Reducing decision time for on-orbit operations with virtualized ground stations and machine learning

Carmen Reglero Andres
Amazon Web Services
Eloy Salcedo de Zarraga
Amazon Web Services
Shayn Hawthorne
Amazon Web Services
Margaret Cote
Amazon Web Services
Nicholas Ansell
Amazon Web Services
Yudhajeet Dasgupta
Amazon Web Services

1. ABSTRACT

Decision time for on-orbit operations is shrinking, driving military and commercial space system operations to minimize latency by generating insights from geospatial data more rapidly, so space systems can react faster and execute defensive operations as soon as satellite threat indicators are detected. This responsiveness drives the need to increase efficiency, agility and responsiveness through virtualized solutions. Software-based ground segment architectures connected to the antenna systems where the data is ingested reduces provisioning times and costs and increase security, while maintaining high availability. Cloud services enable the virtualization of the ground segment at the edge and optimize space operations and tasks, such as tipping & cueing, with cloud data analytics and machine learning services for real-time monitoring and more timely detection of irregular activities. This paper explores how ground station virtualization and cloud services can improve space operations and reduce latency to analyze insights and autonomously transmit defensive satellite command to satellites during a single pass.

2. INTRODUCTION

There are more than 5,000 active satellites orbiting Earth [1], a number that is expected to triple over the next decade [2]. Spacecraft operations complexity has also rocketed with increased orbital activity. Satellite operations have shifted from transmit-and-receive functionality to complex operations such as tipping and cueing and re-configuring satellites communication payload on demand. Weather satellites and Earth observations are able to capture and provide information in almost real time in a way that was impossible to imagine some years ago.

Operations in space are typically controlled on the ground, so scaling ground stations is vital to support the growth in space. However, scaling ground stations has become challenging due to the lack of interoperability between ground systems and the limited software-based solutions. Spacecraft need ground stations to download their space data, receive telemetry, tracking, and commanding (TT&C) functions, and to deliver connectivity. Managing the new scale and complexity in space with traditional ground station solutions can be a challenge.

The growth in space operation and number of launches has also increased the probability of catastrophic collisions between debris, probability that will keep increasing in the coming years. The European Space Agency (ESA) estimated in 2021 that 36,000 objects roughly 10 cm or larger are orbiting Earth, and only 13% of these are actively controlled [3]. And more than 1,000,000 objects larger than 1cm and smaller than 10 cm are orbiting the Earth as of July 11th, 2022 [4]. Due to their high speed in orbit, every object threatens satellites and spacecrafts, generating an increased complexity of space missions. The lower earth orbit is especially endangered [5].

To adapt and take decisions faster, ground stations need to increase flexibility, simplify their operations, and embrace software-based solutions. Software-based solutions are possible only with digital signals, but most of the antenna systems today use analog intermediate frequency (IF). Digitizing the IF signal to and from the antenna is the first step for the digital transformation of ground stations. In this process, the analog radio frequency data is converted into digital and vice-versa.

Once the satellite data is digitized, it can be processed using software defined radios (SDR) hosted on generally available cloud services at the edge and in AWS Regions, and data can be analyzed and applied machine learning (ML) and artificial intelligence (AI) models. Complex decisions, such a collision avoidance with a debrief not actively controlled or other threats can be taken faster and without human interaction.

3. SOFTWARE-BASED GROUND SEGMENT ARCHITECTURES

Virtualization of IT architectures can help scale, improve reliability, distribute and store data, avoid costly hardware refresh cycles, and innovate and adapt to change more quickly. With the cloud, users can deploy virtual servers, storage drives, databases and other resources that are tailored to each spacecraft operation needs with a single API call or a few clicks of a mouse. In addition, the ability to manage these resources as code increases the possibilities to automate implementation and configuration steps, reducing the manual effort required to handle hundreds or thousands of satellite contacts.

Virtualization makes it simple to build reliable architectures, plus it increases the ability to recover from failure automatically. For example, you can launch several Amazon Elastic Compute Cloud (Amazon EC2) instances using a single API call to provision redundant signal processing paths in a virtualized ground segment with ease. In the cloud, replacing a virtualized function that exhibits anomalous behavior can be as simple as terminating an underlying instance and replacing it with a new one. The cloud also improves reliability by integrating automated failover functions directly into control planes.

The same gains can be achieved in traditional satellite ground segment architectures as IF signals are digitized closer to the satellite antennas. Once analog radio frequency (RF) data is digitized, i.e. converted from an analog radio frequency signal to a digital form, it can be simply distributed to other locations over long distances with no data loss or signal degradation, or stored for later processing. As they are virtualized, these functions can be hosted on general purpose hardware that is ubiquitous, and end users can delegate the responsibility of the hardware onto the cloud services provider.

By getting rid of the undifferentiated heavy lifting of procuring and maintaining hardware, traditional and New Space companies can invest more resources on reducing the timelines from mission design to launch. And for vendors of digital signal processing (DSP) solutions, a software-based approach enables building modular products and react more quickly to the needs of satellite operators, as building and distributing software is much simpler than doing the same with hardware.

The reference architecture depicted in Figure 1 shows how AWS services can support multiple stages of ground segment virtualization in the cloud. In this section, we focus our attention on stages 2-4 as those more intimately connected with the acquisition, demodulation, and decoding of the satellite data.

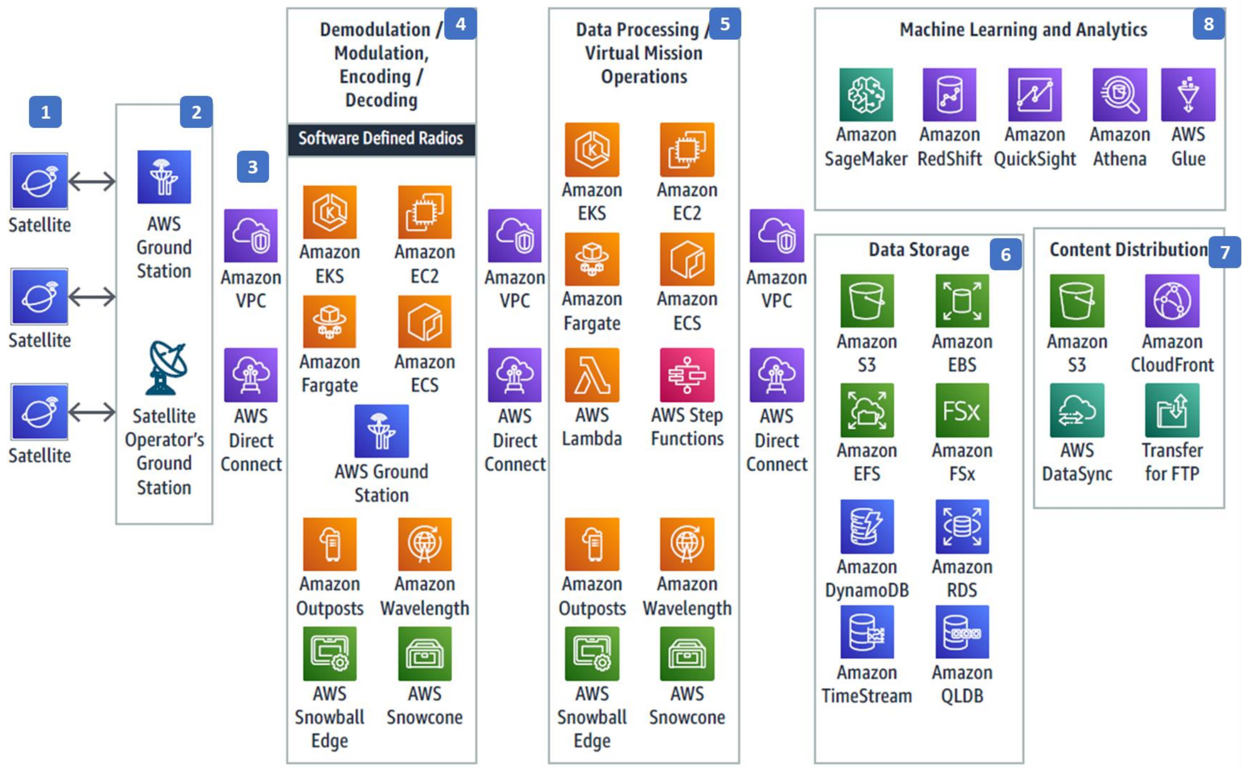

Fig. 1. A reference architecture using AWS services that covers several stages of ground segment virtualization [25]

Stage 1 represents satellites orbiting the Earth communicating with ground stations using radio frequency (RF) communication links.

In stage 2, radio signals are received from the satellite (downlink) or transmitted to the satellite (uplink) by the ground station's antenna system. When receiving data, the analog RF is amplified to compensate for path loss, digitized into a digital intermediate frequency (DigIF) stream, then passed on for demodulation and decoding. The key component which enables the translation between analogue and digital is the digitizer. When transmitting data, DigIF is received from the modulation and encoding stage (Stage 4), converted to an analog RF signal, amplified, and then transmitted to the satellite. For this stage, satellite operators can choose to use AWS Ground Station, their own antenna system and third-party antenna systems. AWS Ground Station [27] offers the ability to digitize radio frequency signals as part of the managed service, while third-party ground stations may need to introduce digitizers capable of translating between the analog and digital radio frequency domains. Existing ground segment operators may choose to use their own ground stations, or use AWS Ground Station to augment their existing footprint.

In stage 3, DigIF data is transported between the ground station's antenna system (Stage 2) and the (de)modulation and (de)encoding stage (Stage 4) differently depending on the ground stations used. If using AWS Ground Station, data travels along a resilient, high-bandwidth, low-latency, fully AWS-managed network directly into customers' Virtual Private Clouds (VPCs). If using third-party ground stations, data can travel over an AWS Direct Connect [28] network link provisioned to the nearest AWS regional point of presence for low latency.

In stage 4, during downlink operations, the DigIF stream received from Stage 2 is demodulated and decoded into raw satellite data streams (e.g. Earth observation data received). During uplink operations, the opposite occurs: data streams (e.g. commands) are encoded and modulated into DigIF streams, then sent to Stage 2 for transmission to the satellite. Software-define radios (SDRs) are a core component of the architectures that virtualize the ground segment. SDRs package complex digital signal processing (DSP) algorithms, previously available only in specialized hardware, into software that can be hosted in the AWS Cloud. Commercial SDR solutions are available on the AWS Marketplace which have been tested in AWS and integrated into virtualized ground segment architectures by AWS customers.

Satellite operators can use multiple types of compute services available in the cloud to process raw satellite data and derive useful information and intelligence. Amazon Elastic Compute Cloud (Amazon EC2) provides secure, resizable compute capacity in the cloud. Amazon Elastic Container Service (Amazon ECS), Amazon Elastic Kubernetes Service (Amazon EKS), and AWS Fargate help with the deployment, management, and scaling of containerized applications in the cloud. Edge offerings like AWS Outposts can locate software modems, such as SDRs, closer to the antennas in 3rd party teleports without the burden of managing the underlying compute and storage infrastructure. Occasionally-connected and tactical environments can leverage portable devices within the AWS Snow Family, like AWS Snowball edge appliances and AWS Snowcone.

The raw satellite data streams obtained from the demodulation and decoding processes can feed downstream applications which can also be built within the AWS cloud as described in the next section.

4. DATA ANALYTICS AND MACHINE LEARNING (ML)

Space companies already use data analytics and ML models to process satellite imagery data and take decisions more rapidly. ML enables space companies to derive results quicker, something that’s critically important when it comes to catastrophic events and in conflict situations.

For example, Fireball International facilitates automated, real-time detection and tracking of wildfires with ML and AI platforms in AWS [23]. Fireball International receives and processes more than 2.5 million images and over 30GB of satellite data per day. It stores the images on Amazon S3. These images are then aggregated to detect and predict wildfire locations, using ML models, which Fireball International continuously builds, trains, and deploys. Through its use of ML, Fireball International can detect wildfire outbreaks and predict how they might spread, with zero false positives. By automating this process on Amazon Elastic Kubernetes Service (EKS), Fireball International is able to scale its ML applications when required, allowing it to derive insights from image data quickly and accurately. With Amazon Simple Notification Service (Amazon SNS) and Amazon Simple Email Service (Amazon SES), it can alert frontline responders five times faster than the average time it would take for the relevant authorities to be alerted through human reporting.

HawkEye 360 provide another example of how to use of ML models to speed up data analysis [24]. HawkEye 360 is a commercial radio frequency (RF) satellite constellation data analytics provider, that uses Amazon SageMaker Autopilot to rapidly generate high-quality AI models for maritime vessel risk assessment, maintain full visibility and control of model creation, and provide the ability to easily deploy and monitor a model in a production environment. HawkEye 360’s data lake includes a large volume of vessel information, history, and analytics variables. With such a wide array of RF data and analytics, some natural data handling issues must be addressed. Sporadic reporting by vessels results in missing values across datasets. Variations amongst data types must be considered. Previously, data exploration and baseline modeling typically would take up a large chunk of an analysts’ time. After the data is prepared, a series of automatic experiments is run to narrow down to a set of the most promising AI models, and in a stepwise fashion from there, to select the one most appropriate for the data and the research questions. For HawkEye 360, this automated exploration is key to determining which features, and feature combinations, are critical to predicting how likely a vessel is to engage in suspicious behavior.

Fig 2 shows an architecture that brings satellite data from the space through the antenna, into the AWS cloud [22], and triggers Amazon SageMaker to accelerate latency between the data reception and the decision.

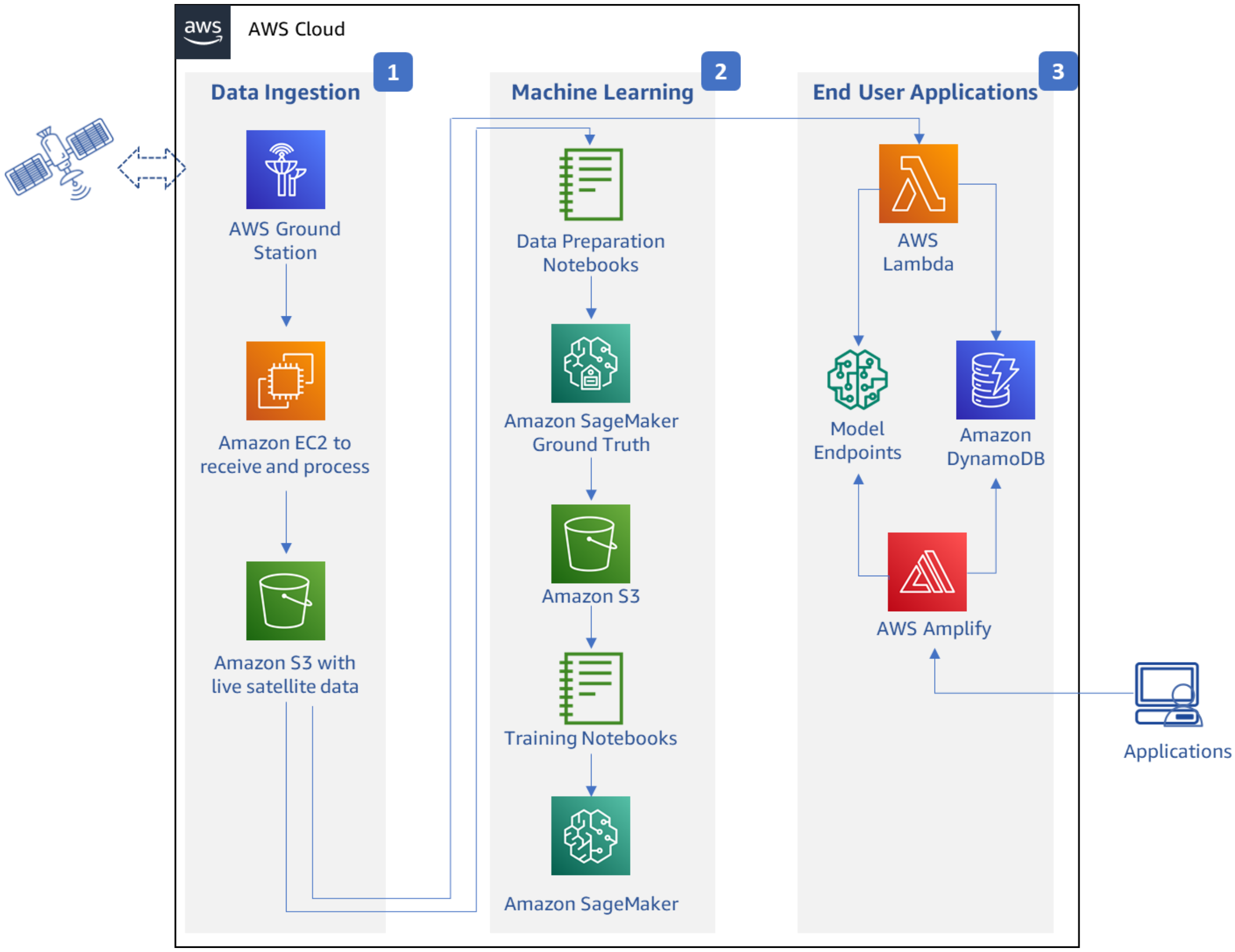

Fig 2: Reference architecture for detecting extreme events from space

Fig. 2 presents an architecture that brings satellite data from the space through the antenna system into the AWS Cloud. This architecture leveraged Machine Learning stack getting deployed at the front end. In Stage 1, the data is received in an AWS Ground Station antenna system, that digitize the IF over IP. With AWS Ground Station, the users can demodulate and decode the data using SDR running on Amazon EC2 instance. Once the satellite data is demodulated and decoded, it is stored an AWS Simple Storage Service (S3) bucket.

Once the data is ingested into the Amazon S3 bucket in Step 1, the satellite data is used by Amazon SageMaker in Step 2. Amazon SageMaker enables to build ML models and gets them ready to connect to the training data, and to select and optimize the most suitable algorithm and framework for the application. Amazon SageMaker includes hosted Jupyter notebooks that simplify to explore and visualize the training data stored in Amazon S3. In order to retrain your models in Amazon SageMaker, users create a workflow that use the satellite image from the Amazon S3 bucket to continuously improve the training data and iterate on the models.

Ingesting the data into an AWS S3 bucket in Step 1 triggers an event to invoke a Real-time Inference AWS Lambda function in Stage 3. AWS Lambda is a serverless, event-driven compute service that lets users run code for virtually any type of application without provisioning or managing servers. This AWS Lambda function calls the Amazon SageMaker endpoints to do inferences on satellite data and store the results in an Amazon DynamoDB Database. With AWS Amplify, users can connect their end-user applications to their backend. AWS Amplify [26] is a set of purpose-built tools and features that lets frontend web and mobile developers quickly and easily build full-stack applications on AWS, with the flexibility to leverage the breadth of AWS services as your use cases evolve.

5. LEVERAGE GROUND SEGMENT VIRTUALIZATION, DATA ANALYTICS AND MACHINE LEARNING TO REDUCE TIME TO INSIGHTS

In this section we'll explore how to combine the virtualization of the ground segment and data analytics and machine learning models in the AWS Cloud to reduce time to space insights.

Satellite operators need a mission operation center that controls satellite communications and processes mission data. Based on the growth of the spacecraft operations complexity and the increased orbital activity, satellite operators need a scalable, flexible and extensible mission operations center that evolves with the business requirements.

With the satellite operation center in the AWS Cloud, satellite operators have the ability to scale up and down as needed to support unique traffic peaks, such as survivability scenarios, collision avoidance and periods with increased satellite contacts. Additionally, satellite operators need highly-available, secure and resilient architectures to ensure their satellites can be operated.

Satellites and satellite control systems generate terabytes of data per day from TT&C and payload data. Virtualizing ground stations and using machine learning models based on spacecraft data is relevant to generate autonomous decisions and reduce time for on-orbit operations.

Fig. 3 presents a reference architecture for satellite operations center using AWS, where the tasking orders for satellite commanding are sent by machine learning models, trigger by satellite data or scheduled events.

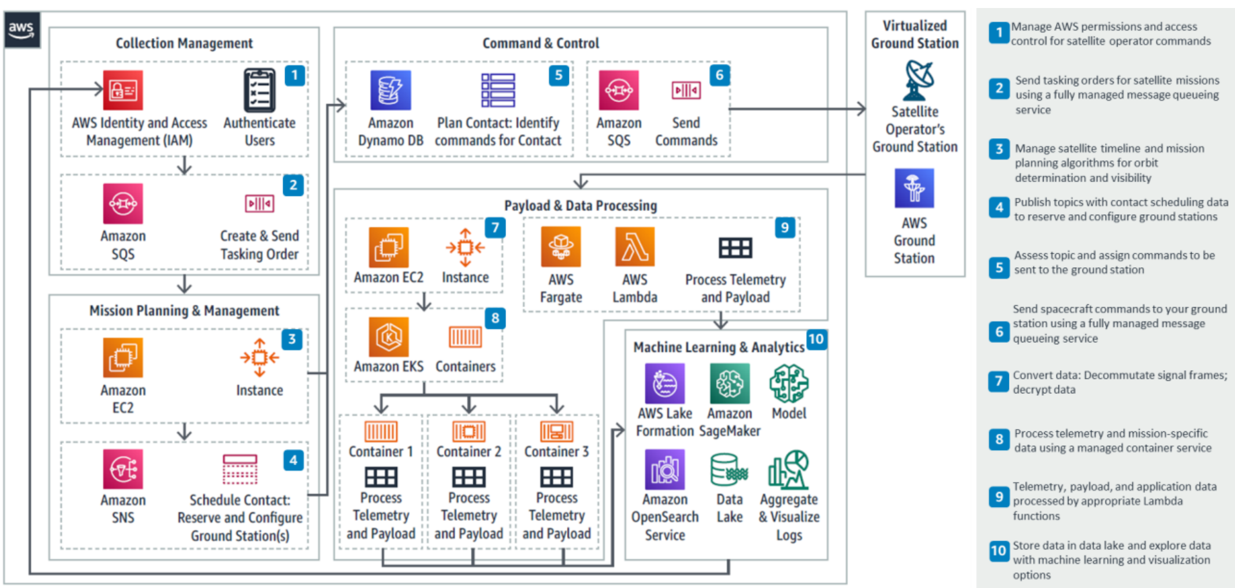

Fig. 3: AWS reference architecture for satellite operations center

Stage 1: Manage AWS permissions and access control

The system needs to validate permission and access control, so only the right users and applications can send commands to the satellite. Users and applications can be authenticated using AWS Identity and Access Management (IAM) [6]. With IAM satellite operators can securely manage identities and access to AWS services and resources. With IAM, satellite operators can set and manage guardrails access controls for their workforce and workloads, and manage identities across single AWS accounts or centrally connect identities to multiple AWS accounts. With IAM, satellite operators can specify who or what can access services and resources in AWS, centrally manage fine-grained permissions, and analyze access to refine permissions across AWS.

Stage 2: Send tasking orders

Once the workloads have been identified and granted access and permission, satellite operators can use Amazon Simple Queue Service (Amazon SQS) [7] to create and send tasking orders. SQS is a fully managed message queuing service that enables satellite operators to decouple and scale microservices, distributed systems, and serverless applications. SQS eliminates the complexity and overhead associated with managing and operating message-oriented middleware, and empowers developers to focus on differentiating work. Using SQS, satellite operators can send, store, and receive messages between software components at any volume, without losing messages or requiring other services to be available.

Amazon SQS is a tool that decouples components of the satellite mission operations center. It is a queueing service to send messages between components, and it does not require the other component to be available. This supports use cases such as sending high-volume commands to the RF equipment in preparation for a satellite contact, and asynchronous processing for tasks such as ML models based on historical logs collected from the satellite.

Stage 3: Manage satellite timeline and mission planning

Mission planning use cases orchestrate and operate satellite contacts, including activities to deconflict and manage the timeline of satellites' contact, and the radio-frequency (RF) equipment being used. Maintenance activities to assess RF antenna obscura and update the telemetry and commanding databases are also included.

Satellite operators can manage satellite timeline and mission planning algorithms for orbit determination and visibility using Amazon Elastic Compute Cloud (Amazon EC2) [8]. Amazon EC2 is a web service that provides secure, resizable compute capacity in the cloud. Satellite operators can scale capacity within minutes with SLA commitment of 99.99% availability. Amazon EC2 offers choices in processor, storage, networking, and operating system, with a wide range of instance types to help meet different compute need. Amazon EC2 accelerated computing instances, such as Amazon EC2 P3 instances [9] with up to 8 NVIDIA® V100 Tensor Core GPUs and up to 100 Gbps of networking, and Amazon EC2 F1 instances [10], with field programmable gate arrays (FPGAs), use hardware accelerators, or co-processors to perform functions more efficiently than is possible in software running on CPUs.

Stage 4: Reserve and configure ground stations

Once the management of satellite timeline and mission planning algorithms for orbit determination and visibility has been done in stage 3, satellite operators can publish topics with contact scheduling data to reserve and configure their ground stations.

Satellite operators can use Amazon Simple Notification Service (SNS) [11] to schedule contact, reserve and configure ground stations. Amazon SNS is a fully managed messaging service for both application-to-application (A2A) and application-to-person (A2P) communication. Using Amazon SNS topics, the publisher systems can fanout messages to a large number of subscriber systems, including Amazon SQS queues, AWS Lambda functions, and HTTPS endpoints, for parallel processing.

A key consideration for satellite mission operators is the latency of data and messages. Satellite commanding and telemetry require low latency to support safe operations. Amazon SNS and Amazon SQS support different use cases and can be used together. Both of these services support best-effort ordering, or first-in-first-out (FIFO). FIFO topics are recommended for satellite commanding and telemetry processing, as the order in which commands are received by the satellite and the telemetry is processed is important.

Satellite operators can increase security with message encryption and privacy. Amazon SNS provides encrypted topics to protect your messages from unauthorized access. The encryption uses a 256-bit AES-GCM algorithm and a customer master key (CMK) issued with AWS Key Management Service (KMS) [12]. Amazon SNS also supports VPC endpoints via AWS PrivateLink [13], so you can privately publish messages to Amazon SNS topics, from an Amazon Virtual Private Cloud (VPC) subnet, without traversing the Internet.

Stage 5: Send commands to the ground station

The satellite operation center sends command and control instructions to the satellite to instruct or direct the spacecraft. The command and control instructions are sent through the RF equipment. Satellite operators have the flexibility to use AWS Ground Station, third-party ground stations and, to operate their own RF equipment.

Satellite operators can plan satellite contacts and identify the right command for each of the satellite contacts with Amazon Dynamo DB [14]. Amazon DynamoDB is a fully managed, serverless, key-value NoSQL database designed to run high-performance applications at any scale. DynamoDB offers built-in security, continuous backups, automated multi-Region replication, in-memory caching, and data export tools.

Stage 6: Send commands to the spacecraft through the ground station

Satellite operators can send commands to the selected ground stations using Amazon Simple Queue Service (Amazon SQS). With Amazon SQS, AWS manages all ongoing operations and underlying infrastructure needed to provide a highly available and scalable message queuing service. Amazon SQS queues are dynamically created and scale automatically so satellite operators can focus on building and growing applications quickly and efficiently. To keep sensitive data secure, satellite operators can use Amazon SQS to exchange sensitive data between applications using server-side encryption (SSE) to encrypt each message body. Amazon SQS SSE integration with AWS Key Management Service (KMS) allows to centrally manage the keys that protect SQS.

Stage 7: (de)modulate, (de)code and (de)crypt spacecraft data

The SDRs (de)modulate and (de)code the satellite data to securely transmit the maximum amount of information in the smallest bandwidth possible, based on satellite link performance and antenna size at both ends. The satellite data can be (de)modulated and (de)coded close to the antenna system with SDRs running on AWS edge services for low latency, and in the preferred AWS Region.

Launching a software version of traditional modems with a single API call or a few clicks of a mouse can increase operational agility, enabling satcom services providers to deploy new networks faster, and with backwards compatibility with the satellite networks already deployed.

The SDRs need to perform digital signal processing to support TT&C and payload data services. Amazon EC2 provides secure, resizable compute capacity, offering choices in processor, storage, networking, operating system, and purchase model, with a wide range of instance types to help meet different compute need.

Stage 8: Process telemetry and mission-specific data

Satellite operators can use Amazon Elastic Container Service (Amazon ECS) [15] to process telemetry and mission-specific data. Amazon ECS is a fully managed container orchestration service that makes it simple for satellite operators to deploy, manage, and scale containerized applications. With Amazon ECS satellite operators can plan, schedule, and execute batch computing workloads across the range of AWS services.

For telemetry and mission-specific data processing, satellite operators can also use Amazon EKS [16], a managed container service to run and scale Kubernetes applications in the cloud, to help with the deployment, management, and scaling of highly secure and reliable containerized applications in the AWS Cloud.

Stage 9: Process telemetry and mission-specific data with serverless solution

To build and run applications to process telemetry and mission-specific data without managing servers, satellite operators can use serverless options [17]. AWS offers technologies for running code, managing data, and integrating applications, without managing servers. Serverless technologies feature automatic scaling, built-in high availability to increase agility and optimize costs. These technologies also eliminate infrastructure management tasks like capacity provisioning and patching.

Satellite operators can use AWS Lambda [18] and AWS Fargate [19] to scale up and down telemetry and mission-specific data processing without manage any server. With AWS Lambda, satellite operators can process spacecraft data without provisioning or managing infrastructure. AWS Fargate is a serverless compute for containers that removes the operational overhead of scaling, patching, securing, and managing servers.

Stage 10: Apply data analytics and machine learning to spacecraft data

Satellite operations generate and collect terabytes of data every day on satellite performance, command history, telemetry processing, operator actions, anomaly resolution, and more. This data may require data translation for other components of the mission operations center to make sense of it. Satellite operators can use a data lake implementation that offers the flexibility and agility to ingest, store, find, process, and analyze both structured and unstructured data. A data lake ensures a single source of truth for all components of the mission operations center, and avoids duplicate and contradictory data.

Amazon SageMaker [21] is a fully-managed service used to quickly build and train ML models and deploy them into a production-ready hosted environment. Mission operations for satellites generate several gigabytes of logs per satellite contact; users with multiple satellites easily generate over a terabyte of logs per day. It's difficult for humans to parse that volume of data and thoughtfully draw conclusions. ML models can analyze logs for a variety of scenarios, such as degrading satellite performance or historical contact analysis and based on the results, take quick actions and send command to the satellite.

Satellite operators can use visualization tools with Amazon OpenSearch Service to explore data, run analytics, and visualize trends, to keep improving the data modeling and the machine learning models.

6. CONCLUSION

In the last years, there has been a sharp rise in the number of spacecraft in orbit and the probability of catastrophic collisions with debris. Scaling ground stations with software and cloud services is vital to support the growth in space and to reduce decision time for on-orbit operations. Leveraging the ground station virtualization and machine learning and data analytics cloud services, satellite operators can generate insights from satellite data more rapidly with the different cloud architectures presented in this paper

7. REFERENCES

[1] https://www.statista.com/statistics/897719/number-of-active-satellites-by-year/

[2] https://www.euroconsult-ec.com/press-release/new-satellite-market-forecast-anticipates-1700-satellites-to-be-launched-on-average-per-year-by-2030-as-new-entrants-and-incumbents-increase-their-investment-in-space/

[3] https://vision.esa.int/when-debris-disaster-strikes/

[4] https://www.esa.int/Space_Safety/Space_Debris/Space_debris_by_the_numbers

[5] https://www.pwc.fr/fr/assets/files/pdf/2020/12/en-france-pwc-main-trends-and-challenges-in-the-space-sector.pdf

[6] https://aws.amazon.com/iam/

[7] https://aws.amazon.com/sqs/

[8] https://aws.amazon.com/pm/ec2/

[9] https://aws.amazon.com/ec2/instance-types/p3/

[10] https://aws.amazon.com/ec2/instance-types/f1/

[11] https://aws.amazon.com/sns/

[12] https://aws.amazon.com/kms/

[13] https://aws.amazon.com/privatelink/

[14] https://aws.amazon.com/dynamodb/

[15] https://aws.amazon.com/ecs/

[16] https://aws.amazon.com/eks/

[17] https://aws.amazon.com/serverless/

[18] https://aws.amazon.com/lambda/

[19] https://aws.amazon.com/fargate/

[21] https://aws.amazon.com/pm/sagemaker/

[22] https://www.youtube.com/watch?v=u9i7UvyIooo

[23] https://aws.amazon.com/solutions/case-studies/fireball-international-case-study/

[24] https://aws.amazon.com/blogs/machine-learning/hawkeye-360-uses-amazon-sagemaker-autopilot-to-streamline-machine-learning-model-development-for-maritime-vessel-risk-assessment/

[25] https://aws.amazon.com/blogs/publicsector/virtualizing-satellite-ground-segment-aws/

[26] https://aws.amazon.com/amplify/

[27] https://aws.amazon.com/ground-station/

[28] https://aws.amazon.com/directconnect/