2022 CSPP USER'S GROUP, UNIVERSITY OF WISCONSIN MADISON

Using CSPP in Event-Driven Workflows on Amazon Web Services

Rus Healy
AWS Ground Station Specialist SA
Aerospace & Satellite
AWS

Agenda

Why CSPP for Event-Driven Architectures?

AWS Ground Station Primer

What the AWS Aerospace & Satellite Team Built, and Why

Deep Dive—How we Architected and Deployed CSPP

Opportunities

How You Can Try It

About AWS for Space

Why CSPP for Event-Driven Architectures

Why CSPP for Event-Driven Architectures?

CSPP, specifically VIIRS SDR and Polar2Grid, are straightforward to containerize

CSPP is much faster than IPOPP in the same compute environment

CSPP works well with AWS services such as EC2, EBS, and CloudWatch logs

Open-source code

Because containers are cool

Scalability, performance, cost optimization

AWS Ground Station Primer

AWS Ground Station Primer

AWS managed service that connects customers to their satellites in S band and X band

Downlink satellite payload and telemetry data directly into AWS

Uplink satellite telecommand and control (S band)

Uses the AWS global network

Centrally process downlink data from many global AWS Regions

AWS Ground Station Primer

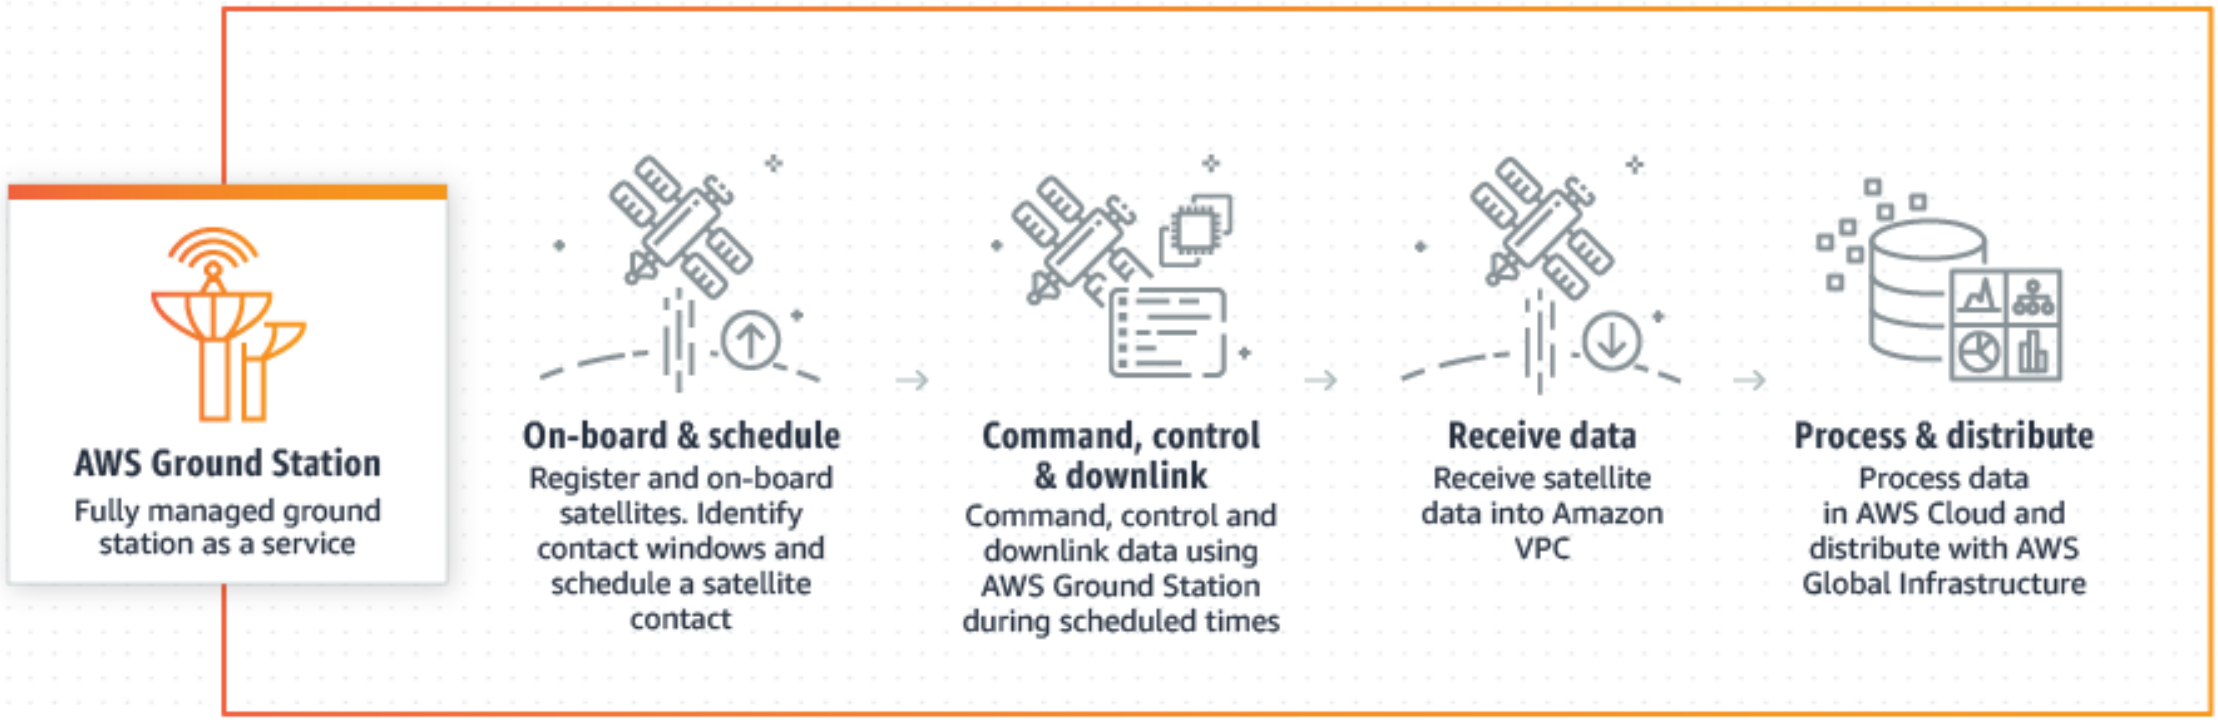

What the AWS A&S Team Built, and Why

Event-driven architecture for image processing

Processes JPSS-1 Level 0 output to Level 2B/3

Processing is triggered by the arrival of new L0 product

Output triggers AWS Machine Learning service to create cloud overlay

Showcases “the art of the possible” using real satellite imagery

The architecture is the basis for a workshop (more on this later)

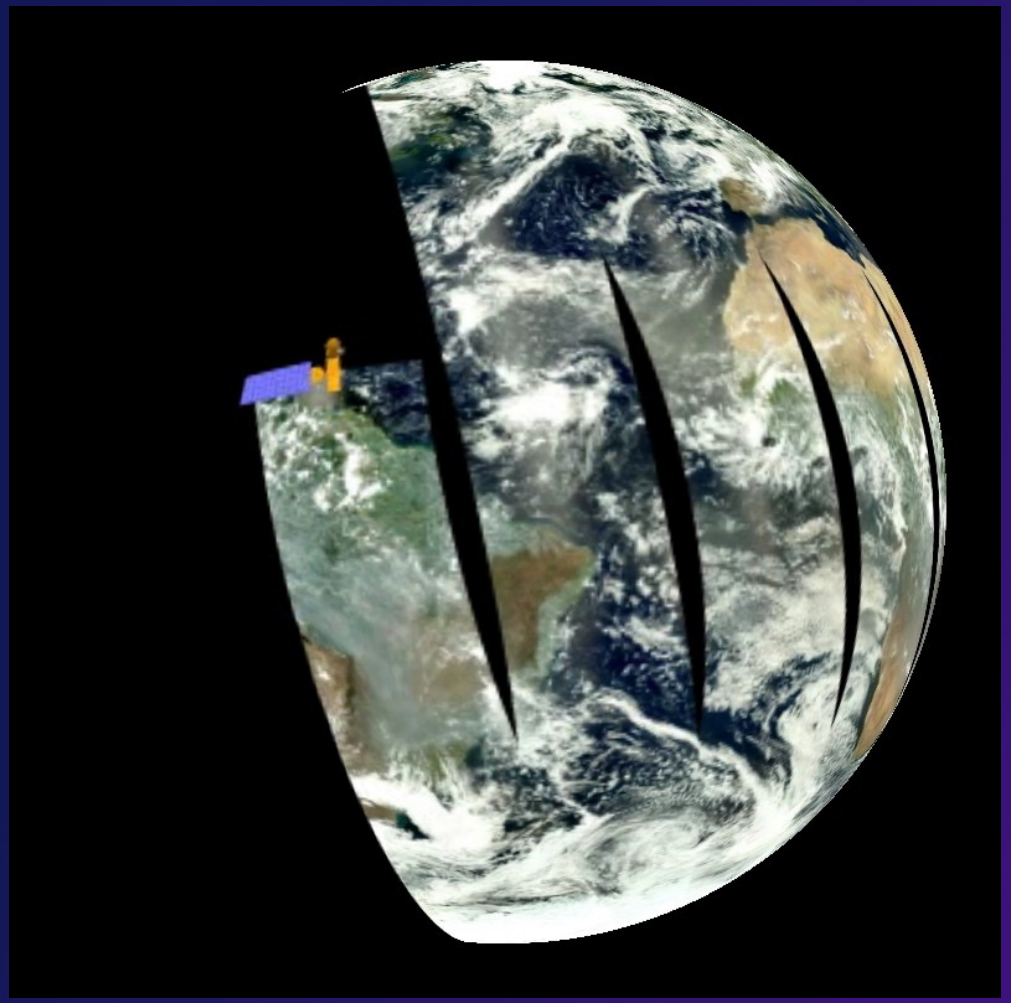

Deep Dive—How We Architected and Deployed CSPP

Design Choices and Deployment

Built on Amazon EC2 compute instance running CentOS and Docker

- Choice: EC2 over Fargate because we wanted to use 8 CPU cores to process images more quickly

Five tarballs installed on the container

Hard part/complexity: attached disks

Easy part: Lookup tables and tiles compressed well

- Which is a challenge later!

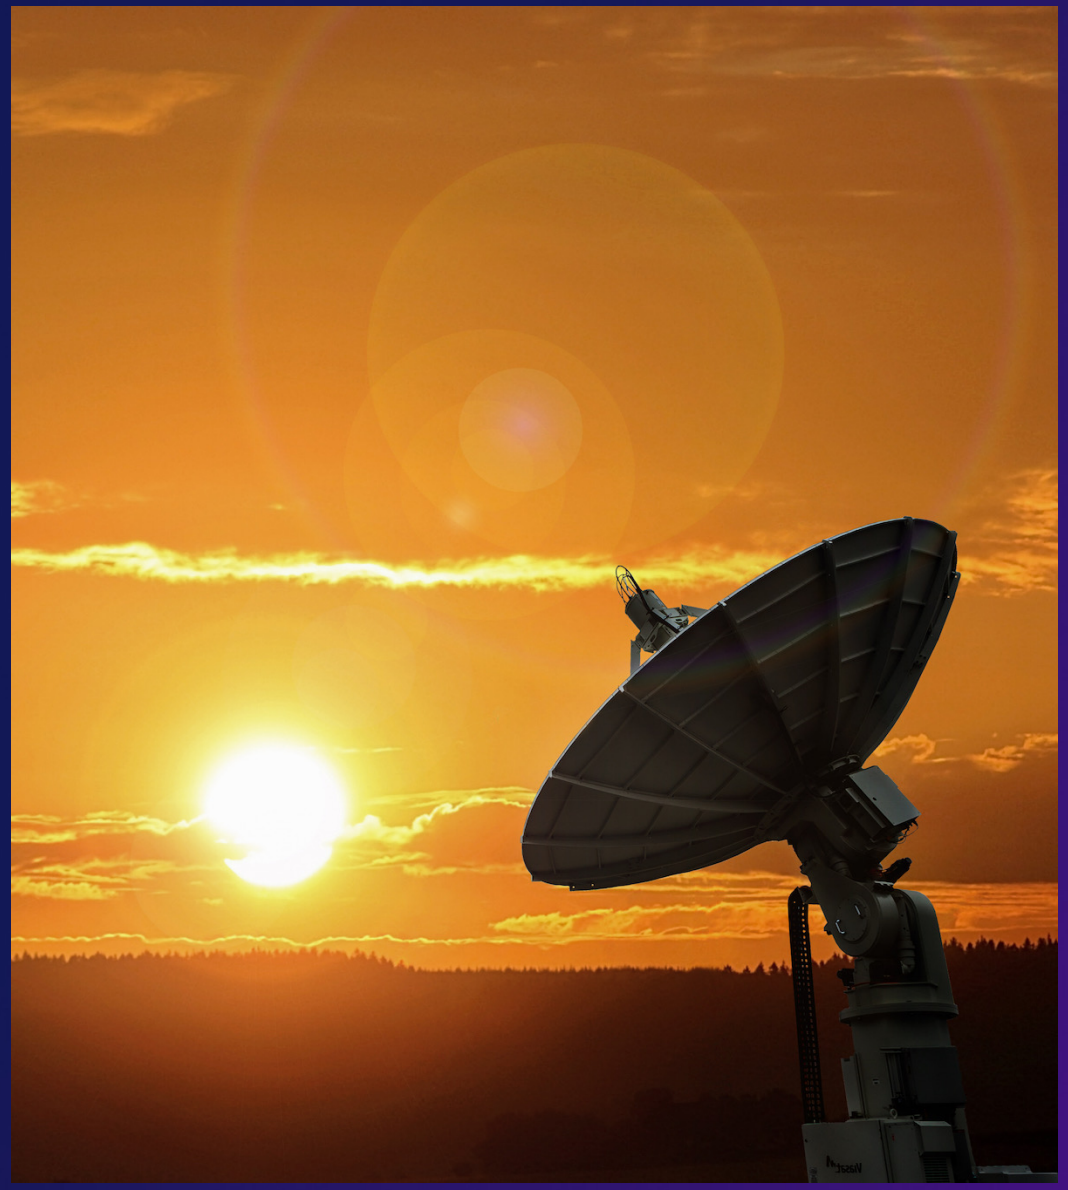

Deployment Deep Dive

- Ancillaries (spacecraft parameters, orbital information)—the software tries to download these on the fly
- Lookup tables—more than 100 GB after decompression
  - Rather than install these on the container, we put them on an external volume manually
  - There’s no way to automate that process as part of the Docker build process
  - LUTs are stored on an Amazon EBS snapshot
    - The Amazon Batch job creates an EBS volume from the snapshot at launch time, and attaches that volume to the container at a specified mount point

Deployment Deep Dive

We wind up with a stripped-down container with only the application binaries

Storage options that would be more optimal:

- For a production system, today: Amazon Elastic File System (EFS)
  - Fast, persistent storage that can be mounted by multiple containers simultaneously
  - EFS would allow for faster deployment and much better scalability
  - LUTs could be updated at any time, rather than updating a volume and taking a new snapshot as is done with EBS snapshots
- Future, more optimal cost-wise: Amazon S3
  - Ideal for object storage and considerably less expensive
  - Today CSPP doesn’t support the S3 API

Event-Driven Architecture: Data Flow

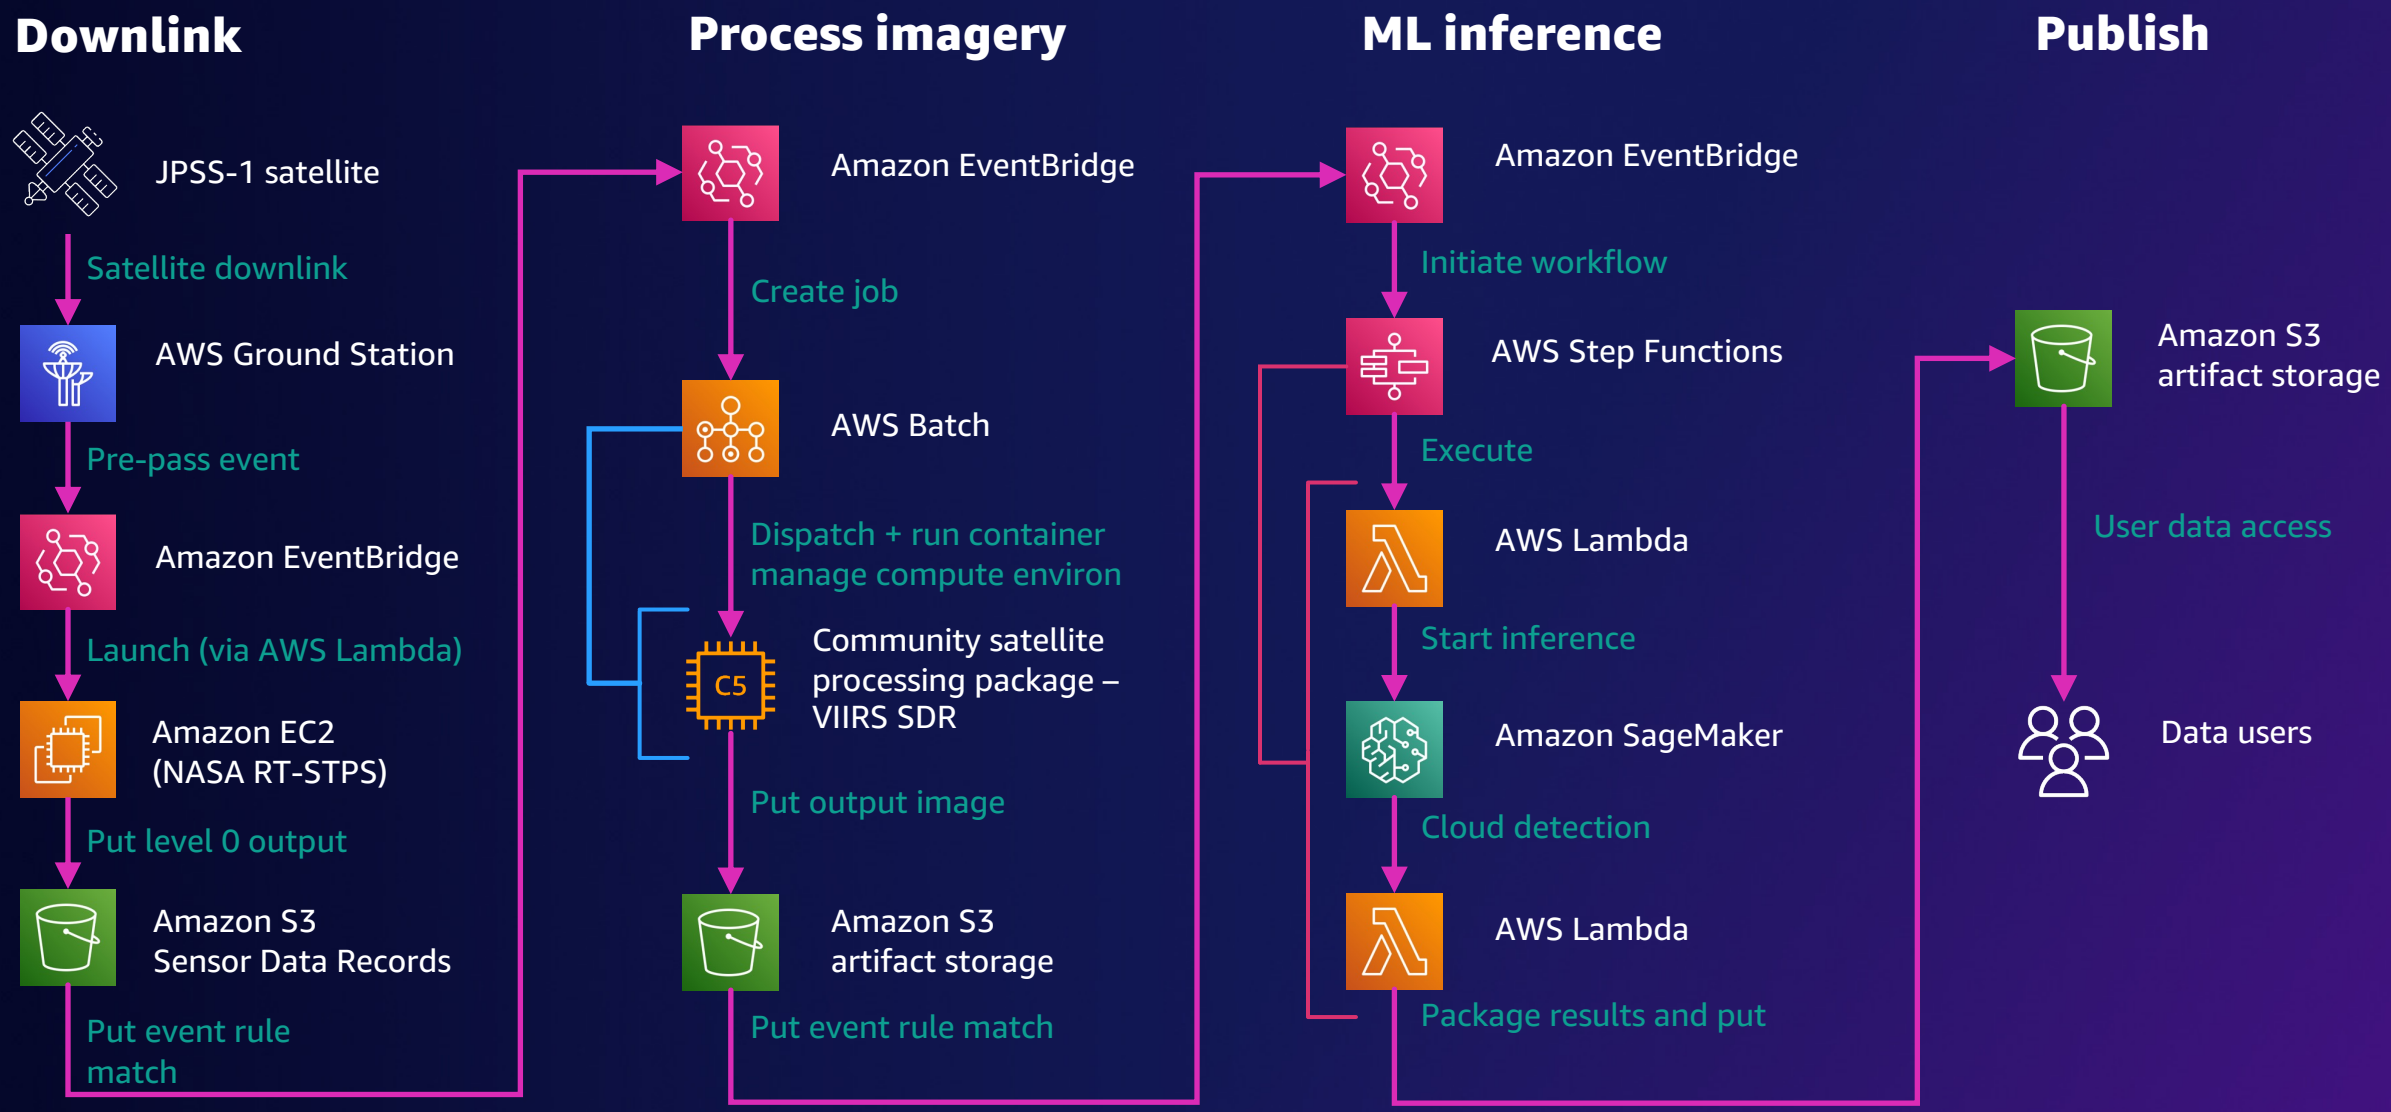

Event-Driven Architecture: Logical View

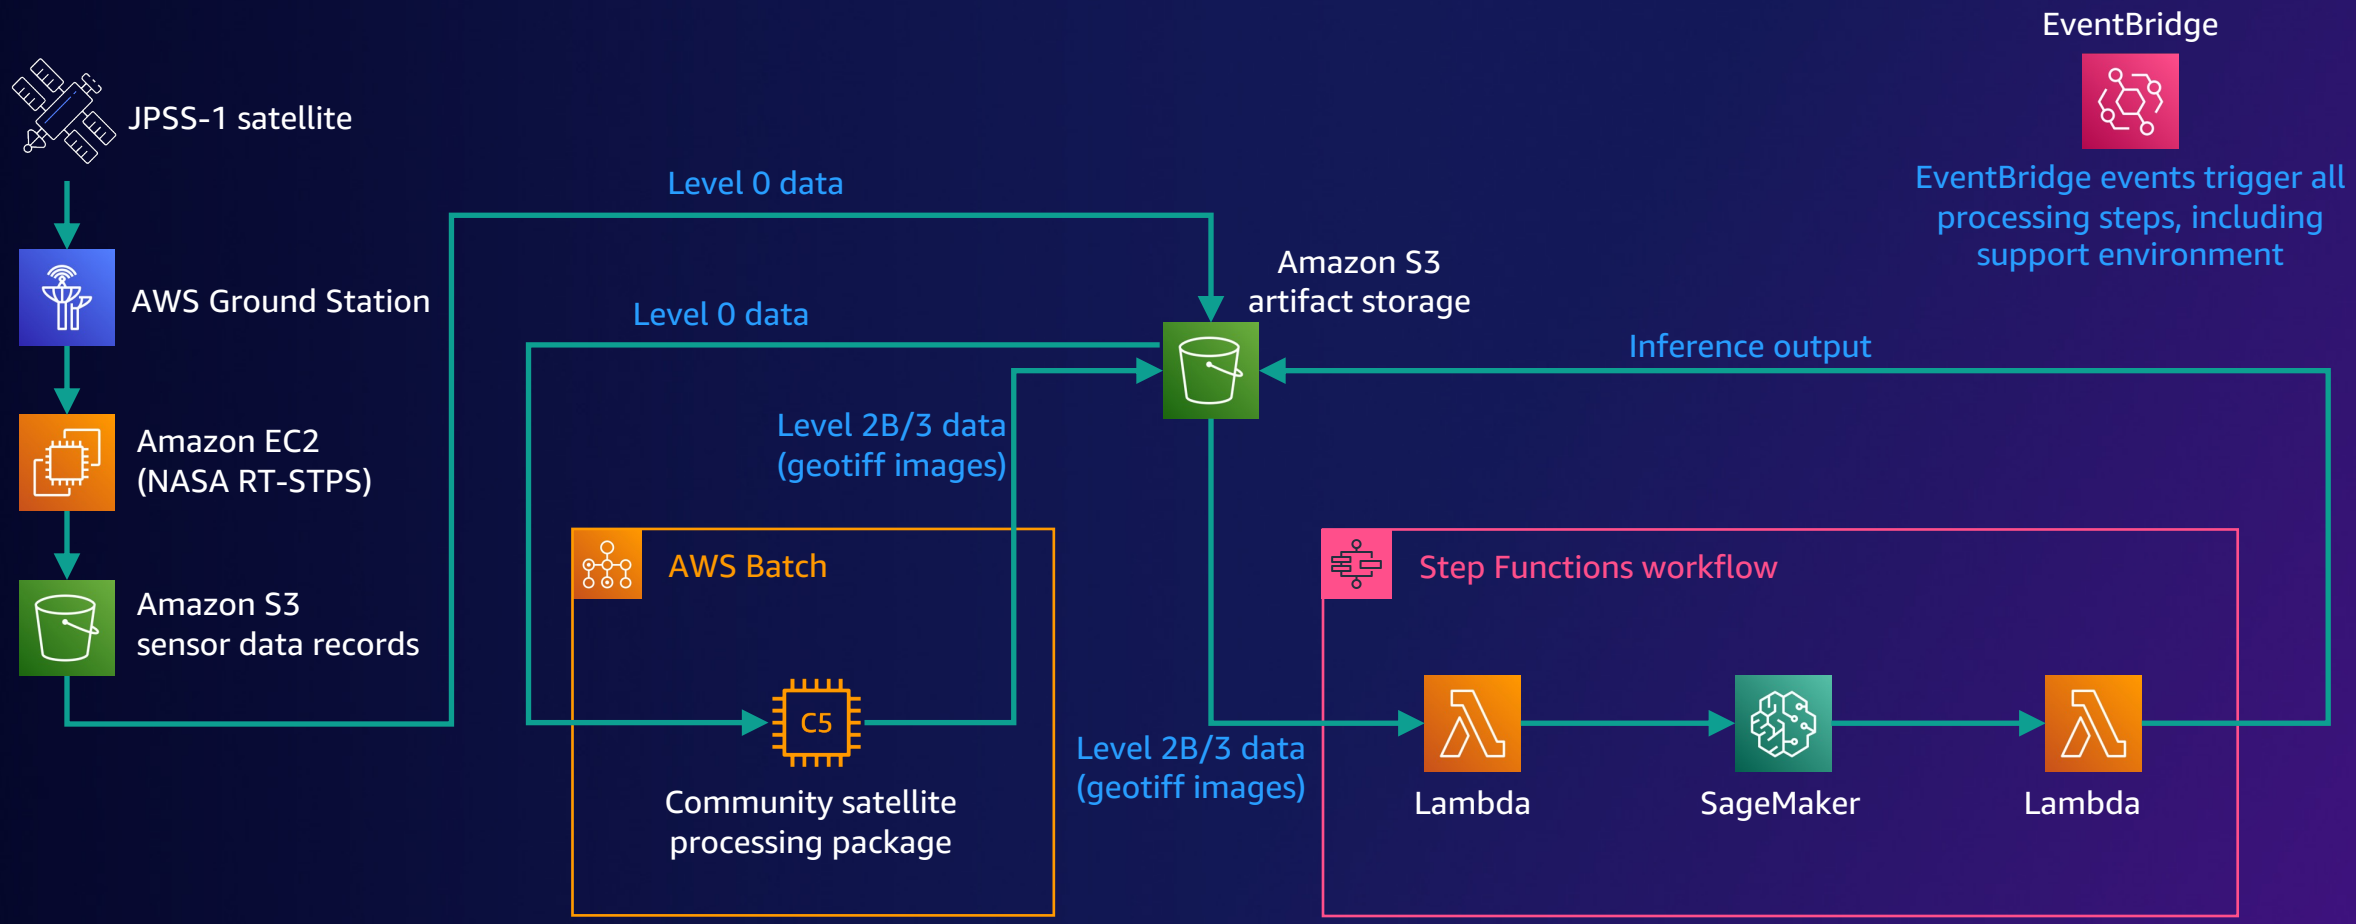

What the Workshop Delivers

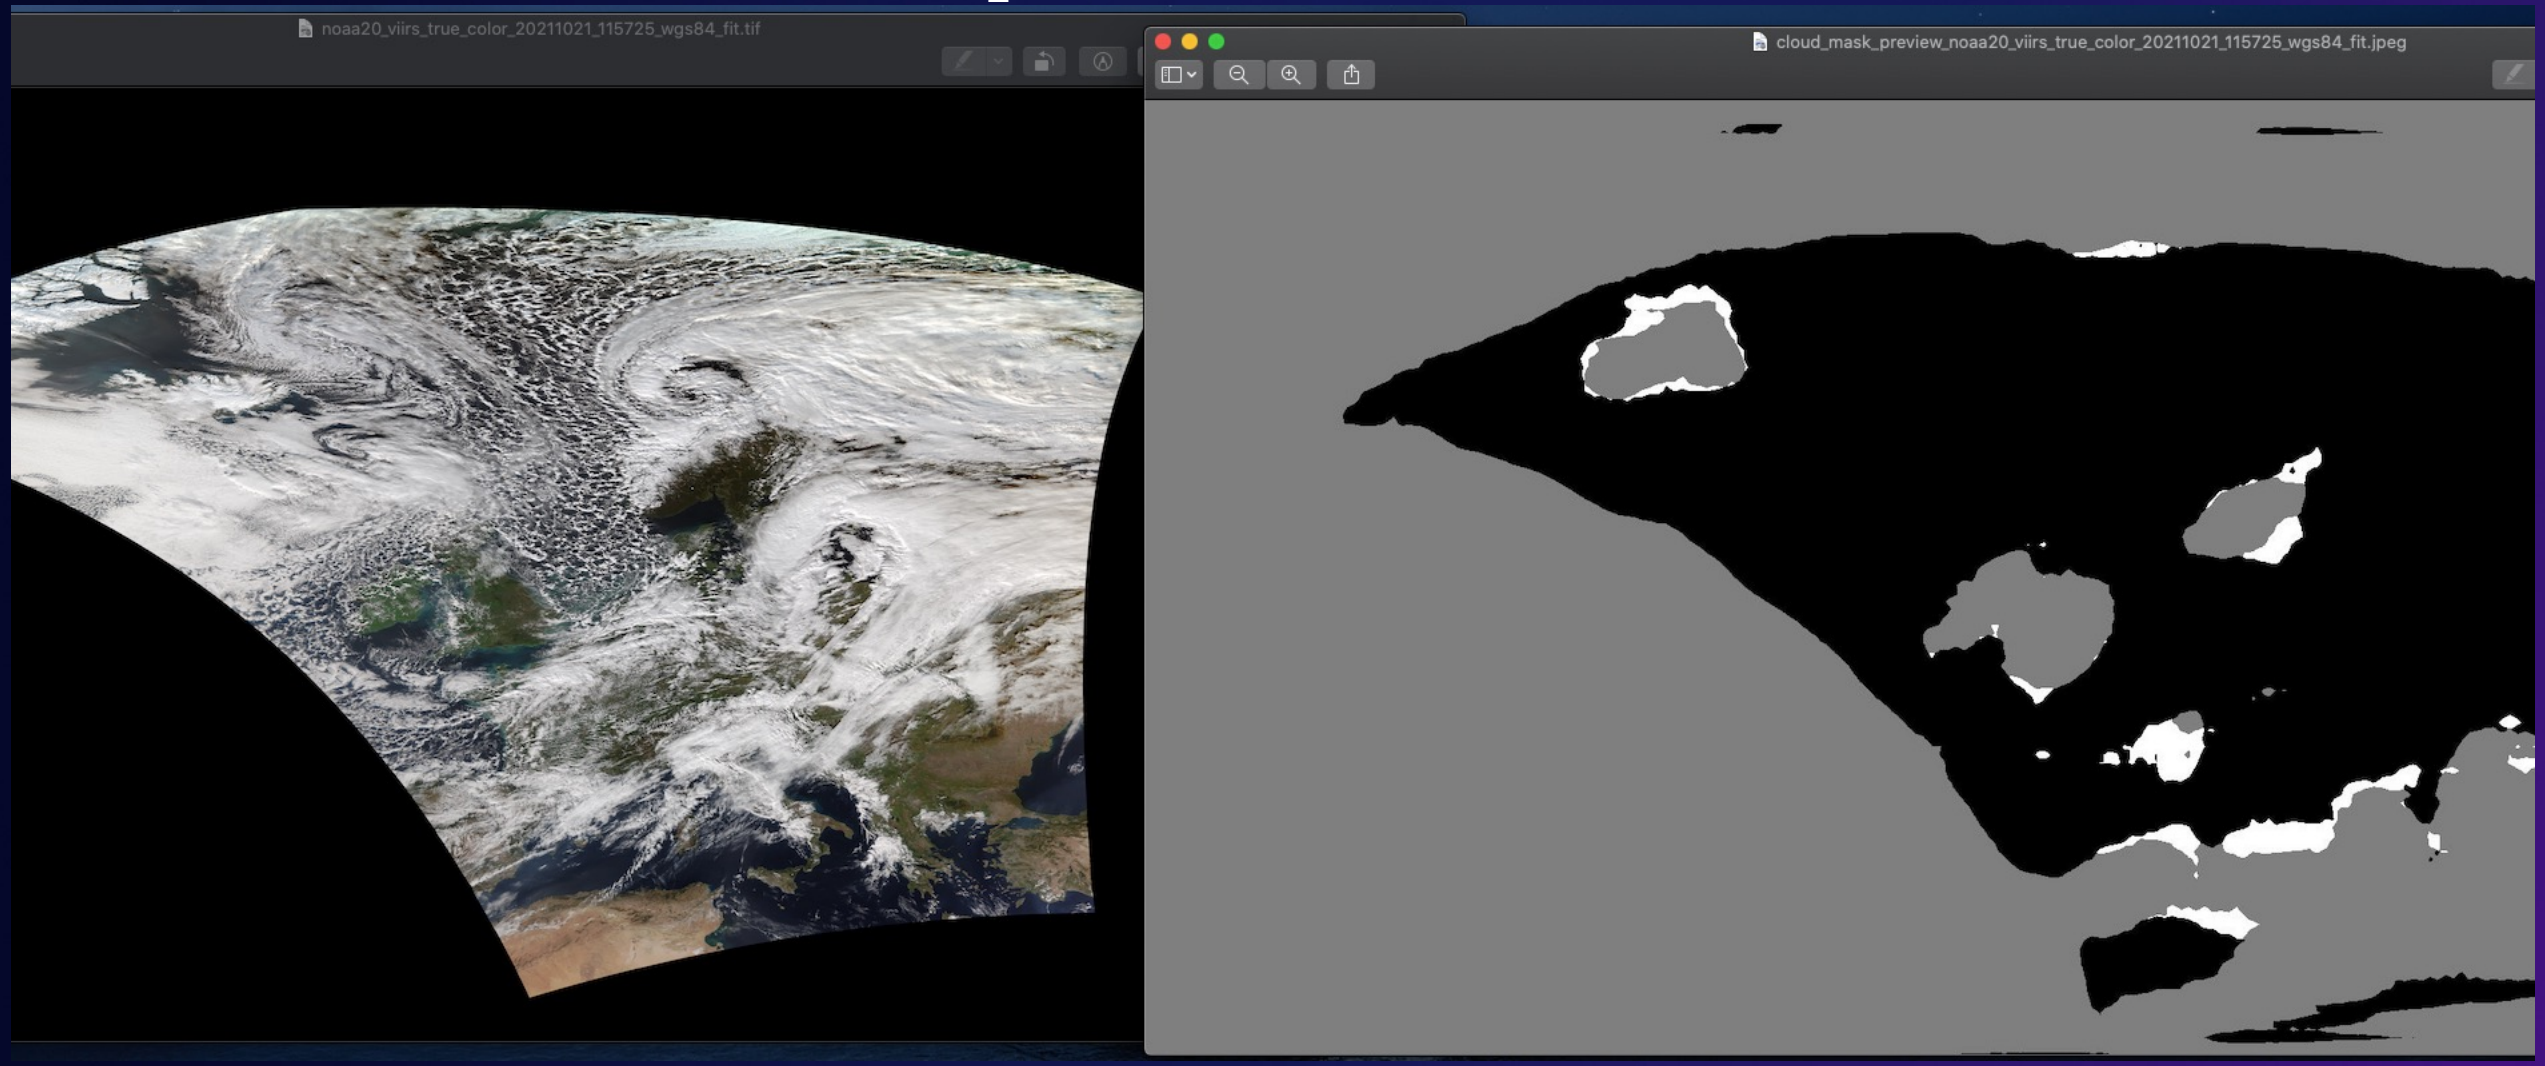

Opportunities

Opportunities

- Moving the storage to Amazon Elastic File System (EFS) is an option
  - Allows multiple containers to mount the same file system
  - The software would have to support it
- AWS Batch jobs run against each container
  - To make it really scalable, a Batch job could be created for each granule and process one granule per container. How parallel they run would be based on the size of the cluster.
  - Is there a post-processing “gather” process that would constrain the architecture or its performance?
- Better understand what viirs_sdr is doing under the hood

How You Can Try It

How You Can Try It

Visit the Workshop website, get familiar with building

- Step-by-step instructions and details

Coordinate scheduling a workshop for your team

Ask us about additional options

- rushealy@amazon.com

Additional Ways to Explore AWS Ground Station

Theory is great, but hands-on experience is better!

Getting started: Implement an Earth observation pipeline using AWS Ground Station

Advanced: Implement a more advanced Earth observation pipeline

Learn more about solution architectures for DigIF solutions

Connect with your AWS account manager or contact us through our website

"There is no compression algorithm for experience."
Werner Vogels, VP and CTO
AWS

Learn in-demand AWS Cloud skills

AWS Skill Builder

Access **500+ free** digital courses and Learning Plans

Explore resources with a variety of skill levels and **16+** languages to meet your learning needs

Deepen your skills with digital learning on demand

Train now

AWS Certifications

Earn an industry-recognized credential

Receive Foundational, Associate, Professional, and Specialty certifications

Join the **AWS Certified community** and get exclusive benefits

Access **new** exam guides

aws

Thank you!

Rus Healy

rushealy@amazon.com