

NASA SPACE COMMUNICATIONS AND NAVIGATION APPROACH TO CLOUD CAPABILITY IMPLEMENTATION

Justin Huebner

National Aeronautics and Space Administration, 300 E Street NW, Washington D.C. 20546,
202-358-2959, justin.p.huebner@nasa.gov

Philip J. Baldwin

National Aeronautics and Space Administration, 300 E Street NW, Washington D.C. 20546,
202-358-1958, philip.j.baldwin@nasa.gov

Wendy C. Evans

National Aeronautics and Space Administration, 300 E Street NW, Washington D.C. 20546,
202-358-3030, wendy.c.evans@nasa.gov

Deepak Kaul

Agile Decision Sciences, 8800 Greenbelt Rd., Greenbelt, MD 20771, 301-286-1058,
deepak.kaul@nasa.gov

Erica L. Weir

Teltrium, 445 Union Blvd, Suite 309, Lakewood, CO 80228, 720-318-7296,
eweir@teltrium.com

Michael Krumpus

Teltrium, 6406 Ivy Lane, Suite 210, Greenbelt, MD 20770, 612-306-6123,
mkrumpus@teltrium.com

Abstract

The integration of cloud services and capabilities across government and commercial sectors has continued to increase over the last several years. Cloud architectures provide on-demand process, storage, and transfer of data as a service without the need to purchase expensive hardware, infrastructure, and their associated maintenance costs. Most cloud environments charge the user for only the resources they use, making cloud implementations highly scalable and easily adjustable to the user's changing demands.

Within the National Aeronautics and Space Administration (NASA), cloud services have become an essential part of basic IT infrastructure, as well as science data processing, distribution, and management. Specific to space communications, integration began within the Space Communications and Navigation Program (SCaN) and its Near Space Network (NSN). The first completed phase delivers mission science data from the ground station to the cloud, completes level-1 processing in the cloud, and delivers virtual channel-separated data to the mission. This capability currently supports missions such as the Plankton, Aerosol, Cloud, ocean Ecosystem (PACE) mission, and is planned to support the NASA-ISRO SAR Mission (NISAR), Nancy Grace Roman Space Telescope (ROMAN) and the Spectro-Photometer for the History of the Universe, Epoch of Reionization and Ices Explorer (SPHEREx) mission.

As SCaN has progressed in its cloud integration efforts, the focus has expanded to address the network holistically – spanning both the Deep Space Network (DSN) and the NSN – as the SCaN Program Operation Cloud (SPOC). SPOC is intended to reduce overall cost to SCaN and its projects, as well as introduce efficient interfaces and processes. This paper describes the overarching implementation plan for SPOC and key elements thereof, including the first NASA cloud instance utilizing Amazon Web Services (AWS) ProServ, ensuring adherence to federal security regulations, integrating lessons from across the NASA community, and preparing for the future of artificial intelligence in the networks. Progress to date will be discussed, as well as the interplay with other strategic SCaN efforts to shift to an increasingly commercial service posture.

1. Context of Cloud Infrastructure for SCaN

The NASA SCaN Program is responsible for provisioning communications and tracking services for over

one hundred spacecraft users, from the launch pad to deep space. SCA_N's portfolio includes the Deep Space Network (DSN) and the Near Space Network (NSN), the latter of which is comprised of space-based relay assets, government-owned-and-contractor-operated (GOCO) ground network assets, and commercial partner ground network services. In effect, SCA_N ensures connectivity, regardless of how that connectivity is achieved – through GOCO assets or commercial services.

Based on National Space Policy [1], and Agency guidance, SCA_N has set goals to increase levels of commercial direct-to-Earth services, and to migrate from the government Tracking and Data Relay Satellite services (for new missions) to commercial space-based relay services by 2031. The impetus to leverage commercial industry capability comes from both a desire to foster industry growth, and to spend government dollars more efficiently.

As SCA_N seeks to continue this major transition, delivery of the mission data comes into focus. Among the technologies that could provide transformational capabilities to SCA_N and NASA missions, cloud services can arguably be placed at the top of the list. Commercial and private cloud capabilities enable the virtualization of the space-to-ground networks and improve ability to interface and interact with missions in more efficient ways. Benefits include:

- Providing entire services virtually, delivering data to virtual secure mission repositories and eliminating or greatly reducing the need to replicate infrastructure at each ground communications location.
- Simplifying mission support by removing the need for site-based mission unique equipment or processing.
- Removing barriers to scalability by eliminating the exponential growth in dedicated circuits as new providers / provider infrastructure is incorporated into the network.
- Reducing overall cost associated with operating and maintaining physical network infrastructure including facilities, power, etc.
- Reducing complexity in securing the SCA_N infrastructure, centralizing security implementations to virtualized infrastructure and facilitating improved mission security.

2. Overview of SCA_N Networks and the First Steps Toward Cloud Integration

The Near Space Network and Deep Space Network represent a complex set of ground and space assets designed to facilitate connectivity between missions in space and their operators and data end-users.

The NASA-owned network's ground stations are located at the White Sands Complex (WSC) in New Mexico; U.S. McMurdo Antarctic Station; Wallops Flight Facility (WFF) near Chincoteague, Virginia; North Pole, Alaska; and the University of Alaska in Fairbanks, Alaska. The network also purchases services from commercial providers in Hawaii, Norway, Sweden, Singapore, South Africa, Australia, Antarctica, and Chile. SCA_N also operates the Launch Communications Segment, a series of antennas near Kennedy Space Center (KSC) to support launches of NASA human exploration missions, including Artemis using the Space Launch System (SLS) and Orion.

The space-based relay component of the NSN is composed of geosynchronous (GEO) Tracking and Data Relay Satellites (TDRS), and three ground stations: the primary station at White Sands Complex, New Mexico, and two smaller stations in Guam, and Blossom Point, Maryland.

The three DSN complexes are stationed at approximately 120-degree longitudinal separations around the globe to provide continuous coverage for mission spacecraft spread throughout the solar system. Two of the complexes are in the Northern Hemisphere, in Goldstone, California, and Madrid, Spain; the third is in the southern hemisphere in Canberra, Australia.

This legacy network architecture relies on direct connections between communications networks and missions. Data transport is established via one-to-one connections between various Service Link Providers (SLPs) and Mission Operations Centers (MOCs) located worldwide, as shown conceptually in the left portion of Figure 1.

As a first step toward migrating away from this point-to-point and less flexible architecture, SCA_N completed a pilot project called the Data Acquisition Processing and Handling Network Environment

(DAPHNE). DAPHNE was implemented in the AWS cloud and deployed for use by several newer missions as an integral part of more capable ground stations in Alaska, Chile, Virginia, and Svalbard. Data was captured at the ground stations and delivered to AWS S3 buckets. Level 1 processing was performed in an automated manner then delivered to AWS S3 buckets for further processing by the missions. This first demonstration of cloud in the NASA/SCaN context was critical, as it illustrated the benefits of migrating 90% of the infrastructure from on premises to the cloud. With this in mind, we pivot to a discussion of the long-term vision.

3. SCaN Program Operation Cloud (SPOC) Vision

The SCaN Program Operation Cloud (SPOC) is intended to provide a single logical interface for fulfilling customer mission scheduling, data transport, and network management needs through government and commercial communications services providers. It will offer a virtual secure private mission cloud environment where data processing (as applicable) and distribution is geographically distributed and strategically located near the missions' Science Data Segment (SDS) and Mission Operations Center (MOC). This architecture is scalable with a modular approach, which eliminates the need for mission-specific hardware at ground stations.

The future state for SCaN cloud services architecture, shown in the right panel of Figure 1, connects SLP antennas to a cloud provider such as Amazon Web Services (AWS). In this architecture, each SLP chooses how to connect to the cloud, removing the burden from NASA and giving the provider freedom to make appropriate business choices. In its partnerships with commercial providers, NASA stands ready to collaborate and provide guidance as applicable to ensure integration is successful.

This flexibility will streamline the integration of new commercial providers and international space program partners with SCaN operations. Similarly, all MOCs using the SPOC will connect, as well as a notional SCaN Network Operations Control Center (SNOCC), all using the most appropriate method given their location and network capabilities. The cloud architecture also provides the ability to standardize the mechanisms and patterns for scheduling and control interfaces as well as the data delivery pipeline to missions. The result is a virtualized environment for ground station operations and data delivery involving the SLPs supporting space missions, the MOCs, and the SNOCC.




Figure 1: Legacy architecture with point-to-point connectivity (left); Cloud implementation approach (right).

Architecturally we can view the implementation of cloud capabilities for NASA SCaN purposes in two ways as depicted in Figure 2. First, is the cloud backbone at the core of the cloud environment, which provides connection to the physical users, network assets, government assets, and commercial provider assets equally, and enables functionality through applications in/on the cloud: (1) scheduling, (2) data delivery, (3) network management. The second view is contextual. SPOC extends beyond just network operations (shown in the center of the figure)—any project within SCaN will be able to leverage the cloud environment to meet development (left) and/or the monitoring and optimization needs of the program (right). It is critical to the program and its projects to utilize common resources. Rather than projects and sub-projects establishing individual relationships with cloud providers like AWS and Microsoft Azure, SPOC establishes the common environment. This was one of the primary motivators for SCaN to pursue the SPOC. Historically, there has been no NASA group which provides 24/7 support for mission critical operations related to the cloud. The burden has relied on developers, which

is not sustainable or scalable. SPOC grew out of the need to address this issue. Efficiencies can be realized with a single centralized cloud environment including common security policies, budget controls and cloud monitoring. SPOC's goal is to create a single cloud environment which supports mission critical operations as needed and leverage economies of scale to reduce costs.




Figure 2: Illustration of the cloud backbone, application environment and physical connectivity of users and network assets to the cloud.

4. Conceptual Approach to Realize the SCA_N Program Operation Cloud

Virtualization of space communications operations and data delivery has been implemented by several commercial ground-station-as-a-service (GSaaS) offerings like AWS Ground Station. Those offerings abstract away the need to build ground operations by providing the antennas and associated infrastructure. SCA_N's challenge is different in the sense that this infrastructure already exists, but SCA_N seeks to virtualize it for efficiency and flexibility in a cloud environment. This means that the journey is a technology transformation that leads to a similar end-state architecture as commercial GSaaS offerings, but with the additional challenges of dealing with legacy systems that may be difficult or unwise to replace. Further, SCA_N must also address the challenge of creating an environment which enables acting as a hub for a portfolio of diverse commercial service providers (both ground and space-based) as well as a diverse set of users.

When viewed at a high level, this technology transformation is like those faced by many large IT organizations. In the simplest sense, it seeks to remediate a topology problem of having many costly point-to-point integrations by replacing them with a consistent, centralized set of solutions for network operations. A typical approach to this transformation for ground station operations for example, would be to implement each functional area using cloud technologies. For example, implement control, contact scheduling, data delivery, mission data processing, ground station monitoring, etc. using appropriate cloud technologies. However, when concentrating on the functional aspects of the domain, the result is often a collection of bespoke implementations of each, which can easily lead to inconsistency and lack of cohesion in the final architecture.

A better architecture approach to such a transformation is to first decompose it into familiar and tractable IT problems, and then apply typical cloud-based solutions to each one. This should be done without focusing too much on the specifics of the ground operations functions -- at least initially -- to establish an overall cloud architecture that can meet the needs. When the transformation is considered at this more abstract level, it is easier to see that there is nothing particularly unique about the needs; users request/require data, data is provided by a heterogeneous set of providers, the data must be processed, and then delivered to end users. The standard cross-cutting concerns of security, monitoring, robustness

and resiliency, metrics, analytics, and notifications are all needed, just like in any other cloud implementation. Table 1 below maps the domain-specific needs of ground station operations to familiar IT problems and then to typical AWS solutions. *This example illustrates the approach to the problem, rather than the final solution.*

**Table 1: Example Mapping of SCA_N Needs to Familiar IT Problems
and Typical Cloud Solutions**

Domain-Specific Needs		Familiar Domain-Independent Problem	Typical Cloud Solutions (AWS examples)
Cloud Backbone / Adoption	- Connect service link providers (SLPs) to a central location for operational needs. - Connect mission operations Centers (MOCs) to a central point to receive and process data. - Conform to all security policies. Utilize a Federal Risk and Authorization Management Program (FedRAMP) certified environment. - Centrally manage cloud accounts, access policies, and billing.	- Connect a set of heterogeneous, geographically distributed data providers to a cloud environment. - Connect a set of heterogeneous, geographically distributed data users to a cloud environment. - Implement proper identity and access management to all resources. - Ensure accountability for resource usage and control access policies.	- Direct Connect to nearest point of presence, Site-to-Site VPN or Client VPN over Internet, or PrivateLink for providers already in AWS - Direct Connect to nearest point of presence, Site-to-Site VPN or Client VPN over Internet, or PrivateLink for users already in AWS. - Design and implement Identity and Access Management (IAM) policies for users and roles to control access to AWS services and resources. - Implement an AWS organization containing AWS organizational units as needed. Consolidate billing at the organization level. Implement Service Control Policies to manage permissions to resources. Design the organizational policies to work with the IAM policy approach.
Scheduling	Allow mission operations center users to schedule/request satellite contacts with service link providers.	Provide a web user interface to view information and perform functions. Invoke services to retrieve information from one of many heterogeneous service providers. Implement logic and send information to providers.	Implement user front-end as serverless web application served from S3 bucket. Implement APIs for all scheduling functions and expose with API Gateway. Implement API service logic in Lambda functions and use AWS database service like RDS for persistence as appropriate. Send information to service providers via provider-specific interfaces.

Data Delivery	• Transfer received satellite data to a central point for processing. • Transform files of satellite data into more useful formats or end products. • Data Delivery: Process streaming data. • Deliver data products to mission users in the Mission Operations Center (MOC). • Notify users of the delivery of new data. • Automatically perform mission-specific processing upon delivery of new data.	• Transfer files and stream data from data providers to cloud environment. • Perform ETL (extract, transform, load) operations on data automatically. • Process streaming data for purposes of transformation, analytics, etc. • Deliver processed data to destination storage locations. • Generate notifications of various types (email, etc.) based on data transfer events. • Trigger data processing based on the delivery of data.	• S3 buckets for file storage, Kinesis for streaming data. Developers use AWS SDKs, CLI, or invoke APIs directly to transfer data into cloud environment. • Trigger Lambda functions on arrival of data in S3 buckets or use AWS Glue for more comprehensive ETL solution. • Implement Kinesis Data Stream Applications executing on EC2 instances or as Lambda functions to process incoming data. Developers use Kinesis Client Library. Use Kinesis Data Firehose to deliver streams to S3 buckets, databases, or other AWS services. • Write processed file data to destination S3 buckets. Write stream data to S3 buckets, databases, data warehouses like AWS Redshift, or other AWS services using Data Firehose. • Use S3 Event Notifications to detect creation of S3 objects and Simple Notification Service (SNS) to send notifications via email or SMS. • Use S3 Event Notifications to detect creation of S3 objects and trigger Lambda functions to process new data. More complex processing needs can trigger AWS Batch jobs to execute processing on EC2 instances or execute containerized workloads on AWS Fargate (serverless Elastic Container Service).
Network Management and Monitoring	• Send control directives from mission operations centers to service link providers. • Monitor ground station operations from mission operations centers and SCA Network Operations Control Center (SNOCC). • Gather utilization metrics and obtain insight into needs and return on investment. • Maintain separation of mission operations from one another to avoid any unintended disruption and facilitate proper security policies.	• Invoke services to send commands to one of many heterogeneous service providers. If needed, provide a web user interface. • Invoke services to obtain status and other operational characteristics from providers. Monitor the data movement and processing occurring in the cloud environment itself. • Expose the utilization of internal cloud resources for end users. • Implement separation of computing and network resources between user groups.	• Implement APIs to represent control requests/commands and expose with API Gateway. Implement logic in Lambda functions and use AWS database service like RDS for persistence. Send control info to service providers via provider-specific interfaces. User front-end can be implemented as serverless web application served from S3 bucket. • Implement APIs to retrieve operational status and expose with API Gateway. Implement any service logic in Lambda functions. Invoke provider-specific interfaces to retrieve operational data. Use CloudWatch for cloud application health monitoring and resource monitoring. • Use CloudWatch for cloud resource monitoring. Use AWS financial management tools like AWS Budget and AWS Cost and Usage Report to track and manage costs. • Implement distinct virtual private clouds (VPCs) for each mission. Design the approach carefully considering the identity and access management (IAM) and organizational policies (see below).

This type of mapping exercise, which SCA Network will be executing in earnest in the months ahead, demonstrates that different domain-specific functions are often the same or similar IT problem that resolve to similar cloud solutions. For example, satellite contact scheduling and monitoring/control

functions are very different from a domain-specific viewpoint but are actually very similar from the domain-independent IT viewpoint and can be solved with very similar if not identical cloud solutions. In fact, the left-hand functional column could be replaced with functional needs from a completely different industry and still map to the same familiar IT problems with standard cloud solutions.

Another advantage of this approach is that it helps mitigate the risk of vendor lock-in with a particular cloud provider. This decomposition and mapping of functions allows one to replace the right-hand column with the analogous cloud services from a different cloud provider. Figure 3 below illustrates an example AWS cloud architecture that can be used to implement the functions required for virtualization of ground station operations but does not reflect any architectural decisions made yet by the SCan team. The final architecture will, by necessity, address the need for additional resiliency – through regional diversity, redundancy and backups, secondary providers, and continuous monitoring.




Figure 3: Cloud architecture for virtualization of ground station operations.

5. Current Activities and Status

Some progress has been made in the major categories of capability as follows.

5.1 Cloud Backbone and Adoption

The SPOC landing zone has been established on an AWS platform. The SCan team and SPOC have completed the third-party assessment and expect to receive Authorization to Operate in September of 2024. This provides the foundation for all future work. The team is also building the end-to-end concept of operations, that captures the full operational flow from initiation of service request through to data delivery.

Next steps: Begin connecting NASA mission locations to AWS using Direct Connect to the nearest point of presence (POP). This will establish secure, dedicated fiber lines into the AWS cloud. Engagement with local telecom service providers is required for the "last mile" of this terrestrial connectivity. Begin connecting a single NASA ground station location (e.g., a Wallops antenna) to the AWS cloud using either AWS Direct Connect or, more simply, AWS VPN.

5.2 Scheduling

The activity currently underway is developing a proof of concept for the scheduling capability and mirror legacy network operations in the virtual environment. The goal is to complete the development of the cloud-based scheduling function in the next year.

Next steps: Integrate the scheduling function. The scheduling function will be deployed in the operational environment with a single ground station location (e.g., Wallops), but not used operationally.

5.3 Data Delivery

Next steps: Identify a candidate mission for cloud implementation. Provision AWS Virtual Private Cloud (VPC) for the mission and provision users within the Identity and Access Management policy. Build data receipt, transformation, and delivery of data using the AWS services described above within the mission VPC. Implement any mission-specific data processing using appropriate AWS services.

5.4 Network Management and Monitoring

Next steps: Implement the use of AWS CloudWatch to provide appropriate monitoring of all AWS services and data flows. Connect the SCaN Network Operations Control Center (SNOCC) to these monitoring services.

6. Looking Ahead – Considerations and Challenges

As SCaN continues to build on the SPOC environment and works toward the end-state-vision, the team will be actively monitoring wider industry as well as reviewing our own efforts to capture and integrate lessons learned and best practices. Although widely used in commercial industry cloud services integration is not fully adopted in the government primarily due to policy limitations. One challenge SCaN faces is advancing operational capability in the cloud while navigating policy limitations and the pace of change. One of the underlying concerns from a government policy standpoint is the security of data and data transfer. In the envisioned architecture, cloud-native security tools will be implemented. The system security for SPOC provides the inherent foundational work and controls for any project using SPOC, and includes features like single sign on. SCaN is also collaborating with the NASA Office of the Chief Information Officer (OCIO) to demonstrate that the cloud-native tools meet the intent of Continuous Diagnostics and Mitigation (CDM).

Security, among other considerations, will be particularly important as NASA embarks on an expanded human space flight presence at the Moon and beyond. While not ideal, robotic science missions can tolerate network outages, but in the HSF domain, this puts astronaut safety at risk. As SCaN architects the cloud implementation, encryption, regional diversity and backups, diversity of suppliers, and built-in reliability measures will all be important to consider,

Adapting to the cloud environment will require operational changes, including how we staff and train personnel supporting SCaN and the networks. Ensuring the program has the right expertise, and that current personnel with highly unique skills in space operations and networks gain the additional foundational knowledge about cloud infrastructure will be important to the long-term sustainability of the architecture.

As SCaN works toward the end-state vision, each step in migrating resources will be analysed to quantify costs and benefits and inform decisions. Utilization data will also contribute to forward planning and decisions, ensuring that use of the cloud capabilities is maximized and costs are properly controlled.

7. References

- [1] National Space Policy of the United States of America. 09 December 2020, <https://www.space.commerce.gov/policy/national-space-policy/> (accessed 08.15.2024).