17th International Conference on Space Operations, Dubai, United Arab Emirates, 6 - 10 March 2023.

SpaceOps-2023, ID # 557

NASA's Implementation of Cloud Services for Human Space Flight

Philip Baldwin[a*], Wendy Evans[a], Justin Huebner[a], Monica Saraf[a], Salem El-Nimri[b], Erica Weir[c], Daynah Rodriguez[c]

[a] *National Aeronautics and Space Administration, 300 E Street NW, Washington D.C. 20546,* philip.j.baldwin@nasa.gov, wendy.c.evans@nasa.gov, justin.p.huebner@nasa.gov, monica.saraf@nasa.gov
[b] *National Aeronautics and Space Administration, Goddard Space Flight Center, 8800 Greenbelt Rd., Greenbelt MD 20771,* salem.f.el-nimri@nasa.gov
[c] *Teltrium, 6406 Ivy Lane, Greenbelt, MD 20770,* eweir@teltrium.com, drodriguez@teltrium.com
* Corresponding Author

Abstract

Cloud services are a tried-and-true technology used throughout the United States government agencies, including the National Aeronautics and Space Administration (NASA). With reliable results and infrequent downtimes, cloud services allow for secure remote access, customizability, and streamlined monitoring options, creating an environment for better data integrity and availability. As NASA increasingly migrates functions to the cloud, the Space Communications and Navigation Program (SCaN) program has been investigating how this capability can be leveraged to provide communication services to its mission users and customers. Currently, missions such as NASA-ISRO Synthetic Aperture Radar (NISAR), the Plankton, Aerosol, Cloud, ocean Ecosystem (PACE), and the Nancy Grace Roman Space Telescope (RST) are planning to incorporate cloud into their data delivery architecture. However, SCaN is looking to expand further. This conversion to using cloud services allows for greater availability of mission data for both robotic and human space flight (HSF) missions. The SCaN program and the Near Space Network (NSN) are working to consolidate resources and create a cloud environment suitable for the entirety of the SCaN program network architecture. SCaN is finalizing its cloud architecture and soon will expand its cloud services. New services will adhere to federal regulations including the Federal Risk and Authorization Management Program (FedRAMP), which is built upon National Institute of Standards and Technology (NIST) documentation. While keeping in mind these security requirements, an auxiliary objective of the cloud integration is to ensure the most cost-efficient solution for providing a scalable, robust, and resilient system. Using cloud services, NASA will gain access to better centralized monitoring and management features, along with customizable services on a pay-per-use plan. With the ever-growing NASA mission data volume needs, maintaining ample storage space is another major constraint. Processing and storing such large amounts of data, on the order of terabytes a day, requires dynamic processing capability, which is inherently a strength of cloud computing. By routing this data from ground stations through the cloud, there will be greater ease of access for both SCaN and the user community. Artificial intelligence and other built-in cloud functions can also enhance efficiency, improving data processing time and thereby also allowing for better data availability. As we look to the future of cloud services, NASA will continue to leverage capabilities that will benefit NASA's ability to provide cost-effective communication services. This paper further outlines the evolution of cloud use by SCaN in the context of Human Space Flight.

Keywords: cloud services, space communications, networks, networking, NASA

Acronyms/Abbreviations

Acronym	Definition
AWS	Amazon Web Services
CMD	Command
CSCI	Computer Software Configuration Item
DAPHNE	Data Acquisition Processing and Handling Network Environment
DSN	Deep Space Network
FDF	Flight Dynamics Facility
FedRAMP	Federal Risk and Authorization Management Program
FISMA	Federal Information Security Modernization Act
HSF	Human Space Flight
IF	Intermediate Frequency
LEGS	Lunar Exploration Ground Segment

Acronym	Definition
M&C	Monitoring and Control
MOC	Mission Operations Center
NASA	National Aeronautics and Space Administration
NEN	Near Earth Network
NIKA	Near Earth Network (NEN) Initiative for Ka-Band Advancement
NISAR	NASA-ISRO Synthetic Aperture Radar
NIST	National Institute of Standards and Technology
NRHO	Near-Rectilinear Halo Orbit
NSN	Near Space Network
OCIO	Office of the Chief Information Officer
PACE	Plankton, Aerosol, Cloud, ocean Ecosystem
PoP	Point of Presence
RST	Roman Space Telescope
SCaN	Space Communications and Navigation
SDR	Software-Defined Radio
SDS	Science Data Segment
SLP	Service Link Provider
SNOCC	SCaN Network Operations Control Center
TLP	Terrestrial Link Provider
TRK	Tracking
UME	User Mission Environment
UMGS	User Mission Ground System

1. Introduction

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. SCaN'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 ground network assets, and commercial partner ground network services. In effect, SCaN ensures connectivity, regardless of the physical mechanism.

Based on National Space Policy [1], and Agency guidance, SCaN has set goals to increase levels of commercial direct-to-Earth services to near 100% by 2024, and to migrate from the government Tracking and Data Relay Satellite services (for new missions) to commercial space-based relay services by 2030. The impetus to leverage commercial industry capability comes from both a desire to foster industry growth, and to spend government dollars more effectively. As SCaN seeks to make this major transition, delivery of the mission data comes into focus. Among the technologies that could provide game-changing capabilities to SCaN and NASA missions, cloud services can arguably be placed at the top of the list. Commercial and private cloud capabilities create the potential to virtualize the space-to-ground networks and to interface and interact with missions in more efficient ways. Benefits include:

- *Providing entire services virtually*, eliminating the need to replicate infrastructure at each ground communications location
- *Delivering data to virtual secure mission repositories*, eliminating the complicated, and therefore costly, terrestrial communications infrastructure and the need for site-based mission-unique equipment or processing
- *Simplifying SCaN's technology infrastructure*, creating opportunities for improved total cost of ownership and lower recurring engineering costs
- *Reducing complexity in securing the SCaN infrastructure*, centralizing security implementations to virtualized infrastructure and facilitating improved mission security

In January 2021, SCaN took on a study to address the possible use of cloud services to: (1) better understand and identify drivers for cloud data delivery, from the mission community and from SCaN itself; (2) investigate services to be offered via cloud, including real-time streaming and off-line file delivery; (3) assess options for interfaces and data delivery standards; and (4) create a SCaN cloud architecture that is extensible for not only Near Space and Deep Space robotic missions, but also human spaceflight support. The results of this study helped develop a roadmap for SCaN cloud implementation starting with the NSN's near term planned usage all the way to how SCaN will support future lunar missions.

2. Legacy vs. Future Architecture

The 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. For example, referencing Figure 1:

- The MOC for Mission '1' may be supported by SLP 'A,' while the MOC for Mission '2' may be supported by SLP 'B.'
- Even though SLP 'B' may be capable of supporting the MOC for Mission '1' (e.g., in case of a spacecraft emergency), such support will not be possible without a pre-established one-to-one connection and support agreement between SLP 'B' and the MOC for Mission '1.' Further, in most cases, this approach requires that mission specific hardware is deployed at the SLP for local processing and support.

Although functional, this architecture features several disadvantages, including:

1. Adding new service providers requires establishing new interfaces, with missions and/or providers having to account for physical and computer resources for each addition
2. Missions are directly impacted by provider outages; recovery requires establishing / using alternate direct connections
3. Significant time/effort associated with sustaining mission-unique assets
4. One-to-one connections increase the complexity of managing security as Interconnection Security Agreements are required for all unique connections. As the number of one-to-one connections scales up, so does the security challenge.

NASA has been using the current approach for the past 60+ years, as the availability of global networking services was limited in the early decades of space exploration. However, over the last 10+ years, the growth of cloud services has accelerated worldwide, presenting an opportunity to transition away from one-to-one connections.

The SCaN Cloud architecture is intended to provide a single logical focal point for fulfilling customer mission scheduling and data transport 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 cloud approach, which eliminates the need for mission-specific hardware at ground stations.

The future state for cloud services architecture, shown in Figure 1, connects SLP antennas to a cloud service such as Amazon Web Services (AWS) via the nearest cloud point of presence (PoP). 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. Similarly, all MOCs that plan to leverage a cloud service may choose how to connect to the cloud PoP nearest to them. In addition to the participating SLPs and MOCs, a SCaN Network Operations Control Center (SNOCC) is also connected to the cloud service. Thus, logical connections are established between the SLPs supporting space missions, the MOCs, and the SNOCC. The infrastructure that connects all the communications network assets and customer nodes is built within a virtual network manager in the cloud.

---

[*] In the context of NexTEra, this paper will refer to a ground station as a Service Link Provider (SLP).

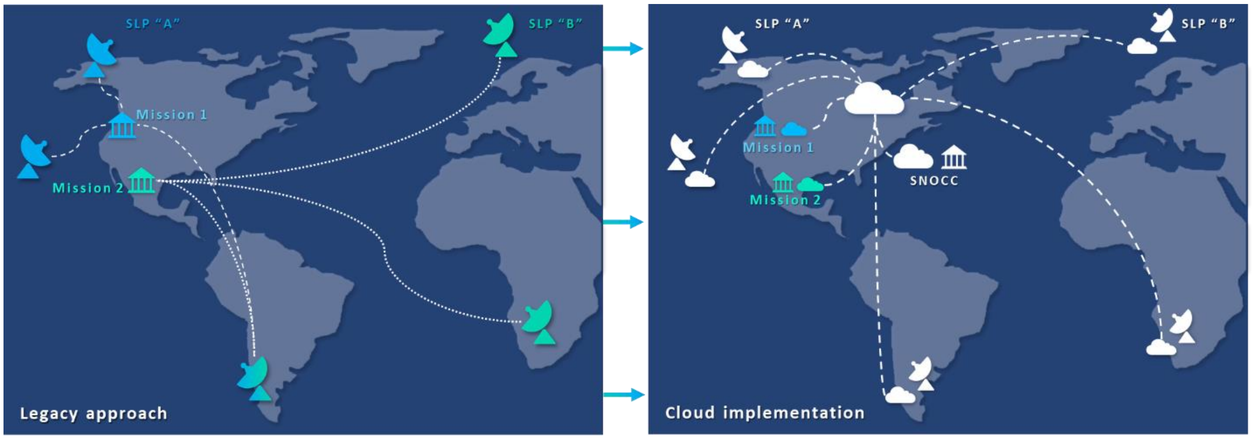

Figure 1: Legacy architecture with point-to-point connectivity (left); Cloud implementation approach (right).

In this approach, Mission 1 and Mission 2 can now access more antennas, as compared with the legacy approach of point-to-point connections between a MOC and its supporting ground stations. Missions maintain fewer interfaces, but rather will have a single logical interface through the SCaN Cloud. This in turn means establishing a connection with a new provider does not require a new interface nor the associated additional physical and computing resources. SCaN will then, from a macro level, via the NSN or DSN manage services across providers, resulting in seamless service in the event of particular ground site outages. In addition, the SNOCC does not need to be directly connected to the SLPs: the cloud services connection architecture facilitates flexibility in the provisioning of services between communications network assets and MOCs and obviates geographical restrictions on network management capabilities such as the SNOCC.

3. Modular Architecture Approach

Implementation of the SCaN Cloud architecture will be extensible and leverage modularity. This approach intentionally leverages modern software architecture forms to break the old monolithic software into loosely coupled software modules. By switching to a modular architecture, SCaN will be able to increase the resilience of systems, as each module establishes a connection with other modules but without dependencies (i.e., the failure of one module will not result in a failure of other modules). This allow for the SCaN cloud system to continue to function and sustain many failures without bringing the whole infrastructure into a halt – as is the case with legacy monolithic architectures.

3.1. First Steps in the Near Space Network

SCaN, via the Near Space Network at GSFC, first implemented this approach with a pilot project called the Data Acquisition Processing and Handling Network Environment (DAPHNE). This project leveraged cloud services using AWS capabilities. DAPHNE was designed to initially support two upcoming missions, NASA-ISRO Synthetic Aperture Radar (NISAR) and Plankton, Aerosol, Cloud, ocean Ecosystem (PACE), and laid the foundation for other missions' use of cloud services. Most recently, the NSN has worked to advance SCaN's Cloud architecture plans with a system called NexTEra[†]. The NSN's NexTEra extends the functionality and capability initiated in DAPHNE by adding more modules to service SCaN's needs.

As shown in Figure 2, the NexTEra cloud environment supports any number of User Mission Environment (UME) modules which are instantiated dynamically, scheduling, and monitor/reporting/control modules, and the user mission ground segments (UMGS) which represent mission operations and science operations centers and associated personnel and capabilities. These virtual modules are interconnected, performing specific functions while remaining independent of each other (loosely coupled). This means that, for example, the Monitoring/Control Reporting module can connect to multiple UMEs within the NexTEra environment, and the failure of one UME will not cause any impacts to the operation of the other UMEs and vice versa.

[†] Nexus, Terra, Astra; meaning to state "Connecting Earth to Stars"

3.1.1 The User Mission Environment and Modular Services

Each UME is a unique carefully preconfigured “virtual private cloud” corresponding to an individual mission where data processing and management are configured, accommodating unique data processing requirements and security considerations. The UME itself is also modular, as shown in the left-hand side of Figure 2. The independent software modules support services and functions required by the mission, including uplink services, data distribution, level 0 processing, or a software-defined radio function. At the front end of the UME is a translation layer, which can accommodate any number of data or frame formats incoming from different ground stations and providers and then translate so that the mission receives this information in their preferred format.

Offering services in a modular fashion is advantageous, saving time and money in implementation. Integrating services like data processing in the UME allows the mission to tap into the benefits of the cloud architecture, and eliminates the need for stand-alone, separate processing, and the resulting complexity of additional hardware/software to interface with, upkeep, and maintain security for. Further, once modules are developed, they can be reused in any UME, eliminating the need for mission-unique development efforts.

During operations, missions benefit from the significant simplification of the approach. For example, return data services in the legacy architecture requires that the mission make changes by interacting with an operator at the ground site or through an operator at the SNOCC. Instead, in the cloud architecture, the UMGS’s MOC sends directives to the relay module, which are in turn routed to the proper SLP monitoring and control (M&C) ground station system, servicing the mission. The SLP M&C module then communicates with the receivers at the ground station to enact the change.

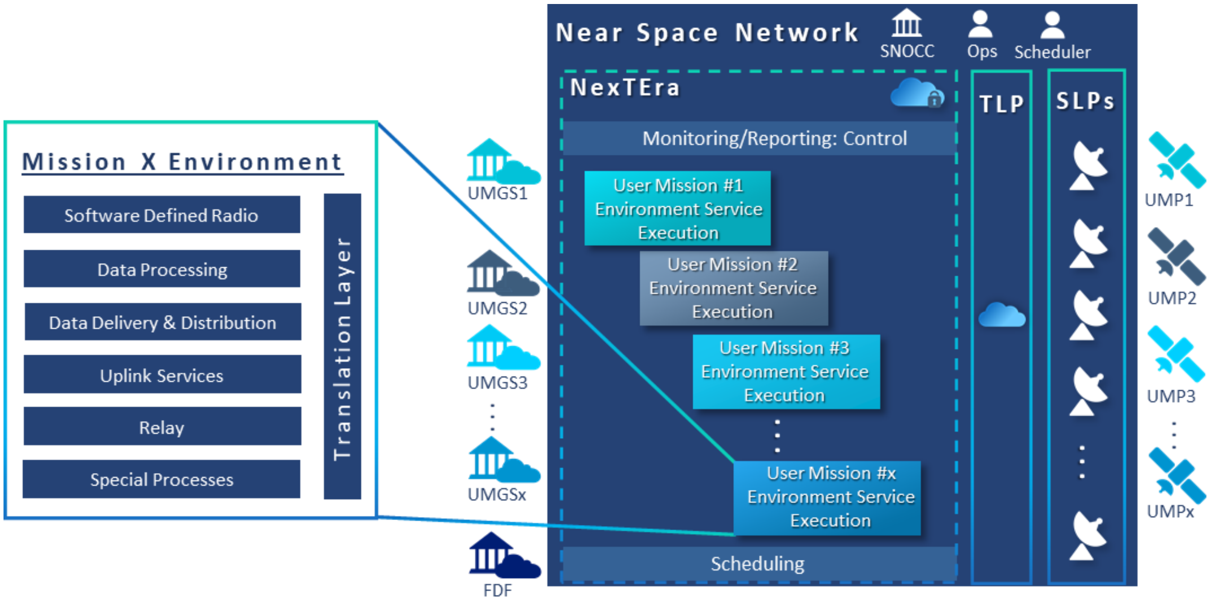

Figure 2: SCaN’s First Steps in Cloud Implementation – Near Space Network Architecture and Modular Mission Environment

3.1.2 Scheduling and Monitoring – Addressing Reliability and User Confidence

Missions desire ease both ease of operation and reliability of service; it is important to have access to a network when needed, and to reliably transmit, receive, and process data. The scheduling module addresses a portion of this need – simplifying a complex process. The scheduling module functions include receiving schedule requests, generate a deconflicted schedule, and sending schedule requests to providers. MOCs and SLP scheduling systems at each ground station connect to the scheduling module. The user views the availability of applicable resources and then selects the necessary timeslot for their mission. The scheduling module algorithm evaluates available slots, selects service times based on available prioritization data, and distributes schedules to the MOCs.

Network monitoring, and the ability respond to changes in the network, are also essential to building user confidence in the reliability of operations. The NexTEra Monitoring/Control/Reporting module interfaces with the User Missions, the SLPs, the Scheduling module, and all the mission environments. The monitoring function is passive; it monitors the status which is pushed from each UME. Similarly, all the independent connections from the functional modules within the SLPs report their respective statuses to M&C. Each mission will be provided with the status of their satellite separate and secure from other missions.

3.2. Modular Architecture Is Synchronous with Commercial Service Objectives

The proposed modular cloud architecture can be deployed using commercial capabilities and does not have to be built by NASA, much the way NASA currently procures commercially provided cloud data services. A variety of providers can participate, providing different modules and/or competing to provide modules. One can imagine a future state in which individual missions could select different vendor-provided modules for use within their UME. Again, the modularity provides independence between missions—a failure of an individual module does not impact other modules or missions. In short, there is significant opportunity moving forward for the commercial sector, to build these modules that are data-driven and cloud-based.

4. SCaN Cloud Services for Human Space Flight

NASA has embarked on the Artemis Program, an initiative to return humans—including the first woman and first person of color—to the moon. Artemis is significantly distinct from Apollo. This new campaign is targeting both a build-up to a sustained operational presence and providing the foundation for future human missions to Mars. Further, the approach to Artemis relies not just on NASA, but also incorporates commercial industry and international partners. SCaN is committed to supporting Artemis, bringing together the right resources and creating an architecture that meets the intense demands of this new phase of exploration. Current architectural plans are described in a separate paper [2] and include new Lunar Exploration Ground Segment (LEGS) stations and Lunar Relay capabilities in addition to the resources of the DSN. The first three LEGS sites will be government-owned/contractor-operated and placed in the following distributed locations to provide full coverage of the lunar region – White Sands, New Mexico in the U.S., and to-be-determined locations in South Africa and Australia. Additional LEGS services are intended to be procured commercially, as are the lunar relay services. As presence on the moon grows, NASA anticipates multiple vehicles in orbit and at diverse locations on the surface of the moon – both robotic and human – taking on increasingly complex operations. In this context, the cloud architecture will provide connectivity from these diverse locations and different providers, reliably returning data to the SCaN cloud, as shown in Figure 3.

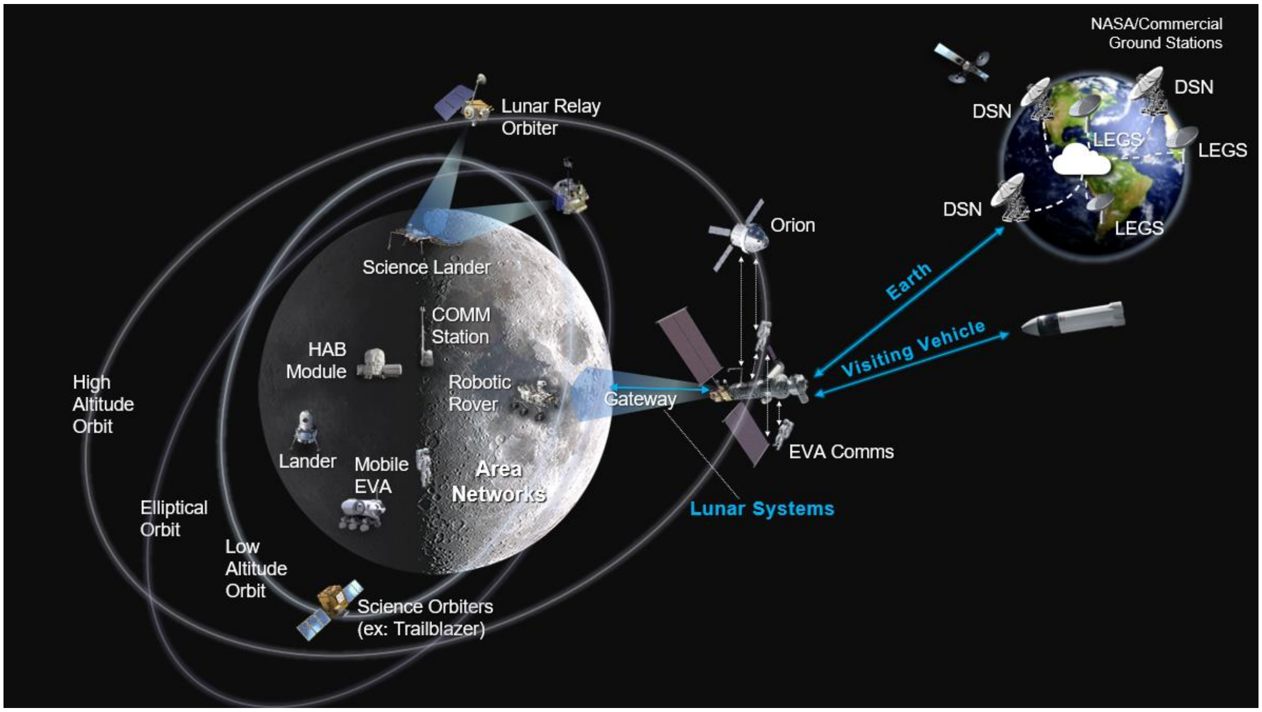

Figure 3: Artemis Conceptual Communications View

From a Human Space Flight (HSF) perspective, this approach is both simple and resilient; it provides the HSF mission elements, like the Orion vehicle, with one logical/digital connection to the cloud, which in turn provides them access to services being provisioned through lunar relay, ground sites, or other NASA assets acting as relay—like the Gateway platform in Near-Rectilinear Halo Orbit (NRHO) around the moon. Services from different locations may be run in parallel when (1) the operations concept for HSF requires it and (2) a best (network asset) source selection can be made. The best source select can be performed from a frame level (best frame select) or at the SLP level (i.e., which SLP to select). The cloud approach provides greater overall reliability, which is key to supporting HSF. The cloud architecture facilitates real-time processing, which is essential to HSF operations, and provides SCaN with the ability to retain and share data securely with the HSF community. This new architecture simplifies the HSF network concept of operation and enables the HSF central command to focus on managing the missions and taking care of the astronauts' safety.

4.1. Unique HSF Considerations

The NexTEra system is a first step in the implementation of the SCaN cloud architecture, however, in order to meet the needs of HSF additional capabilities and enhancements will be needed. There are significant differences in the support of a robotic mission and that of an HSF mission. While scientific robotic missions meet key Agency objectives, unexpected outages can be managed. This may result in the pausing science collection, storing the data locally, or in the case of Voyager 2 (when the only antenna to uplink was down for many months), put into a low data mode with extended communication outages. Such mission management for HSF is not desired and can easily put astronaut safety at risk. Although the Artemis Program will test out astronaut autonomy in preparation for Mars, unplanned and long duration outage pose a great risk.

4.2. Architecting for the Reliability Needs of HSF

Inherently cloud services are highly reliable [3]~~,~~ however, outages are not uncommon. There have been several very public and non-public outages of cloud services [4]. As we move to more commercial services, NASA is inherently losing some level of control over operations relative to the current architecture. Obviously, this transition to a commercial services model will allow for more cost savings and efficiency, but it requires a closer look at how NASA, and in this case SCaN, will manage failures and contingency modes in systems that we do not control. This is directly applicable to a commercial cloud service, where for missions that need higher certainty in terms of availability and reliability, additional capabilities will be required. This of course increases costs.

4.2.1 Regional Diversity and Backups

Cloud providers generally provide services from multiple locations both across the country and the world to ensure that there is not a single point of failure, and to ensure that one location is not overloaded with data. In the United States, for example, user data coming from a particular location can be routed to either the West Region (US West-1, based in California) or the East Region (US East-1, based in Northern Virginia). Both regions are considered "hot backups," as data can flow into either region at any time, with little to no latency. Each region has the capability to independently process data received. If an unexpected outage occurs in one region, data may be moved to the other "hot backup" to prevent a loss of data.

4.2.2 Notional Implementation for HSF / Artemis

Figure 4 shows a notional highly reliable architecture implementation for HSF, using representative Artemis projects/systems elements and service providers. This notional architecture provides redundancy in cloud service regions with hot backups, and also employs a secondary cloud service provider, all with the end goal of maintaining uptime for this critical mission set. The representative flight elements (Human Landing System (HLS), Orion, and Gateway) are shown in the upper-left section of the figure, connecting to a variety of ground stations (such as those of the DSN). In this depiction, AWS is the primary provider of cloud services, with two regional hot backups and associated support for each mission element including virtual private cloud, processing, and storage. Redundancy is provided through this regional, hot-back up approach, in which triggers such as unexpected downtime or scheduled maintenance result in a failover to the other region. Mission control is depicted at the bottom of Figure 4, as the consumer of all the data, regardless of which region receives/processes the data. For HSF, however, even this level of redundancy may not be sufficient to gain user confidence; control still rests with the (single) provider, not NASA.

4.2.2 The Cloud "Watchdog" – Continuous Monitoring Increases Resilience and Mission Confidence

To further mitigate risk and increase reliability, a cloud "watchdog" approach could be implemented – as shown notionally in the upper-right section of Figure 4. The watchdog function requires insight into all regions, and potentially alternate providers' and their associated regional nodes. Implemented/operated by NASA, the watchdog continuously monitors for outages and can make decisions to either fall back to degraded modes within the region, shift to a different region, or a different service provider—depicted in Figure 4 as Microsoft Azure on the right-hand side. Historically, and culturally, the HSF community has operated its network supports with significant human-in-the-loop activity – knowing people are in place to watch, observe, and act on behalf of astronaut safety has been a source of reassurance. For future HSF, with missions to the moon and Mars, this level of human involvement is not feasible given the size and scale of the disparate networks and providers involved. As such, implementing architectural features such as the cloud watchdog is critical to ensuring data availability, scalability, and mission confidence. [Note that depiction of AWS and Microsoft Azure are representative only and do not reflect or proscribe use of these vendors, or in this capacity, by NASA.]

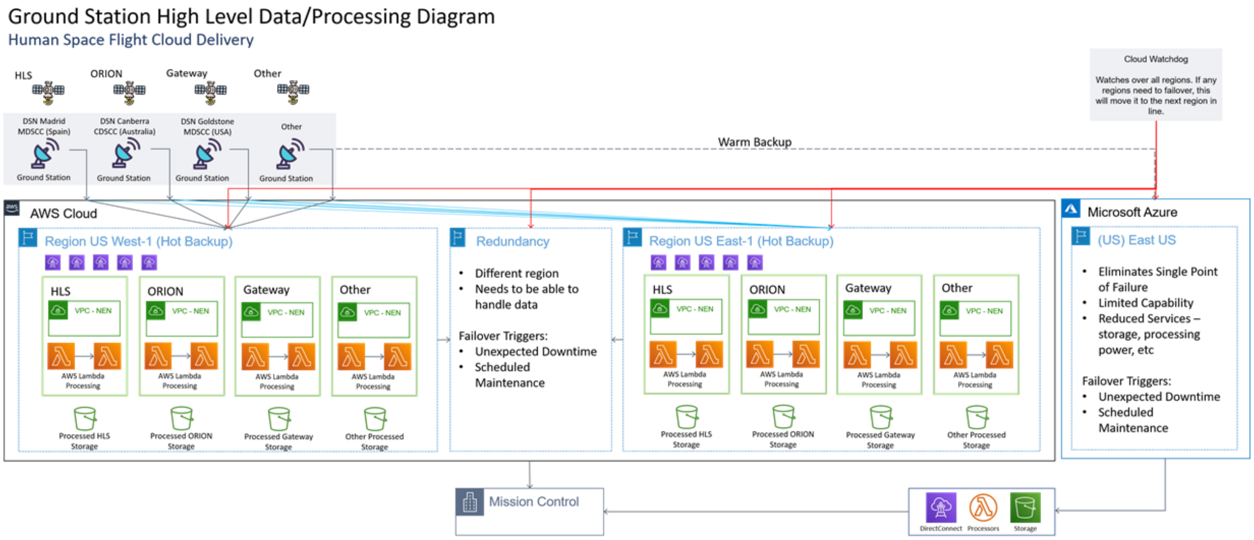

Figure 4: Notional Architecture for HSF Cloud Services

In short, SCaN's proposed cloud implementation will enable a single logical connection with access to network infrastructure, providing a secure, resilient, redundant, highly available, and economic solution. With such bi-directional access, security has to be at the forefront of the plans. The next section discusses the context of security considerations for cloud services.

5. Cloud Services Security Considerations

Security is an ever-present concern for NASA. Whereas the NASA Office of the Chief Information Officer (OCIO) is responsible for the Agency's IT infrastructure, SCaN is responsible for the security of space communications. SCaN's communication infrastructure, inclusive of partner networks, must provide high network security and resiliency.

An advantage of the cloud architecture is the ability to select providers who are Federal Risk and Authorization Management Program (FedRAMP) certified; a process managed by U.S. General Services Administration (GSA). Most cloud providers in the private sector have some level of FedRAMP certification. DAPHNE was implemented in the AWS cloud and NexTEra is continuing to evolve based on AWS capabilities. The current implementation of SCaN cloud services is categorized at the FedRAMP Moderate level but can be deployed at the High level in the future, which will be important with the expansion of critical and HSF missions in the lunar domain[‡]. With continued adoption of cloud services, NASA anticipates a more diverse set of providers will be available to provide FedRAMP compliant space communication services in the future.

6. NexTEra Current Status and Near-Term Next Steps

The DAPHNE cloud platform was deployed for two upcoming missions, NISAR and PACE. DAPHNE was first implemented as a component of the Near-Earth Network (NEN) Initiative for Ka-Band Advancement (NIKA) project, which added Ka-band communications services capabilities to NASA NSN ground stations at Fairbanks, Alaska; Svalbard, Norway; Punta Arenas, Chile; and Wallops Island, Virginia.

The DAPHNE deployment represents NexTEra V1.0. All basic elements of the architecture were implemented in V1.0 except for the uplink relay module. Forward work is ongoing, including implementation of scheduling capability, user mission ground segment transition, and associated translation layer development. The NSN will work with missions on an individual basis to: (a) identify the data needs and processing requirements that will define their UME modules; (2) work through the termination of the direct connections to specific SLPs; (3) remove any mission unique equipment at ground sites; and (4) establish the single logical/digital cloud connection.

[‡] AWS does have a FISMA high level that can be leveraged by SCaN when needed.

8. Conclusion

Beginning with the DAPHNE project, SCaN completed significant work to define and implement a cloud-based architecture for NASA missions in a seamless, resilient, and cost-effective manner. NexTEra V1.0 is ready to provide capability for the first two new science missions, NISAR and PACE.

SCaN is working with both the NSN and DSN to mature the broader cloud architecture approach that encompasses the data delivery needs of both networks, and critically integrates design elements like the hot-backups, watchdog functionality, and common interfaces that support the unique needs of HSF, as well as the connectivity between lunar assets and the array of NASA, partner, and commercial network resources. This will streamline the architecture, providing a more efficient and reliable future for communications services for Artemis and beyond. As the cloud architecture matures, commercial industry capability and innovation will be sought out and leveraged, consistent with Agency and SCaN objectives. The architecture, and the mission users, will benefit from a diversity of competitive suppliers that can provide functional modules and services within the cloud architecture.

References

[1] National Space Policy of the United States of America. 09 December 2020, https://www.space.commerce.gov/policy/national-space-policy/ (accessed 06.09.2022).
[2] P.J. Baldwin, G.W. Heckler, J. Armitage, A.J. Petro, J.W. Mitchell, W.C. Evans, E.L Weir, P.W. Wetherbee, L.M. Toney, NASA's Communications and Navigation Architecture Plans to Support the Return to the Moon and a Sustainable Lunar Presence, SpaceOps-2023, ID # 537
[3] R. Bala, D. Smith, K. Ji, D. Wright, M. Angel Borrega, Magic Quadrant for Cloud Infrastructure and Platform Services, 19 October 2022, https://www.gartner.com/doc/reprints?id=1-2AOZQAQL&ct=220728&st=sb&refid=fc81dabe-57e1-4c46-8d33-cfd3acf1ef08 (accessed 19.01.23)
[4] A. Bednarz, Top 10 Outages of 2022, 13 January 2023, https://www.networkworld.com/article/3685490/top-10-outages-of-2022.html (accessed 19.01.23)