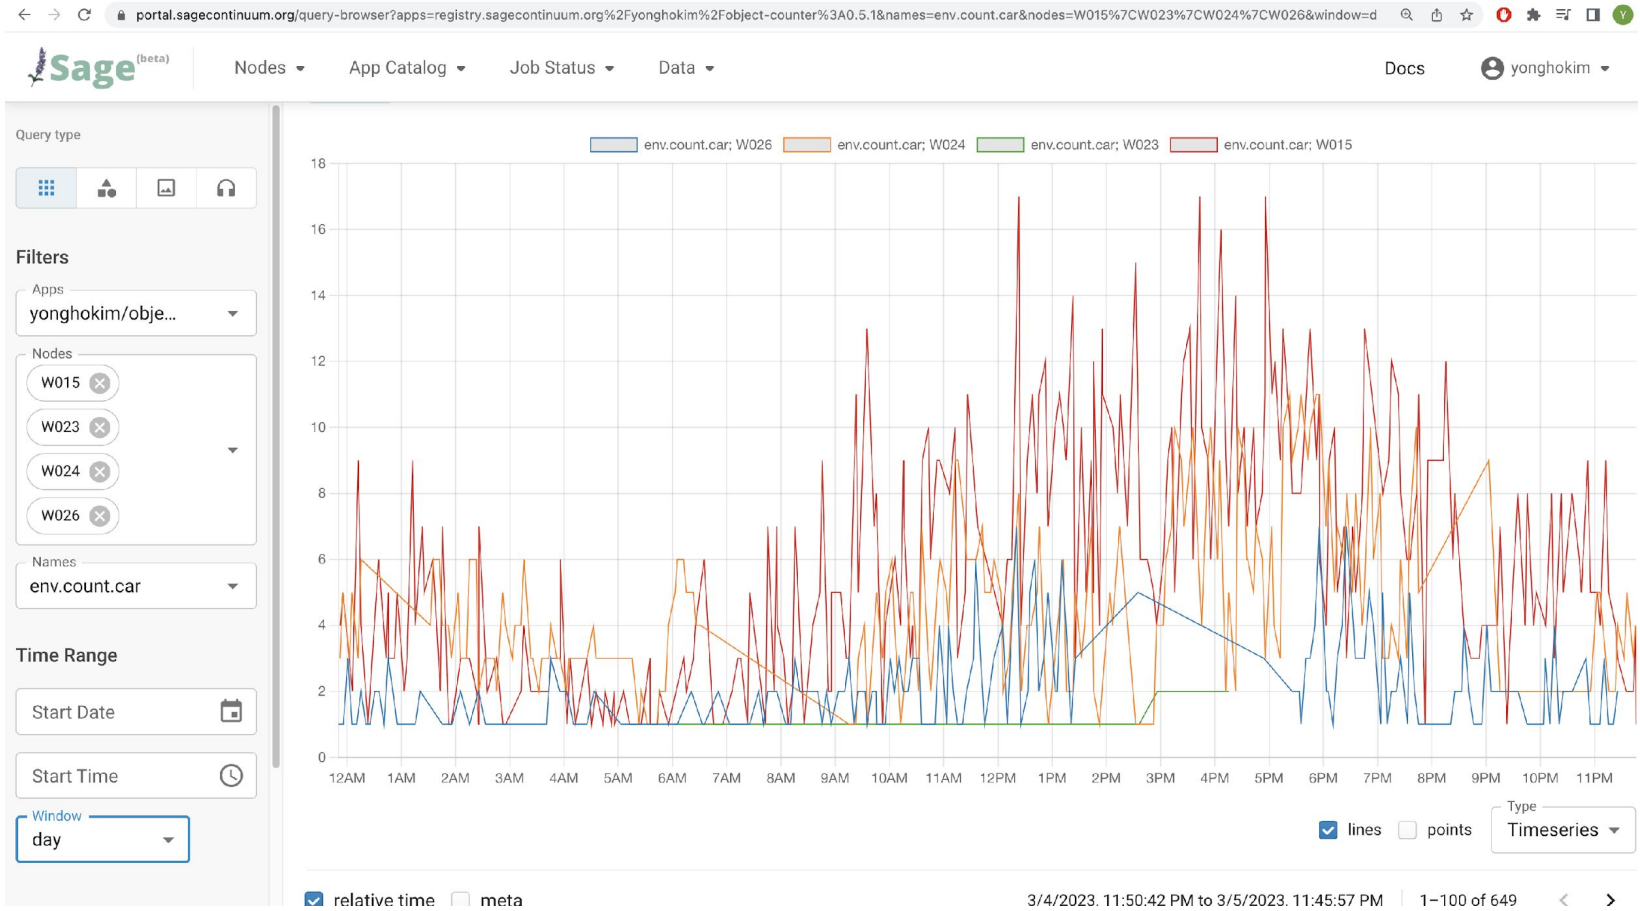

Credit for the Web UI design and implementation: Neal Conrad, Argonne

AI@Edge: Digital Continuum

Sensors

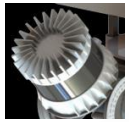

LIDAR

Software Defined Radios

< image credit: Scot Collis, CROCUS project, Argonne National Laboratory

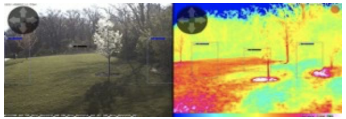

Thermal Imaging

Facilities

NEON tower

NEON mobile deployment platform

Actuators

Dynamic adaptation

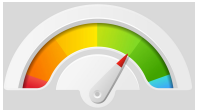

source: Freepik.com

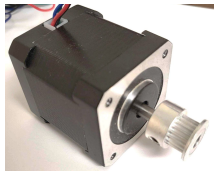

Motors

Edge Computing

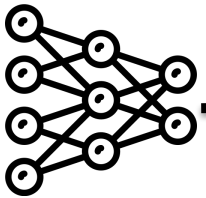

Scientific Data Analysis & Control

Artificial Intelligence
Deep Learning
Inference
Lightweight Training
Autonomous Action

Advanced Networking

New inference (model)
Adaptive controls / steering

Computation

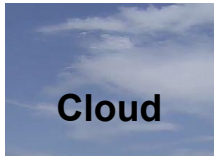

Predictive Sim
Digital Twins
Data Analysis
Machine Learning

^ Image source: Aurora, Argonne Leadership Computing Facility, Argonne National Laboratory (https://www.alcf.anl.gov/aurora)

Challenges for Scalable Edge CI Supporting Science-driven Workflows

- An abstraction to express science-focused workflow execution of applications
  - execution timing and dependencies
- Vertical scaling
  - Self-monitoring and failover mechanisms as edge nodes are resource limited and suffered from dynamic environmental conditions
  - Dynamic parameter tuning supported by applications
- Horizontal scaling
  - Science-level scaling over multiple edge nodes

Waggle Edge Scheduler and Science Goals

- Users submit a job and creates corresponding science goal. The science goal is validated within Waggle context and propagated to target Waggle nodes for execution
- Science goal describes conditions on when and how to run user plugins
- To prevent resource conflicts between user jobs, we schedule user plugins and control their execution on Waggle nodes

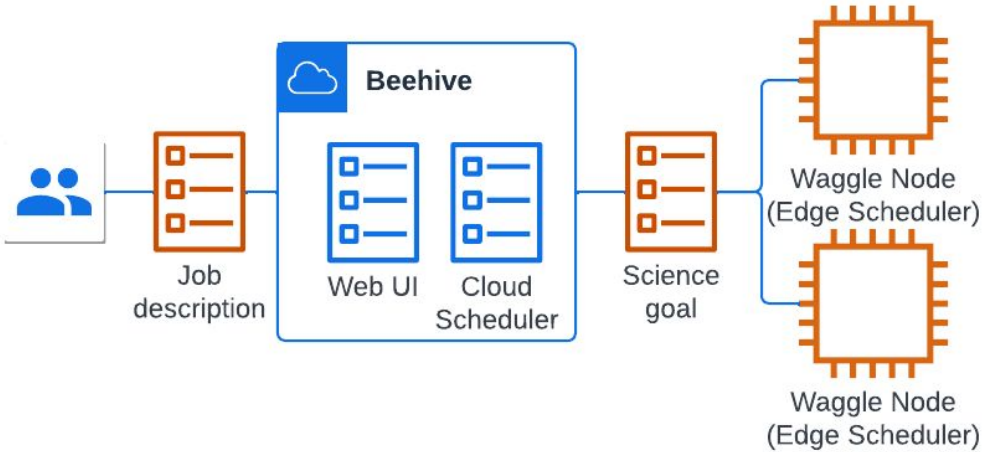

More details: Yongho Kim, et. al, 2022, “Goal-driven scheduling model in edge computing for smart city applications”, Journal of Parallel and Distributed Computing

Life Cycle of A Plugin

- Plugins specified in science goals are downloaded from the edge code repository registry
- The edge scheduler,

  > **promotes** them from Waiting List based on conditions (Waiting > Ready)

  > **schedules** them when the resource is available (Ready > Running)

  > **cleans** them up after execution (Running > Waiting)

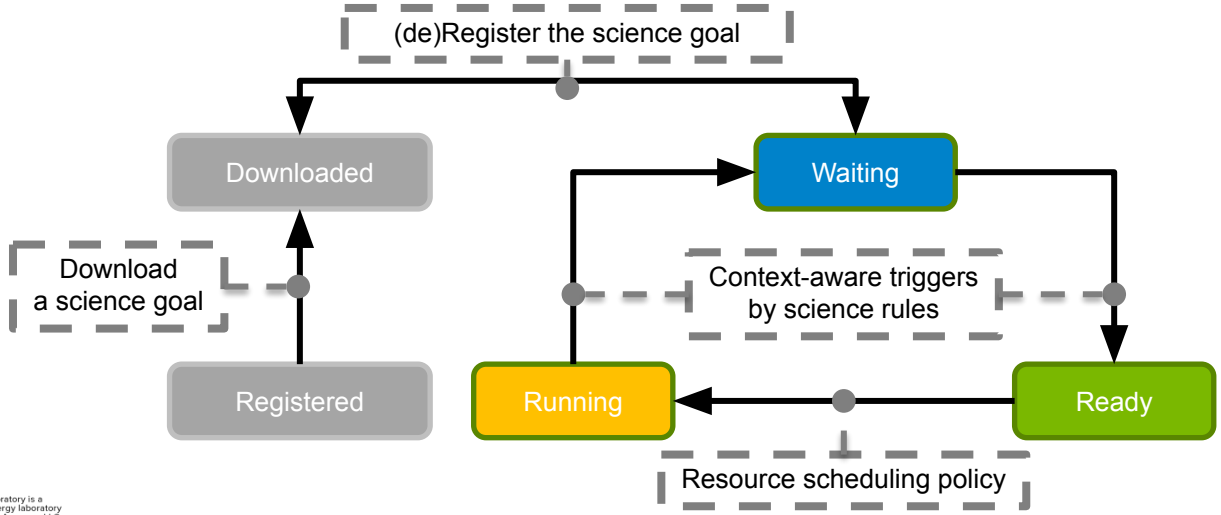

Science Rules for Application Scheduling

- A set of IF-THEN statements instructing the scheduler to server science goal

  - schedule my plugin if it is O'clock
  - schedule my plugin when there are more than 5 cars in the last minute
  - schedule my plugin when the sun rises
  - schedule my plugin if the dependent plugin has run
  - schedule my plugin when averaged noise level exceeds 50 dB

Science rules can also do other actions (*currently in alpha testing*)

- publish "moderate rain" when total accumulation of rain in the last hour is greater than 3 mm
- set my state to "30, 60" for my plugin to continue sweeping the pan-tilt camera

Science rules documentation at https://github.com/waggle-sensor/edge-scheduler/blob/main/docs/sciencerules/README.md

Example Job 1

```
---
name: vto-job
plugins:
- name: image-sampler-left
  pluginSpec:
  ...
- name: image-sampler-right
  pluginSpec:
  ...
- name: object-counter-left
  pluginSpec:
   ...
- name: object-counter-right
  pluginSpec:
   ...
nodeTags:
- WSN
- VTO
scienceRules:
- "image-sampler-left: cronjob('image-sampler-left', '0 * * * *')"
- "image-sampler-right: cronjob('image-sampler-right', '0 * * * *')"
- "object-counter-left: cronjob('object-counter-left', '15/45 * * * *')"
- "object-counter-right: cronjob('object-counter-right', '0/30 * * * *')"
successCriteria:
- WallClock(1d)
```

<< name of the job

<< list of plugins to run

<< specification of the plugin

<< list of nodes (all Waggle nodes under a project)

<< conditions on which the plugins run

<< run every hour

<< criteria on when the job is considered as complete

Example Job 2

```
---
name: vto-video-sampler
plugins:
- name: video-sampler-left
  pluginSpec:
  image: registry.sagecontinuum.org/theone/video-sampler:0.2.4
  args:
  - -stream
  - rtsp://10.31.81.16:554/0/profile6/media.smp
  - -duration
  - 60
nodes:
  W023:
scienceRules:
- "video-sampler-left: cronjob('video-sampler-left', '*/5 12,13,14,20,21,22 * * *')"
successcriteria:
- WallClock(1d)
```

<< plugin Docker image from ECR

<< list of nodes; Waggle node W023 is specified

^ run every 5 minutes in rush hours

Example Job 3

```
---
name: water-detection
plugins:
- name: water-detector
  pluginSpec:
  image: registry.sagecontinuum.org/seonghapark/surface-water-detection:0.0.6
  selector:
    resource.gpu: true
nodeTags:
- WSN
- raingauge
- camera_bottom
scienceRules:
- "water-detector: rate('env.raingauge.total_acc') > 3 and cronjob('water-detector', '*/10 * * * *')"
successcriteria:
- WallClock(1d)
```

<< the plugin requires GPU resource to run

<< list of nodes; all Waggle nodes that have a bottom-facing camera and rain gauge

^ run every 10 minutes if it is raining

- More science rule functions can be found:

  https://github.com/waggle-sensor/sciencerule-checker

Scheduling Policies

- Scheduling policies select the best plugins to run

  > round-robin: selects the most starving plugin since its last execution

  > run-all: selects all plugins triggered by science rules

  > gpu-aware: selects GPU-demand plugins when there is no GPU-demand plugins running

  > and more!

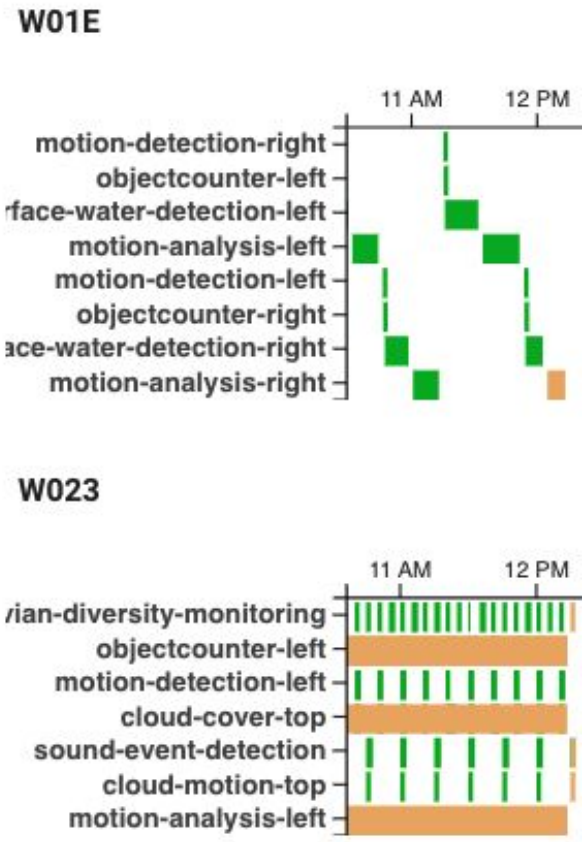

Job status on Sage nodes. The node (W01E) uses round-robin scheduling and the other (W023) is based on science rules-driven scheduling

Automated Vertical Scaling inside A Node

- Edge scheduler needs performance metrics of applications for resource-aware scheduling. Currently all applications are limited by 1000 millicore (1 CPU) and 1 GB memory
- Prototyping performance metrics storage server that constantly monitors performance of application containers
- Application developers can also use this service to ensure the application works on edge hardware

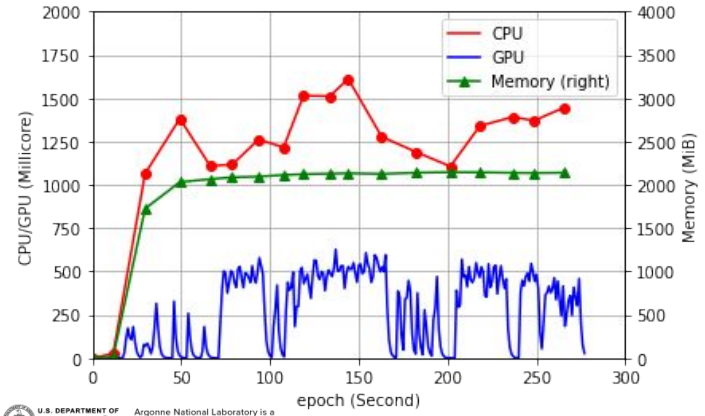

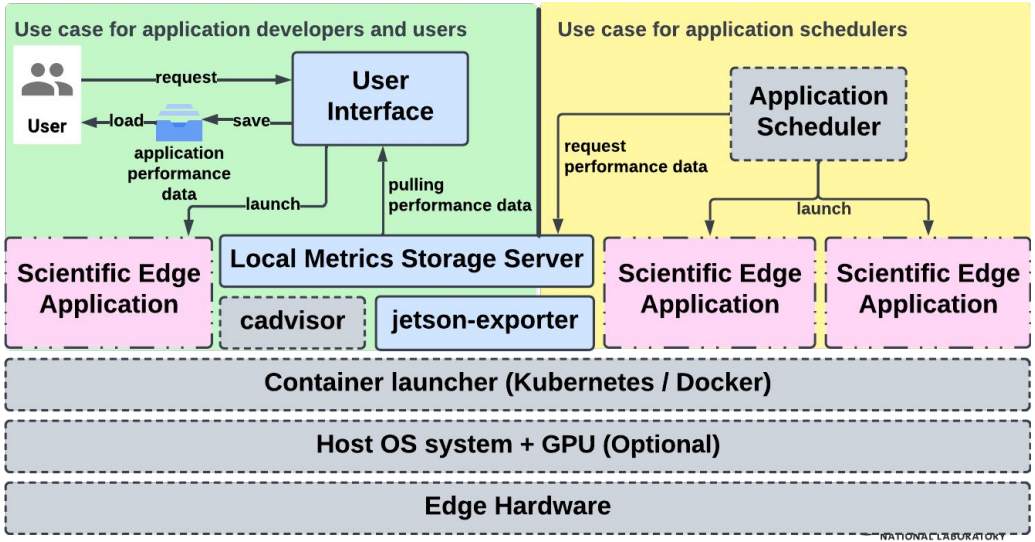

Application Abstraction: Sidecar Approach to Control Applications

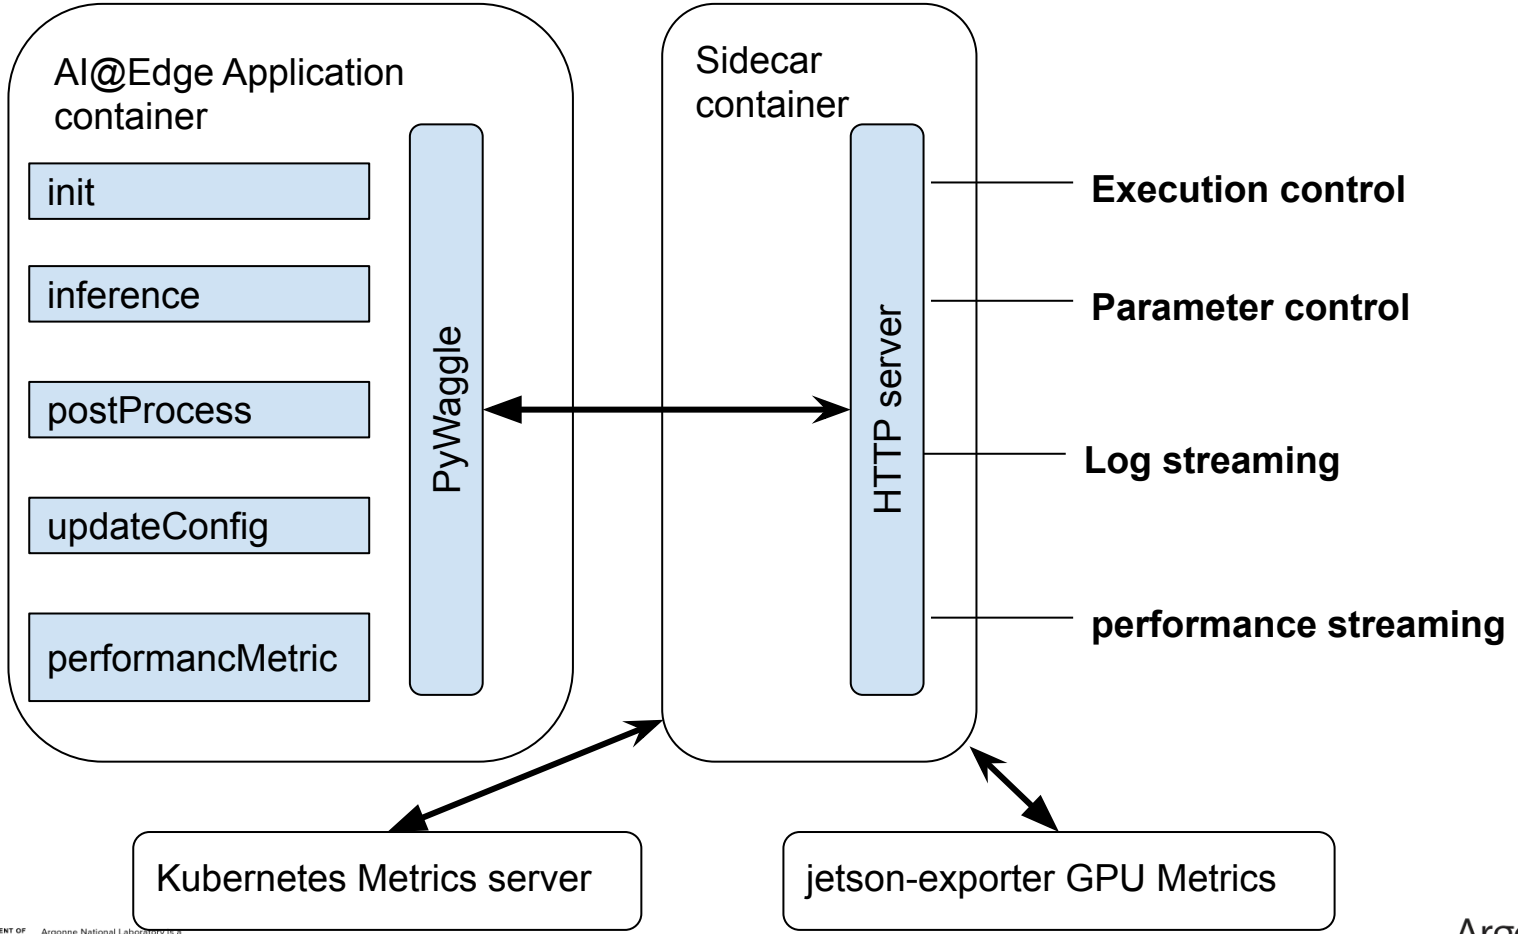

Horizontal Scaling of Jobs across Edge Nodes

- Scientific workflows may specify quantitative requirements, e.g., data volume, number of application executions per node

```
---
name: water-detection
plugins:
- name: water-detector
  pluginSpec:
  image: registry.sagecontinuum.org/seonghapark/surface-water-detection:0.0.6
  selector:
    resource.gpu: true
nodeTags:
- WSN
- raingauge
- camera_bottom
scienceRules:
- "water-detector: rate('env.raingauge.total_acc') > 3 and cronjob('water-detector', '*/10 * * * *')"
successcriteria:
- DataQuantity(1000, 24h)  << specifying desired quantity of measurement to decide scaling
- WallClock(30d)
```

Horizontal Scaling via Cloud and Edge interactions

- Edge schedulers interact with the cloud scheduler / HPC simulations to understand and drive itself towards the global goal
- Edge updates their local measurements
- Cloud triggers HPC simulations and updates science rules in the job
- HPC simulation may send predicted events to edge schedulers to change their scheduling behavior