aws

Architecture Monthly

October 2021
Aerospace

EDITOR'S NOTE

Amazon Web Services (AWS) helps commercial and government customers build and design airplanes and satellites, conduct space and launch operations, manage supply chains, and reimagine space exploration. We have helped build solutions for aviation organizations to design and operate aircraft and equip space organizations to process and transform space data to make it actionable and accessible to customers around the globe. With the AWS Cloud, customers are accelerating their missions, removing barriers to innovation on Earth, and inspiring future generations.

The aerospace and space industries have changed considerably since the early days of jet travel and the Apollo missions. New technology is making space and sky more accessible than ever—startups are reaching skyward and towards the stars with seemingly every day new, innovative ideas.

This month's issue of Architecture Monthly brings you curated content to highlight some of these ideas and innovations from our commercial aerospace, aerospace defense, and space leaders, including the newly launched AWS for Aerospace and Satellite organization. We highlight new and updated technology, practical solutions and their applications, and innovative companies that are working towards exploring and learning about the final frontier. Get inspired!

We'd like to thank our experts, Scott Eberhardt, Worldwide Tech Leader, Aerospace; Shayn Hawthorne, Sr. Mgr., Aerospace Tech Leader; and Buffy Wajvoda, Head of SA - Space & Satellite for their contributions.

Please give us your feedback! Include your comments on the Amazon Kindle page. You can view past issues and reach out to aws-architecture-monthly@amazon.com anytime with your questions and comments.

Bonnie McClure, Managing Editor

Cover: Astra's (astra.com) rocket leaves the launch pad in Kodiak, Alaska. Their affordable and responsive launch services deliver to low Earth orbit as part of their mission to improve life on Earth from space.

TABLE OF CONTENTS

NOTICES

ASK AN EXPERT

Scott Eberhardt: Worldwide Tech Lead for Aerospace at AWS
Buffy Wajvoda: Worldwide Leader for Aerospace and Satellite Solutions Architecture at AWS
Shayn Hawthorne: Space Technology Leader for AWS Aerospace and Satellite Solutions Division

What are the general architecture pattern trends for aerospace in the cloud?

Buffy: The aerospace industry uses AWS to modernize its IT infrastructure, enable digital transformation, and drive innovation. However, they also have some unique challenges that make it an exciting segment to support.

The space and satellite industry, by definition, is focused beyond the confines of Earth. As such, customers in the industry require a global network to communicate back to Earth regardless where their mission takes them. Take for instance Capella Space (Capella). They aim to provide the most frequent, timely, and high-quality synthetic-aperture radar (SAR) imagery products and make them accessible through an intuitive, self-serve online platform.

AWS allows Capella to reach this goal. By using services like AWS Ground Station and our vast 100 gigabytes per second (Gbps) global infrastructure, Capella can downlink their SAR imagery data from all six inhabitable continents and transfer that data back to their main virtual private cloud (VPC).

The aerospace industry also depends on significant computational resources to perform calculations such as computational fluid dynamics, structural stress testing, and Monte Carlo-type simulations. AWS high performance computing (HPC) services allow aerospace customers to complete these calculations in record time.

In fact, Descartes Labs uses AWS HPC to process and understand the world and to handle the flood of data that comes from sensors on the ground, in the water, and in space. In 2021, they achieved a top 40 position for fastest HPC run in the world using AWS HPC and Amazon Elastic Compute Cloud (Amazon EC2) On-Demand Instances. They achieved 9.95 petaflops (PFLOPs) of performance with a peak of 15.11 PFLOPs.

Additionally, Maxar Technologies (Maxar) uses AWS HPC to reduce weather processing by 58%.

> "Prior to using AWS, no one thought any cloud environment was capable of outperforming an on-premises supercomputer in generating numerical weather predictions," says Stefan Cecelski, a data scientist at Maxar. "But with the fast networking speed provided by AWS, we accomplished what many IT experts considered impossible."

AWS storage services are also a differentiator for the aerospace industry. Aerospace customers produce petabytes (PB) of data. This includes data from Earth observation pictures and readings, calculations, engineering artifacts, factory operations, and mission analysis. In addition, the aerospace industry is highly regulated, which creates requirements to keep data for long periods of time. AWS intelligence tiered storage options like Amazon Simple Storage Service (Amazon S3) and Amazon S3 Glacier provide storage with 99.999999999% durability and provide a cost-effective storage method to meet storage and regulatory needs.

Additionally, many aerospace customers are required to comply with International Traffic in Arms Regulations (ITAR) export controls. AWS GovCloud (US) makes this compliance possible and relatively easy.

> "The AWS GovCloud (US) region meets our customer's requirements regarding ITAR data," says Jeff Wright, cloud service senior manager in Lockheed Martin's Enterprise Operations group. "Knowing that it's a separate region that has already gone through accreditations and LM security approvals, coupled with our blueprints on top, provides a lot of confidence."

When putting together an AWS architecture to solve business problems specifically for aerospace customers, what are some of the considerations?

Scott/Buffy: The number one concern we face with aerospace customers is data security. Even commercial companies have ITAR concerns that influence their thinking. In the US, building in AWS GovCloud (US)

satisfies most customers' needs. Outside the US, there is no AWS GovCloud (US) equivalent, so specialized environments must be built. Partners and third parties supply help, such as Thales Solutions for AWS.

The second concern aerospace customers have is around the regulatory environment in which aerospace vehicles must operate. Currently, AWS services, such as AWS IoT Greengrass, AWS Snowball, and AWS Snowcone, operate on aircraft.

AWS GovCloud (US)'s security and compliance features meet many regulatory requirements. For example, architectures can be built on AWS GovCloud (US) that comply with Federal Risk and Authorization Management Program (FedRAMP)) high baseline; ITAR; Export Administration Regulations (EAR); Department of Defense (DoD) Cloud Computing Security Requirements Guide (SRG) for Impact Levels 2, 4, and 5; Federal Information Processing Standards (FIPS) 140-2; IRS-1075; and other compliance regimes.

Aerospace customers use HPC. Most are using commercial off-the-shelf codes that are compiled for specific chip sets. Usually, those chip sets are Intel or AMD based. Another cost-effective option is AWS Graviton2, which powers chips that have been extensively benchmarked for typical aerospace HPC workloads, such as computational fluid dynamics and finite-element analysis. Another recommendation for aerospace HPC workloads is to use Amazon FSx for Lustre or Amazon FSx for Windows File Server.

Building HPC clusters for aerospace customers depends on their size and tolerance for customization. We generally recommend AWS ParallelCluster for smaller customers who need a managed service with available support.

Do you see different trends in aerospace in cloud vs. on-premises?

Scott/Shayn: By using AWS Cloud, aerospace startup companies can reduce their capital costs and do not have to worry about provisioning or waiting to have their IT infrastructure built. They can also see how to scale as their company grows. Most startups use HPC and computer-aided engineering workloads using Amazon S3, AWS Parallel Cluster with Amazon AppStream 2.0, or Amazon WorkSpaces. These are managed services that allow engineers to get what they need, when they need it, without having to worry about provisioning for future workloads. Some startups start their cloud journey using software as a service (SaaS) partners. However, the added cost of a fully integrated SaaS provider can lead them to seek their own solutions on AWS.

Many legacy companies have built supporting software to manage their workflows on their existing infrastructure. They tend to mix and match software from various vendors and use internal solutions to manage data. Because of this, when legacy customers discuss migrating to the cloud, integration is often a significant concern. Understandably, there are some times they do not want to rebuild their workload tools and services for a new platform. So, there is growing interest in containers, specifically Kubernetes. We introduce them to integration services and our Amazon Elastic Kubernetes Service (EKS) to manage their container stacks.

Additionally, some aerospace companies are looking at using data lakes on Amazon S3 so that their data scientists can mine larger stores of data to gather insights they were unable to see before. Use of artificial intelligence (AI)/machine learning (ML) using Amazon SageMaker is growing in aerospace. Data scientists and engineers are increasingly building, training, and deploying ML models as they outgrow their on-premises services. Many companies are also exploring microservices, which simplifies and standardizes coding and makes code more adaptable and scalable.

What's your outlook for aerospace, and what role will cloud play in future development efforts?

Shayn: The outlook for aerospace, and specifically space system development, is enormously positive. The cloud allows companies to reduce their dependency on building undifferentiated infrastructure and focus on generating and analyzing space data and developing space technology.

The space industry is changing at the most rapid rate since the original Apollo missions. Advancements in technology such as propulsion, materials, and the emergence of CubeSats have significantly lowered barriers to entry. Once the sole province of government agencies, space is now accessible to global consortiums, venture capitalists, and university students and researchers.

For example, in 2018, there were 365 satellites launched globally. By 2025, we estimate that number to triple to over 1,000 launches per year. Given these changes and the fact that there are now commercial alternatives, governments no longer have to design, acquire, operate, and sustain all space systems. They can use commercial services to fulfill key missions, including intelligence, communications, military applications, and space exploration.

The AWS Cloud is well suited to space endeavors, which require global networks,

significant computational resources, and the ability to cost effectively store petabytes of data. And there are certainly significant investments going into space technology. According to a Bryce Space and Technology 2018 market analysis, the global space economy was estimated at $360 billion in 2018. A Space Foundation study revised that number in 2019 to $424 billion – huge growth! Through 2024, the US National Aeronautics and Space Administration (NASA) alone is projected to spend between $24-30 billion on Project Artemis to take humankind back to the moon. Goldman Sachs reports the global space economy will exceed $1 trillion by 2040. This is great for AWS space customers and partners because industry analysis of IT spending trends by Deloitte, with a 44 percent year-over-year IT growth rate, forecasts the space cloud spend should reach $3.7 billion by 2023, which shows that the analyst community agrees that the cloud will play an important role in space development.

And what a great development area to be a part of! The world is experiencing a resurgence in space exploration and development that benefits all of humankind. NASA's Project Artemis will land the first woman and next man on the moon in the 2020s, then develop a human encampment on the moon, and follow up with human missions to Mars. The European Space Agency (ESA) and newly space-faring countries like the United Arab Emirates are all conducting missions to the moon and Mars. And many new space constellation operators, including Amazon's Project Kuiper, plan to build and launch thousands of satellites into low Earth orbit to offer global communications services to the entire world. AWS is excited to work with our customers so they can use AWS Cloud services to realize their dreams – to the stars, through the cloud!

Anything else you'd like to add?

Scott: Two areas that are hot with customers are model-based systems engineering (MBSE) and connected aircraft and spacecraft. The former is the subject of the recently released Model Based Systems Engineering (MBSE) on AWS: From Migration to Innovation [whitepaper](#). This whitepaper includes a generic architecture for MBSE.

ABOUT THE EXPERTS

Dr. Scott Eberhardt is the Worldwide Tech. Lead for Aerospace at AWS. He started this role in January 2020 after serving as the EMEA HPC Specialist. In addition, Scott holds a Visiting Reader position at Imperial College.

Buffy Wajvoda is the Worldwide Leader for Aerospace and Satellite Solutions Architecture at AWS. She was a founding leader of the organization in 2020 after serving at the AWS Ground Station Solutions Architecture leader.

Shayn Hawthorne works in the AWS Aerospace and Satellite Solutions Division as the Space Technology Leader. He works with customers to develop services and features that disrupt how space exploration, satellite, and launch operations are conducted and to collect and use space data to improve their missions and products. Shayn is the founder of the AWS Ground Station service.

INTRODUCTION TO AWS FOR AEROSPACE AND SATELLITE

Transforming the future of space alongside our customers

Get inspired. Follow pioneer astronaut Peggy Whitson as she gains first-hand insights from Major Gen. Clint Crosier (Ret.) on how Amazon Web Services (AWS) helps space customers take research and discovery to the next level.

AWS FOR AEROSPACE AND SATELLITE

Explore how NASA's Mars Mission, Maxar Technologies, Fireball International, and Capella Space use AWS to help astronauts, scientists, and everyday heroes do their jobs better. See how AWS is innovating for customers to make the future of space a reality.

Satellite Ground Segment: Moving to the Cloud

This report from Northern Sky Research (NSR) on cloud computing in the satellite world named AWS Ground Station as one of the solutions available to satellite operators looking to run a fully managed cloud-based ground station service. The report highlights the measurability and flexibility that cloud-based solutions provide to the traditionally Capital expenditure (CapEx) heavy investments of satellite ground infrastructure. AWS is touted as offering flexible per-minute access to antennas for self-service scheduling, alleviating the need to buy, lease, build, or manage a fully owned ground segment.

Download the report to learn more

View online

BLOG

## Capella uses space to bring you closer to Earth

by AWS Public Sector Blog Team

[Capella Space](#), a provider of on-demand Earth observation data via synthetic aperture radar (SAR), is going all-in on AWS. Capella runs its entire IT infrastructure on AWS to automate and scale its operations. [AWS Ground Station](#) makes it simple and cost effective for Capella to command and control its constellation and receive its satellite data directly into AWS using a fully managed network of antenna systems located around the world.

"Running our ground infrastructure on AWS allows Capella and its customers to automate and scale operations across the globe in order to minimize latency and reactivity. We have optimized our operations to take advantage of improvements and innovations in cloud computing and this partnership between AWS and Capella lays the foundation for a tight interface between the two companies and their capabilities," said Payam Banazadeh, CEO and Founder of Capella Space.

Earth observation using satellite images has the power to solve some of the world's biggest problems by monitoring change. For example, government agencies can use Earth observation data to get the complete picture when a natural disaster such as a hurricane or flood strikes, and use that information to hasten recovery efforts. Environmental organizations can use the data to monitor changes on the planet's surface, highlight the impact on wildlife, and inform efforts to protect endangered species. Scientists can measure soil moisture to better forecast crop yields and prevent famine.

Near-immediate access to satellite data is critical in these scenarios, yet most traditional ground stations take up to 24 hours to complete tasking, schedule satellite passes, establish uplinks to the satellites, download and process data, and deliver the imagery. Since most traditional satellites are launched
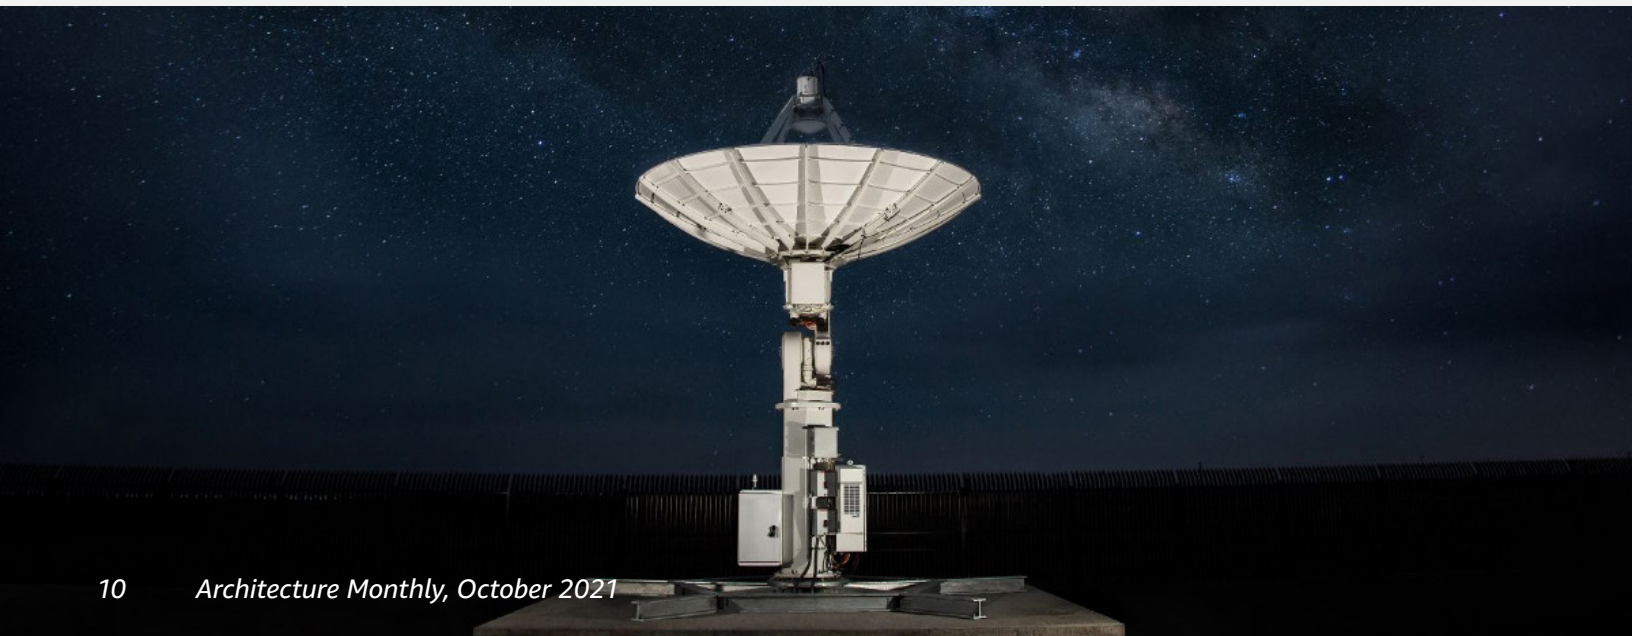
into specific orbits, they are restricted to observing the Earth for only a few hours a day, so they can't accurately observe changes in conditions that require continual monitoring, such as crop health and weather forecasting. They also are stymied by weather issues and darkness, and can't see 75 percent of the Earth's surface. Finally, in order to download satellite data when and where they need it, customers have to invest in multiple geo-located ground stations, which are costly to operate and maintain – or wait for the satellite to pass over a desired location.

Capella's SAR satellites can see through clouds and darkness to collect millimeter-scale resolution imagery 24×7 in all weather conditions, delivering a persistent monitoring capability for customers working in fields where operating conditions change quickly, such as defense and humanitarian assistance. "The ability to use satellite imagery to see virtually anywhere in the world in high resolution and to do so quickly and in rapid succession has never really existed before. Now, you can request satellite imagery easier than ever before and near-real time," said Banazadeh.

Using AWS, Capella provides its customers with access to satellite data within minutes of its capture – far faster than traditional satellite data delivery services. By leveraging and communicating directly with AWS Ground Station antennas located around the world, Capella's Earth observation solution is able to deliver satellite data to customers more cost-effectively than ever before to help customers in agriculture, infrastructure, defense, and disaster response immediately analyze and extract value from their data. "The other piece of the puzzle that AWS is really helping us with is ensuring the flow of data is happening really fast, and to make sure the reactivity is there for rapid delivery," said Banazadeh.

How it works

Capella receives requests from customers for satellite imagery over specific areas of interest (AOIs) via web application or API. Imaging tasks are instantly routed to Capella satellites via a geosynchronous communications relay, and data downlink contacts are scheduled with AWS Ground Station using a scheduling API. Capella satellites orbit over the AOIs and position themselves for the tasks, which range between small millimeter-scale change detection to large area collection. As Capella satellites complete their imaging tasks and continue their orbit, the satellites downlink the SAR data into the AWS Cloud, where it gets processed into high-quality satellite imagery for rapid delivery.

By using AWS, Capella has improved the recency of satellite data by downlinking their raw satellite data to AWS Cloud through AWS Ground Station and making that data readily available to their customers within minutes. As a fully managed service, AWS Ground Station lets Capella Space teams control satellite communications, process data, and scale its operations without having to worry about building or managing its own ground infrastructure. By entrusting its ground station infrastructure to AWS, Capella can focus its staff on innovating with new applications that process, analyze, and distribute satellite data. "At Capella Space, we believe the space-based technologies we're building and our collaboration with AWS will improve security, commerce, conservation, and the well-being of all people through the 21st century and beyond," said Banazadeh.

Learn more about how you can build your aerospace and satellite workflows with AWS. And listen to episode #29 of the Fix This podcast, "To the stars, through the cloud."

Read blog post online

## Run Machine Learning Algorithms with Satellite Data

Use AWS Ground Station to ingest satellite imagery, and use Amazon SageMaker to label image data, train a machine learning model, and deploy inferences to customer applications.
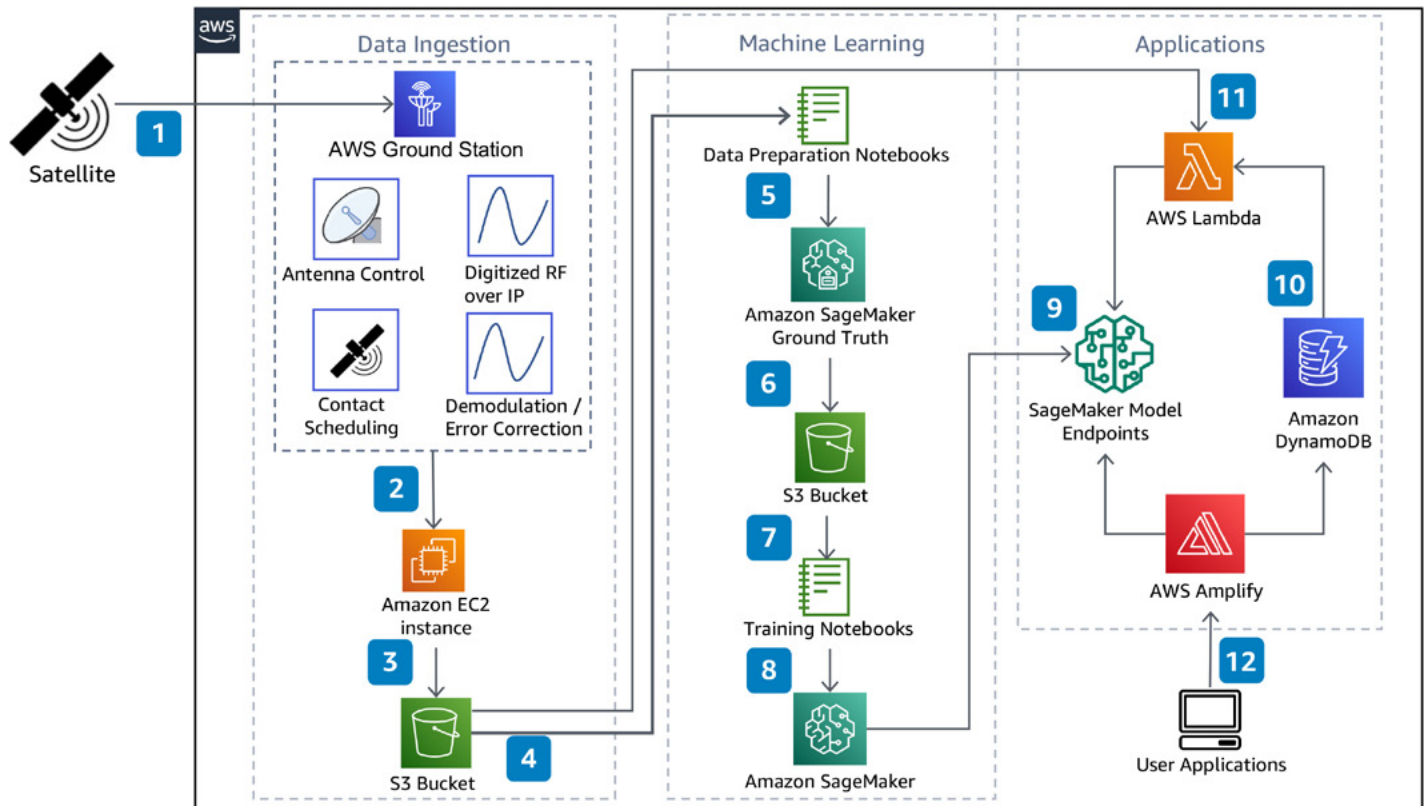
1. Satellite sends data and imagery to the **AWS Ground Station** antenna.
2. **AWS Ground Station** delivers baseband or digitized RF-over-IP data to an **Amazon EC2 instance**.
3. The **Amazon EC2** instance receives and processes the data, and then stores the data in an Amazon S3 bucket.
4. A Jupyter Notebook ingests data from the **Amazon S3** bucket to prepare the data for training.
5. **Amazon SageMaker Ground Truth** labels the images.
6. The labeled images are stored in the **Amazon S3** bucket.
7. The Jupyter Notebook hosts the training algorithm and code.
8. **Amazon SageMaker** runs the training algorithm on the data and trains the machine learning (ML) model.
9. **Amazon SageMaker** deploys the ML models to an endpoint.

10. The SageMaker ML model processes image data and stores the generated inferences and metadata in **Amazon DynamoDB**.
11. Image data received into **Amazon S3** automatically triggers an **AWS Lambda** function to run machine learning services on the image data.
12. Applications interact with **AWS Amplify** to access the ML algorithm and database.

View reference architecture online
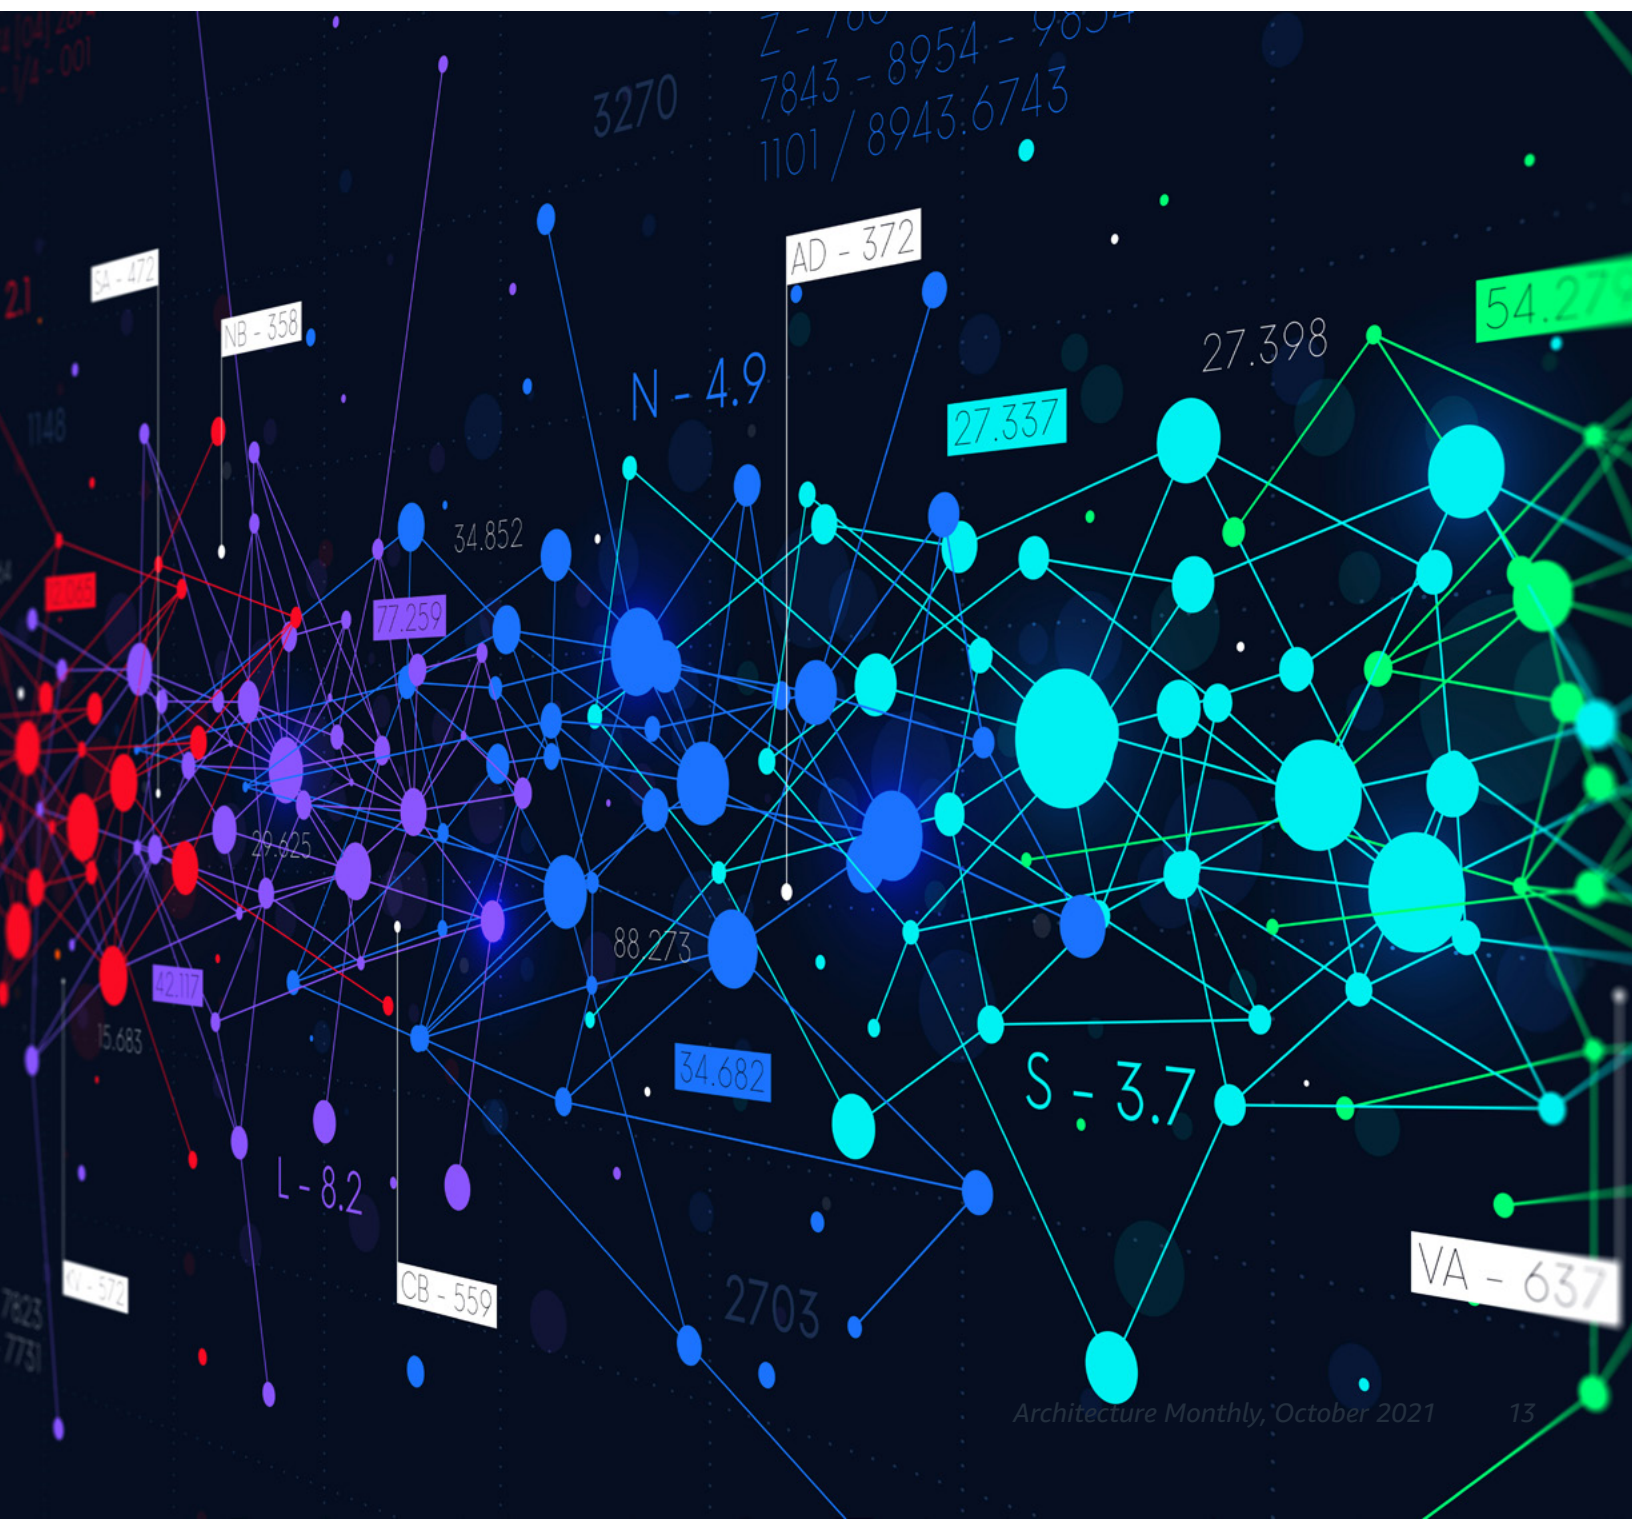
BLOG

UAE Mars mission uses AWS to advance scientific discoveries

by AWS Public Sector Blog Team

On February 9, a new object successfully began to orbit Mars: an uncrewed spacecraft called the Hope Probe. The mission has already returned the first image of Mars, taken by Hope's Emirates eXploration Imager from an altitude of 24,700 km.

Led by the [Mohammed Bin Rashid Space Centre (MBRSC)](), the Hope Probe is the first interplanetary mission for the United Arab Emirates (UAE), the fifth country in history to reach the red planet. It will also be the first spacecraft to capture a complete picture of the Martian atmosphere and its layers during different times of the day and different seasons for one complete Martian year. Once data transmitted by the Hope Probe reaches the scientific teams on Earth, MBRSC will use Amazon Web Services (AWS) advanced technologies to process and analyze the vast amounts of data and imagery to help researchers better understand the Martian atmosphere and its layers.

*First image of Mars returned by the Hope Probe (source: [Twitter]()).*

[Following a successful Mars Orbit Insertion (MOI)](), the next stage in the Hope Probe's journey is the transition from the Capture Orbit to an acceptable Science Orbit in preparation for its primary science operations.

Over the next two years, the probe's primary scientific mission is to study daily and seasonal weather cycles, weather events in the lower atmosphere such as dust storms, and variations in weather in different regions of Mars. It will help scientists attempt to answer questions about why Mars' atmosphere is losing hydrogen and oxygen into space and the reason behind Mars' drastic climate changes.

Scientific instruments mounted on one side of the spacecraft will collect information about atmospheric circulation and capture aerial images of Mars using the visible, thermal infrared and ultraviolet wavelengths. The three instruments are:

- **EXI (Emirates eXploration Imager)** will capture high resolution images of Mars, measuring water ice and ozone in the lower atmosphere and capture colorful images of Mars.
- **EMUS (Emirates Mars Ultraviolet Spectrometer)** will measure oxygen and carbon monoxide in the thermosphere as well as hydrogen and oxygen in the upper atmosphere, which are essential for determining the loss of water from the upper atmosphere.
- **EMIRS (Emirates Mars Infrared Spectrometer)** will measure both surface and atmospheric temperatures, as well as global distribution of water ice, water vapor and dust in lower atmosphere.

MBRSC will use a customized suite of AWS services to host the mission science data center that is responsible for processing the large volumes of data collected by these instruments on the Hope Probe and to convert raw instruments' products into higher level scientific products that will help scientists to perform their analysis and research.
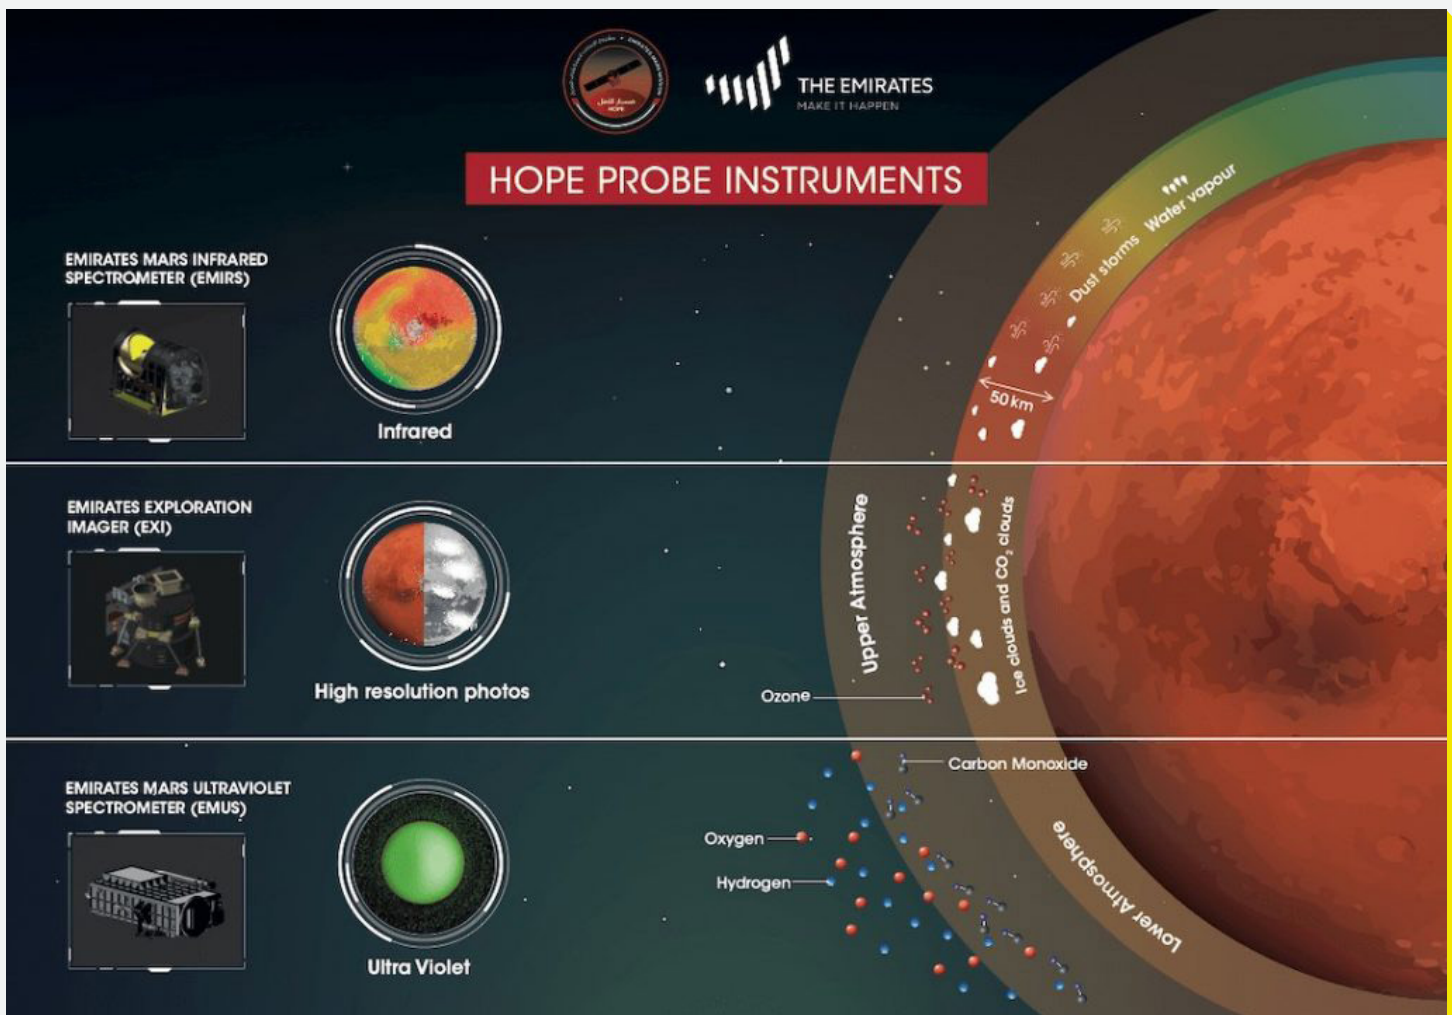
Powering the data journey: From Mars to the scientific community in 20 minutes

What makes this mission different from others is that it is fully automated. This means that the multi-stage process of receiving, validating and indexing raw instrument data, and processing it to generate higher level of scientific products can all be accomplished in less than 20 minutes, compared to 48 hours for other missions not being operated on the cloud. Within this short time, the mission science team would have immediate access to the processed data and perform their scientific analysis.

Omran Al-Hammadi, senior manager, Science Data Section, MBRSC, said, "The Emirates Mars Mission 'Hope Probe' will be the first probe to provide a complete picture of the Martian atmosphere and its layers. It will help answer key questions about the global Martian atmosphere and the loss of hydrogen and oxygen gases into space over the span of one Martian year. The mission science data center is completely deployed on AWS and utilizes different available managed services and infrastructure to fully automate the mission's data management and processing."

"It's exciting to see MBRSC pushing the boundaries of cloud applications for deep space exploration," said Clint Crosier, director, AWS Aerospace and Satellites. "Our team is proud to provide mission-critical support for delivering the vast volumes of data and intelligence from the Hope Probe to the scientific community."

Cloud technology is powering this complex data management, processing, and data dissemination. Scientific instrument data is downlinked from the spacecraft to the Deep Space Network (DSN), and delivered to the Mission Operations Center (MOC). The MOC delivers all science data and any other information required for science data processing to the mission Science Data Center hosted on AWS, where it is indexed, stored and processed for the duration of the mission.

The mission Science Data Center will use AWS Step Function to orchestrate various AWS Lambda functions to check, index, and move files into a master database, and Amazon Simple Storage Service (Amazon S3) for shared storage. Then, data from each of the probe's instruments will be dynamically processed using AWS Batch which utilizes containerized processing software stored on AWS Elastic Container Registry (ECR).

Finally, the data is disseminated to the Emirates Mars Mission Science team and the scientific community via an Amazon CloudFront powered website. The website is hosted in an Amazon S3 bucket for high availability and scalability.

Amazon Cognito powers the site's user authentication and authorization tool, which is important for secure access to the website. Registered researchers will get their own unique URL so there is no unauthorized access to data and no accessibility issues when the data is first made available.

While the orbiter's scientific mission will last for the next two years, the mission intends to provide long-lasting benefits including important insights for improving life on Earth and inspiring the next generation of explorers and innovators.
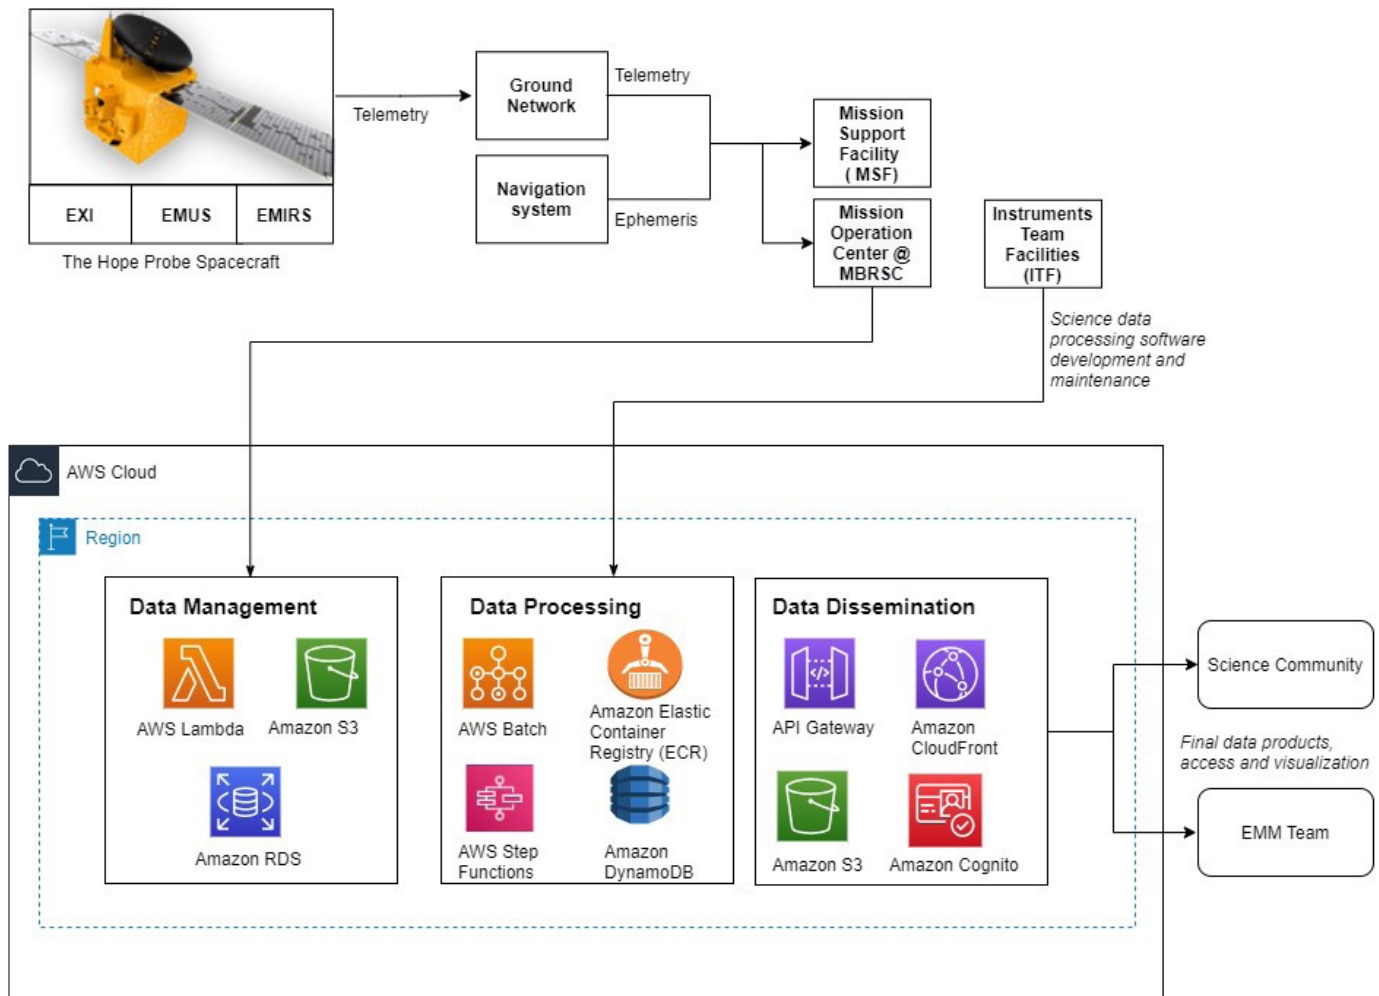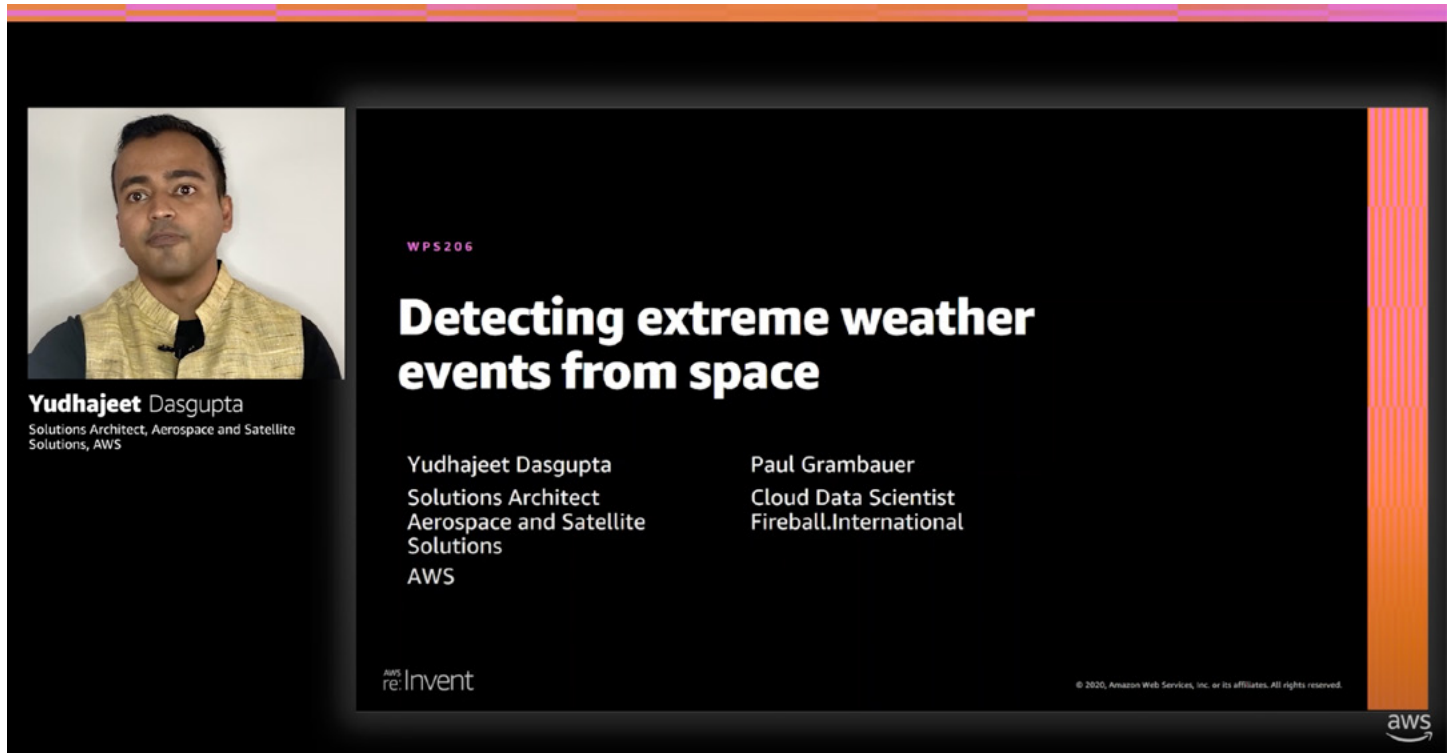
The Hope Probe Spacecraft architecture

Learn more about the Hope Probe mission on the Emirates Mars Mission website. AWS helps commercial and government customers build satellites, conduct space and launch operations, and reimagine space exploration. Read more about how the AWS Aerospace and Satellites team is enabling successful space missions and contact us.

Read blog post online

SOLUTION

AWS re:Invent 2020: Detecting extreme weather events from space

Satellites revolving around the earth collect petabytes of data every day and send them back for ingestion, categorization, processing, and dissemination. While more organizations dedicate resources to environmental monitoring and predictive analytics, it is essential to not just give them the platform to implement their technologies but also share the responsibility in solving the problem. In this session, learn how to use AWS Ground Station, Amazon SageMaker, and data lakes to automate, scale extreme weather event detection, and help build disaster resilience. You also hear how Fireball International delivers early wildfire detection services through multi-sensor rapid wildfire intelligence.

[Watch video online](#)

BLOG

Announcing the AWS Space Accelerator for startups

by Sandy Carter

Helping startups succeed, from inception to IPO, has been an integral part of Amazon Web Services (AWS) since our beginning. Today, more startups build on AWS than any other cloud provider, and many of our customers and partners started with AWS at an early stage. For example, Capella Space began their journey on AWS as a startup to provide customers with access to satellite data within minutes of capture—far faster than traditional satellite data services, and at a lower cost.

As the global aerospace and satellite industry enters a new age, leveraging the power of startups will drive innovation and improve accessibility to space data. AWS understands the importance of creating and nurturing

startup communities, and today, we are proud to announce the launch of the AWS Space Accelerator.

AWS Space Accelerator

The AWS Space Accelerator is a four-week business support program that is open to space startups seeking to use AWS to help solve the biggest challenges in the space industry. Applications are open today and proposals are due by April 21, 2021.

The AWS Space Accelerator will provide technical, business, and mentoring resources to space startups around the globe. AWS is offering this opportunity in collaboration with Seraphim, one of the world's leading

investment groups focused exclusively on the space industry, who will provide business development and investment guidance. AWS and Seraphim will select a cohort of space startups to participate in an intensive, four-week program with AWS Cloud and technical training to help them accelerate research, development, and growth using AWS. Seraphim, has helped dozens of emerging space companies around the world become investment ready through its own space accelerator program.

"Startups provide a catalyst for bold new experimentation in the space industry," said Clint Crosier, director of aerospace and satellite at AWS. "We are proud to announce the AWS Space Accelerator as part of our ongoing commitment to help startups succeed, and to shape the future of aerospace. We look forward to helping the first cohort of companies launch and grow through this new program."

AWS and Seraphim are accepting applications from innovative startups at all stages of maturity working with space technology or space-derived data who have a clearly defined and unique mission. Qualifying missions can include, but are not limited to, earth observation, electronics and robotics, spacecraft launch and delivery, spacecraft hardware and software, launch manufacturing and launch operations, and more.

"AWS is building a world-class accelerator program for space startups tailored around the use of AWS services, and we are excited to work with them," said Rob Desborough, partner, Seraphim. "Combined with Seraphim's unique domain expertise in space investment and acceleration, we hope to attract, select, and accelerate startups who will go on to address some of the world's biggest challenges."

Applications will be judged on several factors, including: the innovative and unique nature of the project, the overall value the solution will bring to the industry, the creative application of AWS to solve problems, and the team's ability to deliver on an identified opportunity.

AWS and Seraphim will select 10 companies to participate in the four-week accelerator taking place in June 2021. Selected startups may receive up to $100,000 in AWS Activate credit, as well as mentoring from space domain and technical subject matter experts with deep experience working on AWS. The program offers collaboration opportunities with AWS customers and members of the AWS Partner Network (APN) looking for cutting-edge technology solutions to their most challenging space problems. It will also bring participants together with space-savvy venture investors for fundraising conversations.

Learn more about the AWS Space Accelerator and how to apply. Learn more about how AWS Aerospace and Satellite supports customers and partners. Learn more about how Seraphim is helping emerging space companies become investment ready.
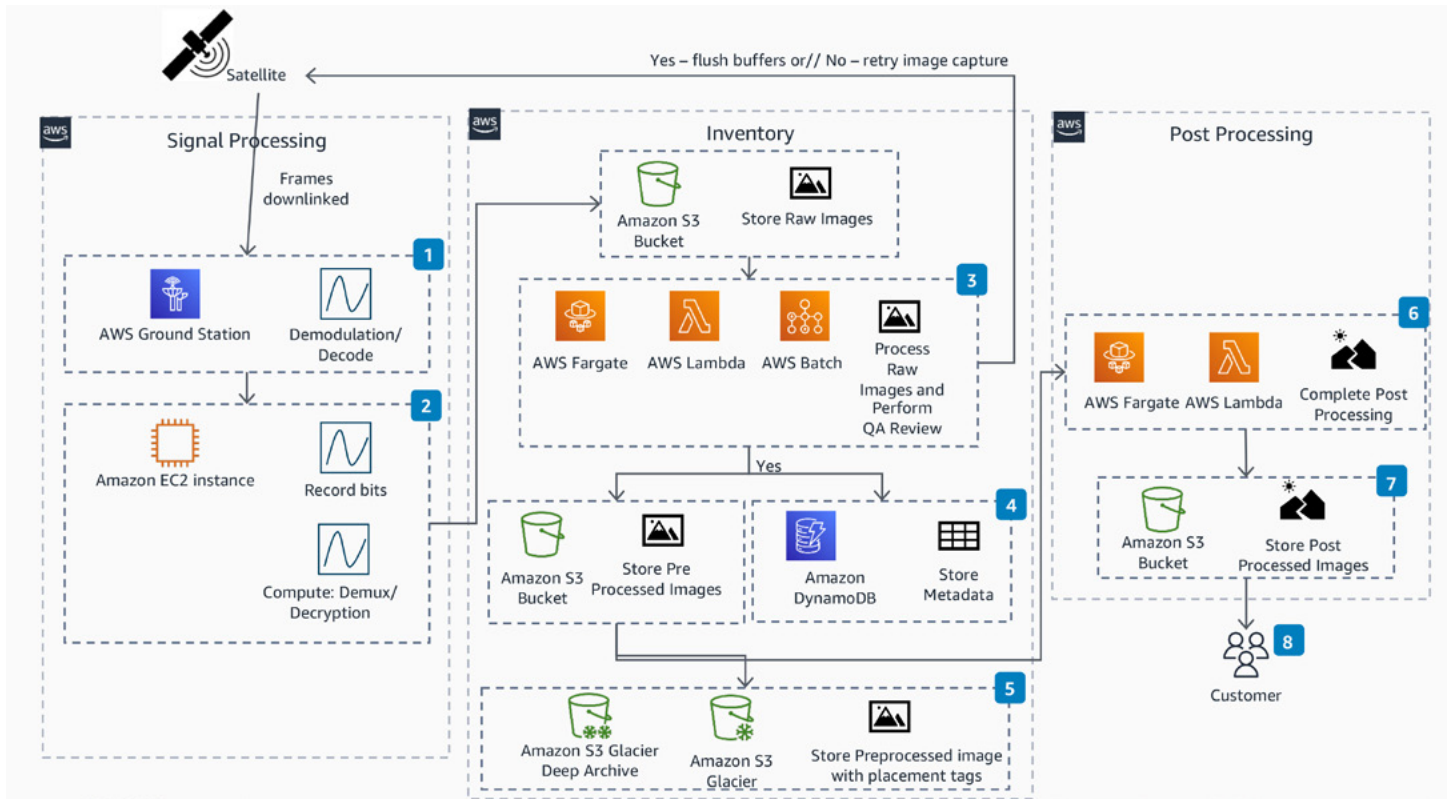
Subscribe to the AWS Public Sector Blog newsletter to get the latest in AWS tools, solutions, and innovations from the public sector delivered to your inbox, or contact us.

Read blog post online

REFERENCE ARCHITECTURE

Electro-Optical Imagery Reference Architecture

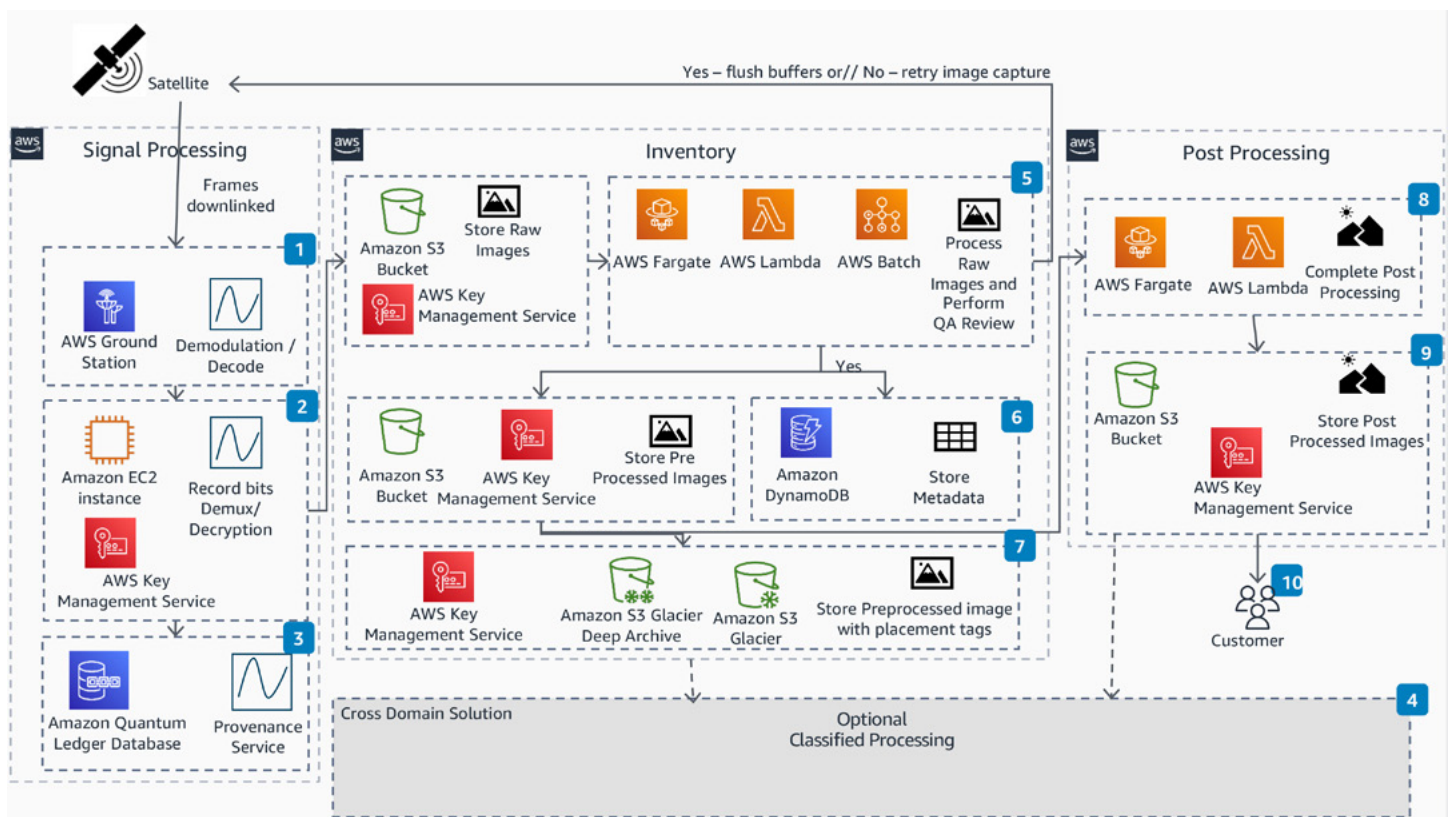Process electro-optical imagery on AWS

1. Demodulate and Decode: Extract baseband waveform from modulated carrier; remove forward error correction
2. Convert into raw sensor data: Decommutate signal frames; decrypt data
3. Process Raw Images: Process Raw Images and Perform QA Review

- QA Review: Confirm Images are sufficient for processing
- **AWS Batch**: Run multiple jobs in parallel
- **AWS Fargate** and **AWS Lambda**:

  - Sensor Correction: Apply corrections for optical distortions
  - Orthorectify: Sensor perspective
  - Georeference: Apply image to spatial grid and assign known coordinate system
  - Generate Thumbnails: Create postprocessed thumbnails for customer purchase

4. Store metadata: Store information on latitude/longitude collection, region collection, time and date of retrieval

5. Storage: Store preprocessed images in a variety of **Amazon S3** services by balancing cost savings and time of retrieval.
6. Post Processing and Analysis: Complete imagery processing

   ◦ Feature Extraction: Identify features in images (e.g. ships)
   ◦ Naming/Tagging of Features: Tag features by name/identification system
   ◦ Time Series Creation: Tag images to sort by time

7. Storage and Dissemination: Final storage of images and analytics for end customer
8. Customer Delivery: Deliver final images to end customers

Classified Processing

Process classified electro-optical imagery on AWS

1. Demodulate and Decode: Extract baseband waveform from modulated carrier; remove forward error correction
2. Convert into raw sensor data: Decommutate signal frames; decrypt data
3. Immutable transaction log: Cryptographically establish provenance and fidelity
4. Optional Classified Processing: Throughout the image processing, move data to the appropriate regions for classified processing.
5. Process Raw Images: Process Raw Images and Perform QA Review

- QA Review: Confirm Images are sufficient for processing
- **AWS Batch**: Run multiple jobs in parallel
- **AWS Fargate** and **AWS Lambda**:

  - Sensor Correction: Apply corrections for optical distortions
  - Orthorectify: Sensor perspective
  - Georeference: Apply image to spatial grid and assign known coordinate system
  - Generate Thumbnails: Create postprocessed thumbnails for customer purchase

6. Store metadata: Store information on latitude/longitude collection, region collection, time and date of retrieval
7. Storage: Store preprocessed images in a variety of **Amazon S3** services by balancing cost savings and time of retrieval .
8. Post Processing and Analysis: Complete imagery processing

   - Feature Extraction: Identify features in images (e.g. ships)
   - Naming/Tagging of Features: Tag features by name/identification system
   - Time Series Creation: Tag images to sort by time

9. Storage and Dissemination: Final storage of images and analytics for end customer
10. Customer Delivery: Deliver final images to end customers

View reference architecture online

AWS Architecture Blog

Cloud architecture guidance and best practices

READ & SHARE
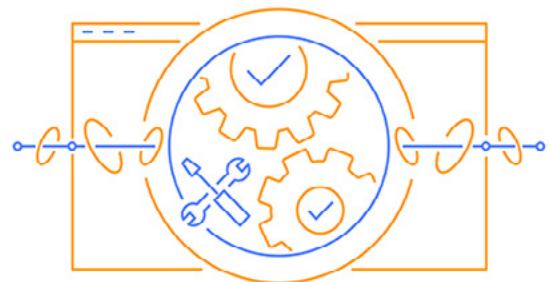
How to Build This

A video series designed for builders of all skill levels to start building with AWS

CASE STUDY

Avio Aero Accelerates Business Growth with HPC Solution on AWS

2019

[Avio Aero](#) is a GE Aviation business and the GE group's companywide center of excellence for mechanical transmissions and low-pressure turbines. It has more than a century's experience in overcoming technological challenges through continuous investment in research and development, strong relationships with leading universities and international research centers, and collaborative partnerships with key operators in the global aviation industry. During a recent period of strong growth, Avio Aero found that its on-premises high-performance computing (HPC) solution delayed response to customer requests.

Avio Aero uses an application called [TRAF](#) to run the computational fluid dynamics (CFD) simulations needed to develop the latest generation of aircraft turbines. Developed in collaboration with the University of Florence, the application is resource-intensive: a TRAF job for any single customer monopolized 100 percent of the company's 256 Intel cores for two to three months at a time.

"To eliminate this bottleneck and better manage the peaks requested, we needed to expand capacity to 800 cores," says Ivano Di Conca, global digital operations leader for Avio Aero. "But accomplishing this expansion in our data center would have cost $110,000 (€100,000) for capacity we don't need year-round. We wanted a solution that could scale up and down as needed."

Using the Latest Technology for Fast Response to Customer Requests

Avio Aero compared cloud providers and identified Amazon Web Services (AWS) as what Di Conca calls "the market leader for cloud HPC solutions." Now, when TRAF demands exceed the company's on-premises capacity, the HPC solution Avio Aero built on AWS spins up [Amazon Elastic Compute Cloud (Amazon EC2) C5 instances](#) to meet the need. Amazon EC2 C5 instances use second-generation Intel Xeon Scalable processors well suited for HPC workloads. Avio Aero deploys these AWS computing nodes in placement groups to reduce network latency and improve performance.

By running on AWS, the company avoids the problem of hardware lock-in inherent to on-premises deployments. "For TRAF workloads, we are always looking for more speed and higher performance," says Di Conca. "One of

"On AWS, we can immediately satisfy customer requests and continue our growth without worrying about capacity."

Ivano Di Conca

Global Digital Operations Leader, Avio Aero

the most important benefits of our AWS HPC solution is that we always have access to the latest Intel processors."

Flexible scalability and pay-as-you-go pricing enable faster results at lower costs. "TRAF scales linearly to about 800 cores, so—on AWS—when we need to reduce the elapsed time, we can accelerate problem resolution by a factor of three without spending anywhere close to what this capacity would cost on premises," says Di Conca.

The cloud HPC solution has transformed Avio Aero's relationship to its customers. "By using an HPC solution on AWS, we're no longer forced to leave customers waiting for results because of limited on-premises resources," says Di Conca. "On AWS, we can immediately satisfy customer requests and continue our growth without worrying about capacity."

To learn more, visit aws.amazon.com/hpc

About Avio Aero

Avio Aero designs, manufactures, and maintains propulsion systems for civil and military aviation. Founded in 1908, the company operates plants in Italy and Poland, and it partners with leading universities and international research centers to develop technological and manufacturing innovations.

About Intel

Intel is the world's leading designer and manufacturer of high-performance processors for servers, PCs, IoT devices, and mobile devices. AWS and Intel engineers have worked together for more than 10 years, building custom hardware to ensure AWS services run on a platform optimized for customer workloads for the best value. Intel Xeon Scalable processors power Amazon EC2 instances to help enterprises drive performance for their compute-intensive workloads.

Benefits of AWS

- Capacity no longer limits business growth
- Avoids on-premises hardware lock-in
- Provides access to latest processor technology
- Runtimes are 3 times faster at lower cost
- Capacity no longer slows response to customer requests

Read case study online

REFERENCE ARCHITECTURE

Connected Aircraft

Onboard flight data collection for fleet-wide analytics and predictive maintenance
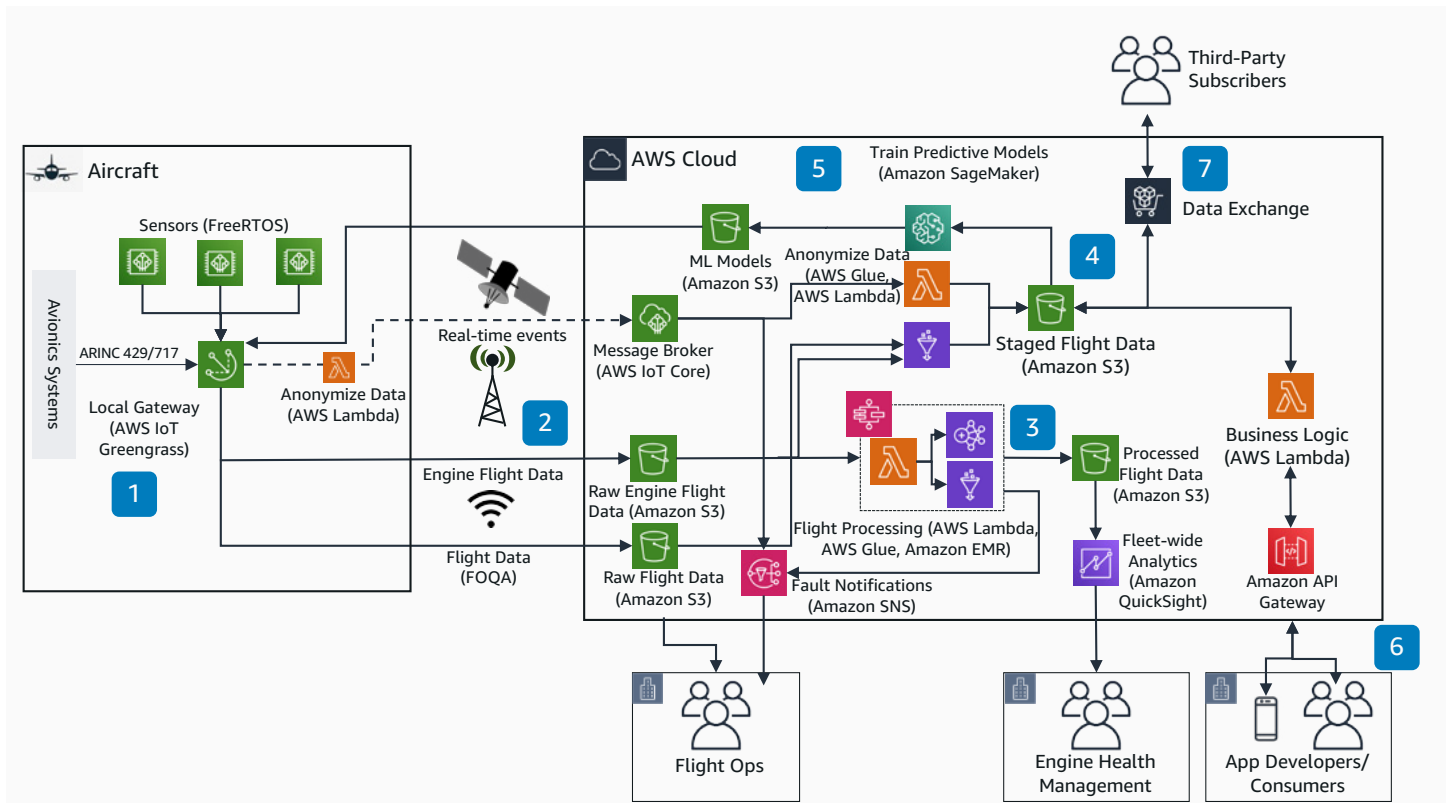
1. Flight (avionics) and telemetry (sensor) data is collected by **AWS IoT Greengrass** running on the flight-data acquisition unit on board.
2. Real-time events are anonymized and sent to Flight Operations on the ground.
   Flight data is offloaded to **Amazon S3** with aircraft at the gate and analyzed by Flight Operations for fuel burn optimization, fault analysis and other use cases.
3. Engine flight data is processed and analyzed with **AWS Glue** and **Amazon EMR** for engine health maintenance.
4. Anonymized flight and fault data is aggregated and stored in a **Amazon S3** data lake.
5. Models trained from aggregated flight and fault data are deployed to **AWS IoT Greengrass** on the aircraft for predictive maintenance
6. App developers build new digital solutions for the connected ecosystem.
7. Anonymized data is offered to third-party developers on a subscription basis with **AWS Data Exchange**.

[View reference architecture online](#)

CASE STUDY

Joby Aviation Uses AWS to Revolutionize Transportation

2021

Joby Aviation is an aviation startup focused on developing an electric vertical takeoff and landing (eVTOL) aircraft. The company's goal is to revolutionize personal mobility through quiet, fast, and clean air transportation. Joby publicly unveiled its five-seater aircraft in early 2020 and intends to begin operating the aircraft as a commercial passenger aircraft in 2024.

To study the flow of air over the aircraft, Joby engineers conduct Computational Fluid Dynamics (CFD) simulations. Joby has focused on using CFD more than traditional aircraft manufacturers might have because it allowed for rapid design iteration as the team explored various aircraft designs and architectures.

Using CFD provided engineers with a great deal of insight into the expected performance of the aircraft. Due to the novel nature of the design, the team would not have much useful data from past aircraft programs on which to base predictions, so the preemptive knowledge the team gleans from these CFD simulations is very valuable. "No one has designed an aircraft like this before. We are breaking new ground," says Alex Stoll, aeromechanics lead at Joby. "The aerodynamics are complex and we've therefore relied heavily on CFD to understand the various interactions across the airframe."

The Joby aircraft has six propellers and can take off in several different ways, each varying in energy efficiency, speed, and noise level. "For example, maybe we can make some adjustments to how the pilots fly the aircraft, helping to reduce our noise footprint, which is essential to delivering community acceptance," says Jeremy Bain, CFD lead at Joby. "CFD is the best way for us to analyze all the options at once and get a clear understanding of what's happening."

Joby engineers rely on high performance computing (HPC) to conduct thousands of complex, compute-intensive CFD simulations that use hundreds of CPU cores each and can take many hours to complete. The company originally ran Simcenter STAR-CCM+ CFD software on an on-premises compute cluster. But when Stoll met another engineer who was running CFD on Amazon Web Services (AWS), he realized that leveraging the cloud could

> *"Using AWS helps us get results from our CFD work faster. In some cases, we get results in 24 hours that may have previously taken a week."*
>
> **Alex Stoll**
>
> *Aeromechanics Lead, Joby Aviation*

improve efficiency, reduce costs, and boost the engineers' productivity by decreasing job turnaround time. "Using AWS helps us get results from our CFD work faster. In some cases, we get results in 24 hours that may have previously taken a week. To run the same analyses at the same scale and with the same turnaround time using on-premises resources, we'd have to have a huge compute cluster running at low utilization, which would be much more expensive to purchase as well as maintain."

Choosing AWS for Elasticity, Agility, Storage, and Security

Today, Joby uses Amazon Elastic Compute Cloud (Amazon EC2) to provide the capacity it needs to run CFD simulations and Amazon Simple Storage Service (Amazon S3) to store and protect its data. The company

also uses AWS ParallelCluster to simplify the configuration, deployment, and management of HPC environments where its simulations run while staying within budget.

Joby recently began using Amazon FSx for Lustre, a fully managed service that provides high-performance, scalable storage for compute workloads. "When we tried to run dozens of simulations at one time, we were reading and writing several gigabytes of data at a time, which slowed everything down," says Bain. "FSx for Lustre eliminated those capacity problems. We can increase the size of our hard drive easily now."

Joby is also using Amazon S3 Glacier to archive its CFD simulations. "We archive everything on Amazon S3 Glacier. It's an effective way to store 100 or 200 terabytes of data indefinitely if, by chance, we need that data in three or five years," says Bain.

Creating New Options in Transportation

By running CFD models using HPC and cloud-based managed infrastructure from AWS, Joby has been able to optimize the design of the aircraft while containing costs. "We can run all kinds of analyses quickly in the cloud, and then scale down when we're not using all that power," says Bain.

"When you're flying an aircraft, you want to have as much analysis under your belt as possible," adds Bain, "and with a goal of entering service in 2024, it's been important to us to work efficiently, which AWS has enabled us to do."

About Joby Aviation

Joby Aviation is a California-headquartered company developing a piloted, all-electric vertical takeoff and landing aircraft which it intends to operate as part of a fast, quiet, and convenient air taxi service beginning in 2024. Founded in 2009, Joby employs more than 700 people.

Benefits of AWS

- Provides the agility and elasticity to run thousands of CFD simulations each month
- Enables team to iterate quickly to design a safe, reliable, and quiet aircraft
- Allows the company to comply with export control regulations

Read case study online
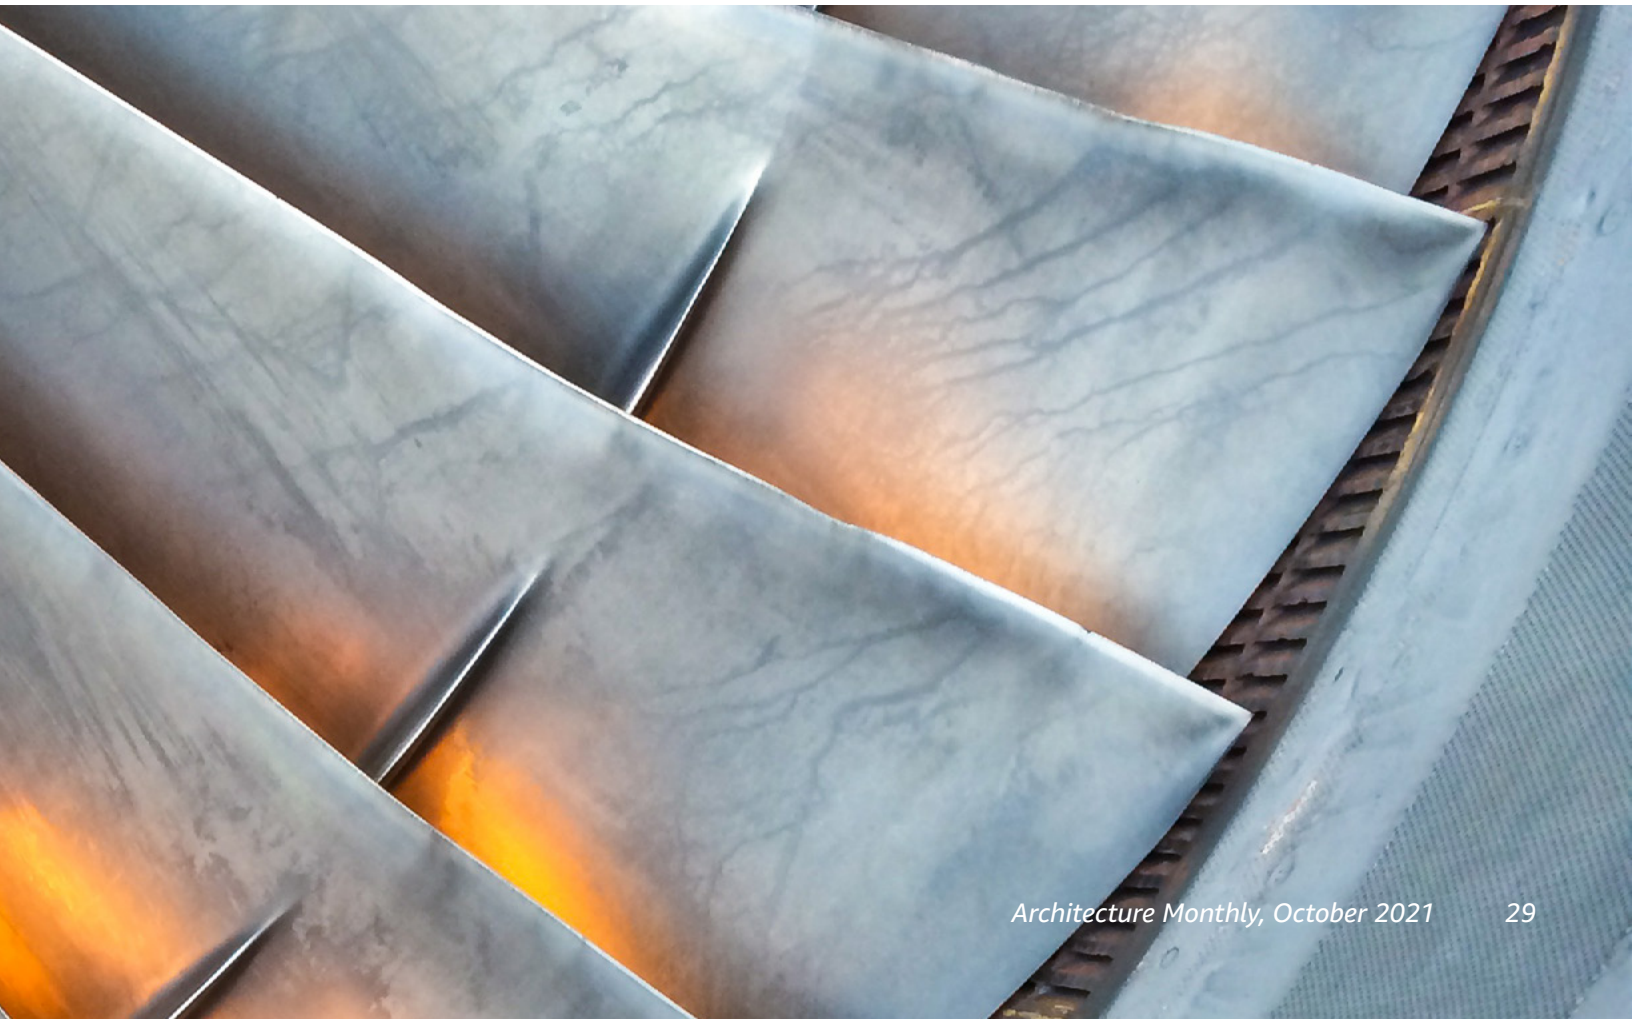
Model Based Systems Engineering (MBSE) on AWS: From Migration to Innovation

September 20, 2021

Abstract

Model Based Systems Engineering (MBSE) is a modern approach to the conventional practice of document-based systems engineering. MBSE benefits from modern cloud computing technologies, microservices, AI/ML, advanced analytics and others.

These technologies not only enable broad adoption of MBSE by engineering organizations, but also go beyond the current prospects of MBSE and bring innovation, flexibility, scalability and cost optimization. MBSE has been recently adopted by aerospace, energy, and automotive customers and growing in other industries where complex products - made of multitude of engineering disciplines and collaboration - are required to design, build, test, sustain and monitor the whole product lifecycle through their lifecycle.

AWS provides both building block technologies and solutions tailored to your needs. This whitepaper addresses both MBSE developers who develop MBSE technologies and MBSE users who use MBSE tools. It also provides introductory information about MBSE and its challenges for newcomers to this technology.

Introduction

## What is MBSE and why do industries start to use?

Today's multi-disciplinary, multi-supplier, and multi-application/tool product development environment encourages customers to adopt **agile practices** under firm deadlines. Agility is especially important in more traditional engineering environments such as **aerospace**, **energy**, and **automotive** industries, where *waterfall product development* has been long practiced and perfected.

As these products are highly complex, following a build-based agile iteration is expensive. Moreover, these industries are highly regulated and exceptionally qualitydriven. Hence, the agility should not bear "haste creates waste" but should bear technologies that allow data-driven decisions, visibility, traceability, simplicity, and preferably automation to "go fast" while "sticking to highest standards".

Complex product development requires inter-disciplinary and inter-supplier/company work. Hence, customers look for ways to be agile in this complex and dynamic environment. At the center of the complexity, we see **Systems Engineering (SE)**.

According to the [INCOSE (International Council on Systems Engineering)](), Systems Engineering is a transdisciplinary and integrative approach to enable the successful realization, use, and retirement of engineered systems, using systems principles and concepts, and scientific, technological, and management methods[1].

The term **system** here is a wide-range definition, including hardware, software, information, processes, people, configurations, supply chains, regulators and geographies. One of the fundamental functions of SE is to conduct the **V**erification and **V**alidation (**V&V**) of the product development and overall lifecycle.
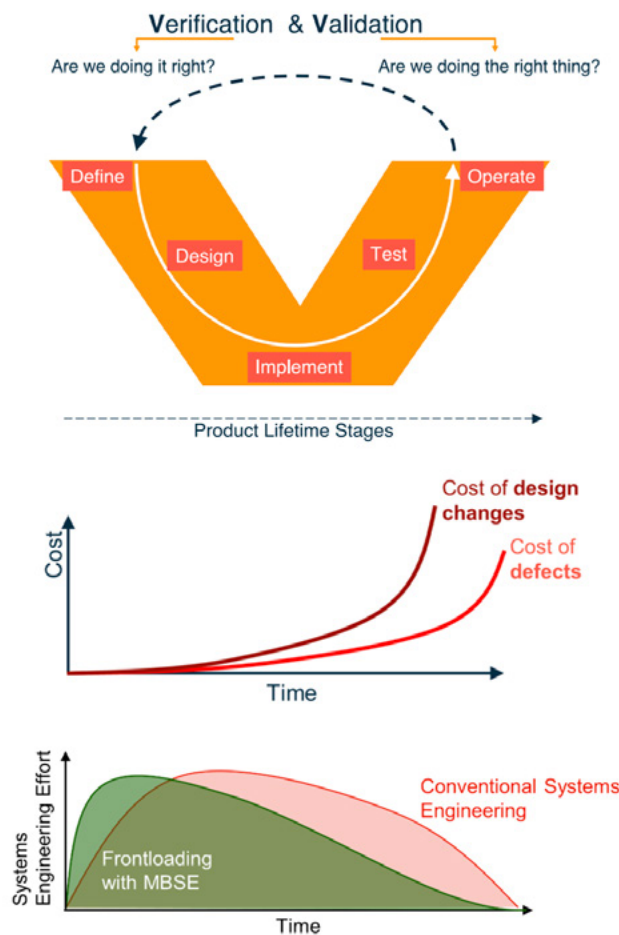
Verification and Validation through a product development, MBSE, traditional SE cost and cost of defects to the project (V&V figures are adapted[2]).

---

1. https://www.incose.org/systems-engineering

2. *Clarus Concept of Operations Archived 2009-07-05 at the Wayback Machine, Publication No. FHWA-JPO-05-072, Federal Highway Administration (FHWA), 2005.*

Above you can see the stages of a waterfall product development method with the V&V diagram.

V&V is designed to follow the product development and the product itself in compliance with the **requirements**. Verification asks "are we building the product right?", while Validation asks "are we building the right product?".

V&V helps minimize late defect discovery by engineering and manufacturing functions orchestrated by systems engineering costs more with issue discovery near the end. It is good to know about serious issues sooner, rather than later.

Traditionally, Systems Engineers follow a **document-based approach**. The challenge is the inability to adapt to changes and complications with collaboration due to a large number of stakeholders and moving parts in the projects that overall slows down the product development process.

Document Based SE   MBSE
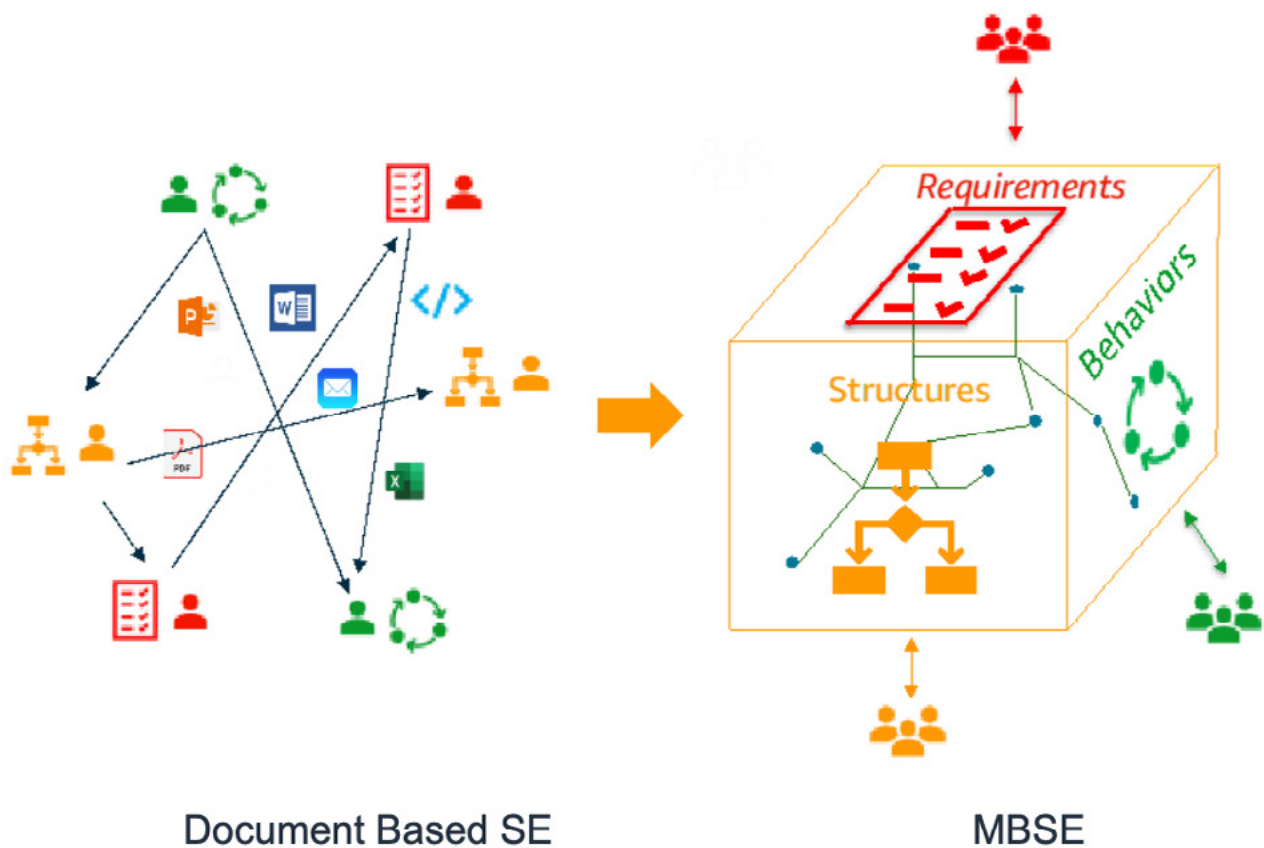
Meanwhile, in the diagram above, you can see on the left side the challenge of a traditional approach.

It often involves lots of meetings, sharing of documents, emails with attachments, spreadsheets, pdfs and models "flying" between stakeholders in a document-based systems engineering setting. It can also mean that defects are prone to be discovered near the end of the project[3].

---

3. Sterman, John. Business dynamics. McGraw-Hill, Inc., 2000.

As we established earlier, late defect discovery costs more. It is far better to become aware of serious issues sooner than later.

Without this kind of approach, you could end up with multiple "sources of truths" for engineers - creating complications such as conflicting data, using inaccurate data, or late-minute correction of source data. This not only damages product quality, but also delays projects as it slows down overall product development.

Finally, with the IoT/Industry 4.0 transition in these industries, there are - and will be - more data. The size of data further creates a burden on an already-complicated situation.
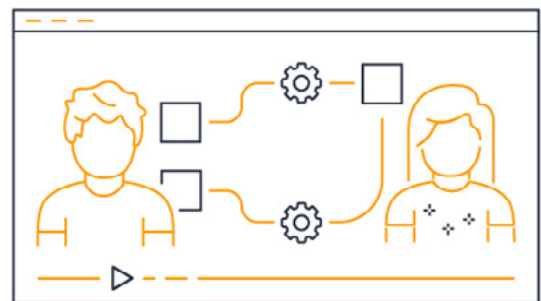
View full whitepaper online
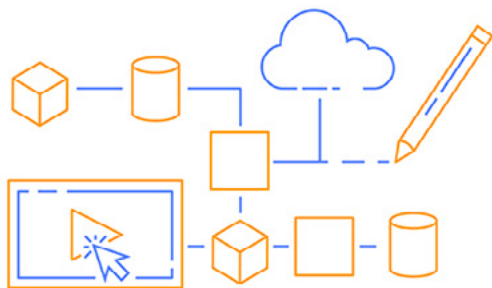
Back to Basics

A video series outlining basic architectural building blocks and best practices

This is My Architecture

A technical video series that showcases unique or innovative cloud architectures

Using Computer Vision for Product Quality Analysis in Plants

Detect and act on product-defect-classification using AWS IoT and AI/ML services.

Use this architecture for camera-based, end-of-line quality inspection; defect-detection using image classification in the cloud; alert notifications; real-time actuation; and root cause analysis using process data and inferred vision results.
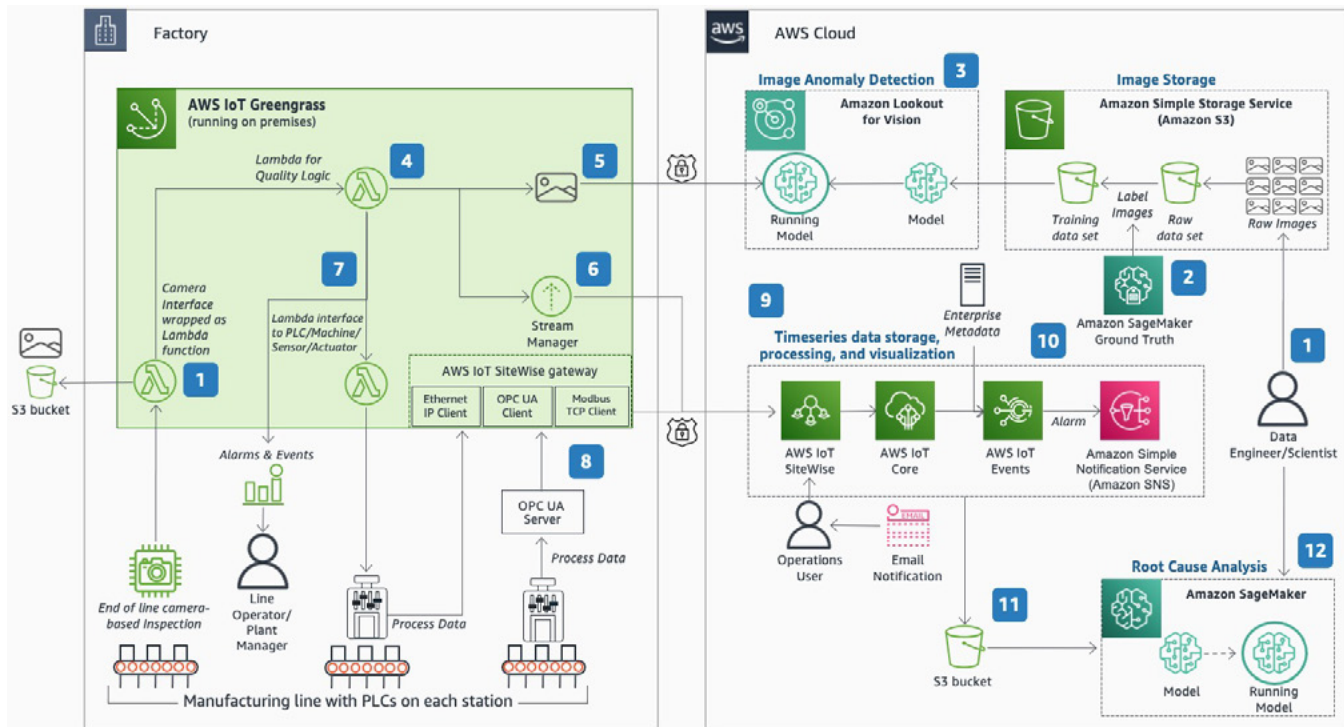
1. Auto-upload training images from the manufacturing line camera to **Amazon Simple Storage Service (Amazon S3)** using a **AWS Lambda** function running on **AWS IoT Greengrass**, or manually batch-upload training images to **Amazon S3**.
2. Use **Amazon SageMaker Ground** Truth to label training images.
3. Train a model using **Amazon Lookout For Vision (Lookout for Vision)** and deploy the trained model for running production inferences.
4. Feed live production images to **AWS Lambda** function on **AWS IoT Greengrass** to perform anomaly detection
5. Present camera image to **Lookout For Vision** for anomaly detection using inference API.
6. Feed inference metadata to **AWS IoT SiteWise** in the cloud for further processing via **AWS IoT Greengrass** stream manager.

7. Perform automated action on machine of concern and/or notify plant personnel from **AWS Lambda** function (using **AWS IoT Greengrass** connector for the **Amazon Simple Notification Service (Amazon SNS)**).

8. Ingest process data into **AWS IoT SiteWise** gateway running on **AWS IoT Greengrass** from machine/equipment using Open Platform Communications Unified Architecture (OPC UA) as the standard protocol. Modbus Transmission Control Protocol (TCP) and Ethernet IP are also natively supported, with the **AWS IoT SiteWise** gateway sending PLC data to cloud.

9. Compute Key Performance Indicator (KPI) metrics from process data in **AWS IoT SiteWise**. Create monitoring and KPI dashboards in **SiteWise Monitor** for operations user.

10. Create events from plant data and enterprise metadata by routing data to **AWS IoT Events** via **AWS IoT Core**, and send out email or text notifications to operations user using **Amazon SNS**.

11. Send process and vision inference data streams to **Amazon S3** for training root cause analysis models.

12. Train and run model inference to pinpoint root cause using **Amazon SageMaker**.

View reference architecture online

VIDEOS

AWS re:Invent 2020: Advancing the future of space in the cloud

Aerospace and intelligence companies are going all in on AWS to automate and scale space operations. Take a deep dive into how Maxar Technologies is using AWS to advance the future of space in the cloud using AWS Ground Station, AWS storage solutions, machine learning, and high-performance computing to predict where clouds and storms will be in order to deliver actionable earth intelligence to the world.
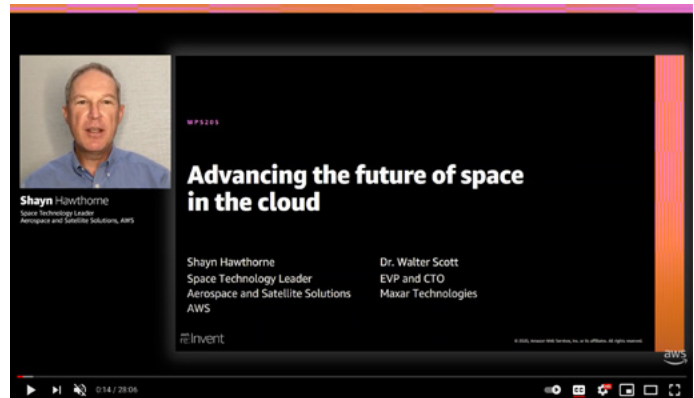

AWS Connected Aircraft Overview | AWS Events

Overview of AWS vision and capabilities for connected aircraft. Learn more about AWS connected vehicle solution at https://amzn.to/3sz645j.

Avio Aero, a GE Aviation Business: Serverless Application to Manage Expense Purchase Approvals

In this episode of This Is My Architecture - https://amzn.to/2GCSsPP, Valerio Lazzeris from GE Avio explains how they developed a serverless application for their finance team to manage expense approvals and purchase orders.
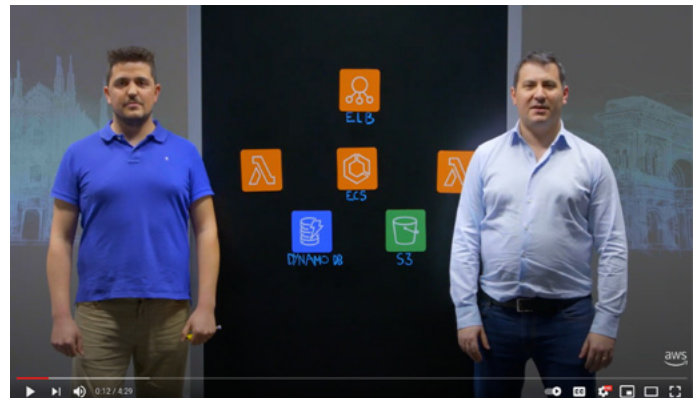
AWS Vision for Model-based Engineering in Aerospace | AWS Events

AWS vision for model-based engineering in the cloud. Learn more about AWS manufacturing at https://amzn.to/3D2BW7d.

Cybersecurity and Compliance for Aerospace | AWS Events

Overview of AWS capabilities for secure workloads. Learn more about AWS GovCloud (US) at https://amzn.to/3k7Rzlt.

Product Lifecycle Management for Aerospace | AWS Events

Overview of AWS solutions for cloud-based PLM workloads in the aerospace industry. Learn more about AWS at https://amzn.to/3xZwnmj.