Onboard Processing for LunaNet Data Services

Jonathan Verville
NASA Goddard Space Flight Center
8800 Greenbelt Rd,
Greenbelt, MD 20771
jonathan.p.verville@nasa.gov

Wesley Eddy
MTI Systems, Inc.
25111 Country Club Blvd, Suite 295
North Olmsted, OH 44070
wes@mti-systems.com

Abstract—**The LunaNet Interoperability Specification (LNIS) has been created and released to the public, enabling international commercial and government lunar systems to have a common baseline to work together in forming complex lunar mission networks. The LNIS specifically defines a set of data service protocols that offer multiple options for different types of network users and providers to work together. This includes real-time frame delivery services similar to those classically used by missions, as well as both real-time and store-and-forward networking that will be important as the number of mission users and complexity of lunar mission operations increases. The NASA Lunar Communications Relay and Navigation Systems (LCRNS) project requires lunar relay satellite services to implement onboard processing and networking capabilities in some ways similar to LEO mega-constellations that offer Internet access, but also going beyond those capabilities to meet unique lunar mission needs (e.g. store-and-forward services, LNIS messaging, etc.). Onboard processing will be necessary for a number of different protocols that are included in the LNIS, including support for data transfer over lunar proximity radio links based on CCSDS framing and CCSDS Encapsulation Packets, networking via IP and DTN Bundle Protocol, a suite of DTN Convergence Layer Adapters, and LunaNet messaging services. This paper provides an overview of our work in contributing to and defining the LNIS service interfaces and LCRNS requirements related to onboard processing for data services. This includes exploration of the key aspects of lunar networking concept of operations, technology trade studies, protocols for trunking, applicability and scope of different protocols, messaging, security considerations, and realistic implementation and deployment constraints. While the technology for new lunar relay networks can build upon the recent advances in LEO constellations, there are unique new challenges related to network management, onboard processing hardware/software system constraints, security, and quality of service for human spaceflight missions. As multiple international systems are planned to be put online in the coming years based on the LNIS and derived data services, this work will have a lasting impact that supports the evolution from single-vehicle user missions into scenarios where many surface and orbital assets are able to collaborate over multi-hop networks operated by diverse providers.**

TABLE OF CONTENTS

1. INTRODUCTION

NASA's Goddard Space Flight Center (GSFC) created the LunaNet architecture, described in 2020 as an integrated architecture to support networking, positioning, navigation and timing (PNT) and science services at the Moon [1]. In subsequent years, the LunaNet architecture spread within NASA, with SCaN and the Moon to Mars Office's adoption of LunaNet. ESA (in 2022) and JAXA (in 2023) subsequently have established partnership agreements with NASA regarding lunar communication and navigation, including LunaNet collaboration.

One of the key objectives for LunaNet is to enable robotic landers, rovers and astronauts on the Moon to have network access similarly easy to use as we are accustomed to with the Internet and wireless access networks on Earth. This includes being able to support a number of different types of data flows, and to automatically route data to different possible destinations across the Earth and Moon. For instance, rovers analyzing samples should be able to send science data through potentially several different relays orbiting the Moon, which can then route that data back to Earth. Meanwhile, astronauts on the lunar surface should also eventually be able to receive real-time alerts (e.g. for harmful space weather events), as well as communicate with each other and Earth via live voice and video streams. Each communications link may offer connectivity to the larger network, allowing data transfers with any other assets on the network.

To make this possible, LunaNet will be based on a mixture of both well-established and newer space communications and networking protocols and standards. Similar to the early Internet, even though the network will start small, since it will provide useful services, there will be incentives for both users and providers to join the network over time. This is designed to facilitate rapid expansion of the network and network-based capabilities at the Moon. The LunaNet basis on open specifications and framework of open standards will allow industry, academia, and international partners to build and operate LunaNet nodes alongside NASA.

The LunaNet Interoperability Specification (LNIS) was created initially by NASA and released to the public as a draft revision 4 (LNISv4) in 2022 [2]. Subsequently, an update, LNISv5 has also been released with ESA and JAXA collaboration, and further detail is already being defined for a future LNISv6 release. The LNIS provides the necessary information for both LunaNet users and service providers to design compatible systems that will be able to participate in the LunaNet network.

The NASA Lunar Communications Relay and Navigation System (LCRNS) is creating a commercially owned and commercially operated lunar satellite relay network, as part of the Near Space Network (NSN) Services (NSNS) contract. Awardees of this contract will be established as LunaNet Service Providers (LNSPs). Under contract from the NSN, these LNSPs will provide LunaNet access to lunar missions over proximity space links. The required capabilities emphasize onboard processing beyond what has been typical for NASA space communications services in the past. The relay satellites will be able to support a taxonomy of services, shown below in Figure 1, that includes data services covering:

1. A real-time frame service, relaying Advanced Orbiting Systems (AOS) data link frames across contiguous sets of links, without adding significant delays. This allows user missions to continue to operate in the traditional fashion where space link frames are exchanged with a communications service provider.

2. A real-time network service, relaying Internet Protocol (IP) packets between a contiguous set of links, without adding significant delays. This allows user missions to interconnect more richly in a network of multiple data flows not just between a single spacecraft and mission control center. This allows more complex operations concepts to be realized, such as those envisioned for "Moon to Mars" architecture [3][4], and support for live voice and video flows with astronauts.

3. A store-and-forward network service, using Delay Tolerant Networking (DTN) Bundle Protocol (BP) and leveraging onboard storage capabilities to provide reliable routed data services, even when contiguous paths end-to-end across the network are not feasible or the end-to-end path includes disruptions or significant delays.

Supporting this set of services implies implementation of several additional protocols beyond just AOS, IP, and BP that must be processed onboard, including CCSDS

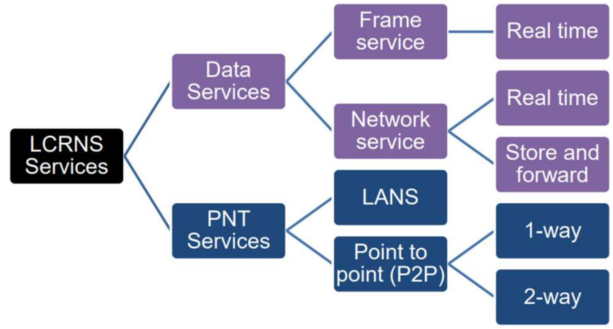

Figure 1 LCRNS Service Taxonomy

Encapsulation Packets and DTN convergence layer protocols. The LCRNS requirements imply that the protocols must be able to support throughput roughly in the 100 Mbps range on a single node (dependent on constellation design). The level of onboard processing will be increased well beyond that for other past NASA operational communication relay systems.

The remainder of this paper describes the LCRNS on-orbit relay functionality in more detail, based on the key system requirements, and the focus on supporting NASA's human spaceflight program as well as multiple uncrewed/autonomously operating missions simultaneously at any points within lunar orbits or on the lunar surface. Section 2 explains some of the historic background that influences LunaNet and our perspectives for LNSP onboard processing. Section 3 provides an overview of the relevant aspects of lunar relay operations, including the similarities and differences from past systems and the key challenges facing LCRNS related to onboard data processing. Based on the required network services, Section 4 describes aspects of the envisioned lunar relay network operations, especially as these relate to operation of the onboard processing systems and protocol implementations. Section 5 includes a deeper investigation of design space options for onboard data processing, focusing on topics such as the hardware/software design options, multiplexing of different user data flows or services within a provider network, store-and-forward aspects, and provider-internal protocol considerations. Section 6 clarifies the benefits of network services and LNSP onboard processing for the needs of the Artemis program. Finally, Section 7 summarizes conclusions and identifies the important future work in this area.

2. BACKGROUND

The LunaNet concept is substantially different from the way that space communications systems have been built and operated in the past, however LunaNet architecture has basis in the history of a number of efforts over the years to bring modern networking technology to bear for space missions. Traditionally, space missions have been able to operate in a mostly point-to-point fashion with communications providers like the NSN offering datalink services (at the frame layer or below) point-to-point between an operations center and the mission spacecraft. Some of the key influences and past developments showing viability of IP and DTN user services, integrating with classical space communications systems and development of onboard network protocol processing include:

- Early demonstrations of IP-based communications over classical space communications services, such as the 2002 CANDOS experiments on the STS-107 Space Shuttle mission using Mobile IP and IP-based application data flows, transitioning over multiple link service providers [5] as well as other demonstrations within the Operating Missions as Nodes on the Internet (OMNI) project scope [6].

- Development of the Space Network (SN) IP Services (SNIS) architecture in 2005, integrating IP networking support to mission users over GEO relay services [7] through NASA's Tracking and Data Relay Satellite System (TDRSS).

- In 2006, NASA's Space Communications Architecture Working Group (SCAWG) considered the future needs for space communications and networking, particularly to support complex human missions. A data networking architecture was produced [8] that included new support for networking protocols, and the key concept of "onramps" to simultaneously offer different varieties of services at multiple layers (e.g. frame, IP packet, and DTN bundle) to different users.

- Adaptation of a Cisco commercial router for use in the Low Earth Orbit environment, integrated onboard the SSTL UK-DMC satellite starting in 2003, and used subsequently to support experiments and demonstrations with IP, Mobile IP, IPv6, IPsec, and the first DTN BP on-orbit experiments [9].

- Integration of IP and DTN BP routing capabilities with three software defined radios attached to the ISS as part of the SCaN Testbed in 2015, including implementation of IP over CCSDS encapsulation, and multiple DTN convergence layer adapters, DTN custody transfer, and CCSDS File Delivery Protocol (CFDP) using DTN [11], with later experiments including DTN network management and DTN security.

- Operation of IP, DTN, and WiFi [12][13][14] technologies on board the International Space Station has informed NASA of the challenges and operational factors for the infusion of these networking features into a space platform.

NASA has sufficient experience with networking tests, demonstrations, and experiments between these and many other prior efforts in order to confidently move into an era where advanced networking becomes an operational capability at the core of an integrated space communications network, rather than confined to temporary experiments or isolated pockets.

3. LunaNet Network Services

Artemis mission elements are designed to rely heavily on onboard IP networks and carry complex sets of simultaneous data flows over their onboard networks. The number of mission elements and complexity of the distributed computing systems throughout the vehicles are driving needs for networked space communications services rather than simple point-to-point services, as multiple elements proceed to be deployed and build out the Moon to Mars architecture [3].

LunaNet follows some design patterns that are similar to aspects of the Internet:

1. Use of open protocol specifications to define the interfaces, both between users and providers, and provider-to-provider. The hope is that this will foster strong interoperability between users and multiple providers that can be selected from.

2. Incentivizing commercial operators to provide network services. LCRNS, for instance, is procuring commercially provided services rather than deploying government-owned/operated systems. Diverse commercial systems can be a part of LunaNet, but also within their networks, providers are free to do things that are proprietary or exclusive to their own implementations.

3. International compatibility, not limited only to mission elements from one country or space agency. While international space agencies have used cross-support services in the past, LunaNet aims to make this support more routine and ubiquitous for lunar missions. Recognizing Metcalfe's Law[1], this will increase the value of the lunar network and the information, science,

[1] The value of a network is proportional to the square of the number of users connected to it.

and other services that can be offered over the top of it.

4. Allowing user choice between multiple protocols and services. Missions will be able to select the type of services (e.g. real-time or store-and-forward) and interface options (e.g. S-band or Ka-band) that are appropriate for their needs. Because the network services provide generic packet and bundle transfer, users will be able to implement their own services over the top of the connectivity provided by LNSPs, similar to the "permissionless innovation" of the Internet [15], which is the ability for users of the network to deploy new types of services and applications making use of the underlying network connectivity without any further coordination or impact to the network architecture.

5. Incorporation of both local and wide-area networking. LunaNet includes a focus on proximity links between relay satellites and the lunar surface, but also includes direct with earth (DWE) communications, and plans for the eventual integration of surface-to-surface wireless networks.

Although NASA has significant relay operations experience from the Tracking and Data Relay Satellite System (TDRSS) acting as a key part of the NSN for decades, the LCRNS relay requirements are considerably different from TDRSS. Unlike a traditional bent-pipe / transparent relay satellite system, LCRNS satellites must be able to modulate/demodulate user space link signals, process the link layer and higher-layer protocols, and make onboard routing decisions based on the contents of frames, packets, or bundles and the relay service configuration.

Mars relay services have also been operated for several years, with some similarities and differences related to LCRNS. Similar to LCRNS, Mars has multiple relay spacecraft that provide services where DWE communications are not feasible due to multiple factors (e.g. link budget, DWE resource loading, etc.). Mars relay services were built organically across multiple missions and organizations over time, with most of the interoperability only at the physical or basic link layer, which limits the onboard processing needed. LCRNS will extend that interoperability to the network layer, and standardize the store-and-forward networking capability and interfaces. LCRNS will differ from Mars relay operations because the onboard processing systems will operate at a higher layer and be able to form a routed network.

Some of the key LCRNS requirements that imply significant onboard data processing capabilities include:

- Real-time data relay services between terrestrial and lunar users, needed to support crewed missions with live voice and video, among other real-time data flows.
- Non-real-time data relay services in order to support missions that can only be served from the far-side of the moon, requiring significant onboard data storage.
- Ability to support multiple concurrent services in parallel, due to the expected density of users at the lunar south pole region, for instance.
- Simultaneous operation of S-band and Ka-band services (along with other PNT services), implying the need to properly route user data to and between the proximity data links and any relay DWE or intersatellite relay crosslinks (either within the same LNSP constellation or between LNSPs). Data flows may also be a mixture of real-time and store-and-forward, complicating the routing determination between currently available links or storage for non-real-time data.
- Measuring and reporting quality of service (QoS) metrics for user data traffic across many simultaneous data streams.
- A diversity of user mission configurations that relays must be able to manage and configure for over time, spanning from the link through to network layers. Unlike terrestrial wireless, there are no initial defaults or ways to automatically negotiate link parameters with the standard proximity space link protocols.
- Lunar surface systems, such as the Human Landing System (HLS) or Commercial Lunar Payload Services (CLPS) landed elements plan to serve as surface access points (APs) or base stations. This means that they may be aggregating IP traffic for multiple end systems and routing through return links such as LCRNS relay satellites. The NASA HLS program has baselined requirements for this function.

Due to these requirements, there are several unique capabilities that the LCRNS relay satellites will include, representing innovations in onboard processing that go beyond traditional space mission communications services.

Multiplexing of data on the DWE links will be needed in order to support the multiple concurrent services and simultaneous set of different services that must be supported. While frequency division multiplexing has been common for Earth-based relays (e.g. on TDRSS space-ground links), for LunaNet, the LNSPs are expected to need to multiplex user traffic at higher layers (above the frequency domain) such as through different types of encapsulation. **Prioritization** of user traffic should be performed while routing and processing it, distinguishing between traffic types (e.g. science data return versus

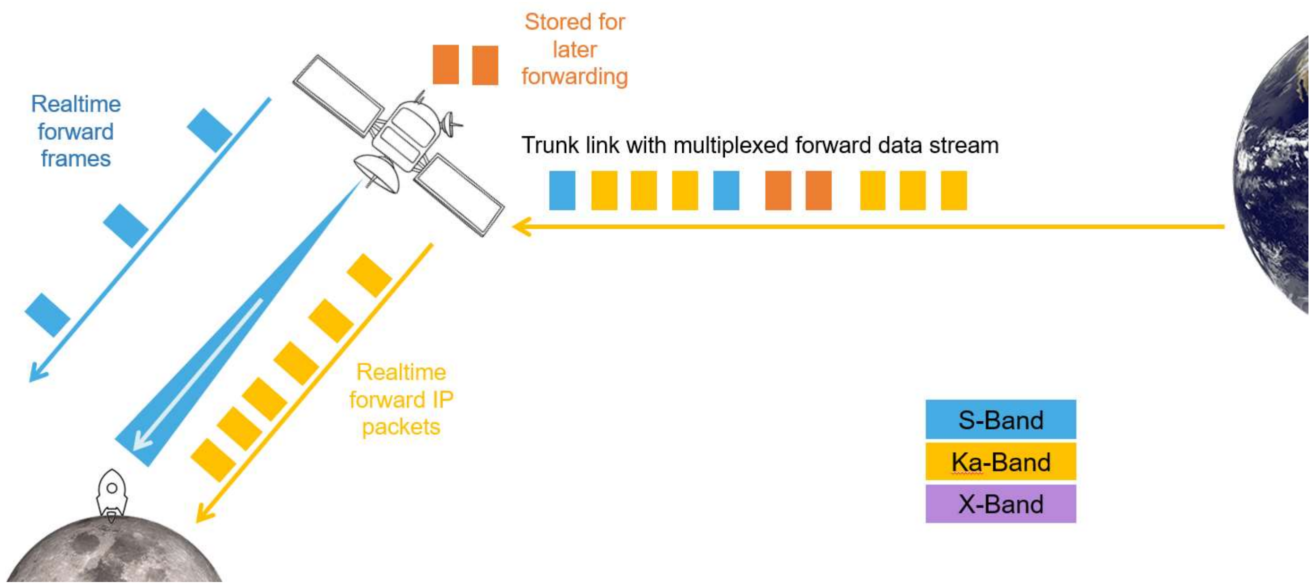

Figure 2 Basic Trunk Multiplexing

critical or emergency traffic) and service interfaces (real-time and non-real-time store-and-forward data with looser latency constraints). Onboard **scheduling** and autonomous initiation of stored data transmission, and updates to routing configurations, according to the planned contacts and arranged services. **Management capabilities** that support the LNSP coordination of protocol stack details and configurations over time to support differing users and services.

4. LUNAR RELAY OPERATIONS

The role of LCRNS is specifically defined within NASA's Moon to Mars architecture, including how LCRNS works as part of a much larger network including other relays such as the Lunar Gateway, and terrestrial DSN sites, NSN ground stations, ESA ground stations, and other provider infrastructure [3] [4].

Serving NASA's Artemis mission users is a primary goal for LCRNS relay services, but the requirements recognize that there may be other lunar mission users, and that the LNSP operating the relays may have its own additional customers or other missions for the relay satellites to perform at times. The relay satellites are not necessarily dedicated, and their use must be coordinated with the LNSP.

Coordination with the LNSP is expected to include traditional exchanges of information needed for service acquisition such as orbit data for antenna pointing & tracking, RF carrier and modulation parameters, etc. New information, relevant to network data services will also need to be coordinated, in order to successfully route data as required. This may include:

- IP address and prefix information pertinent to the user mission and its data flows.
- DTN Endpoint ID information for user mission and the other endpoints of its data flows.
- DTN convergence layer adapter configuration data.
- Agreements on policies for storage of DTN bundle data between or across contacts, since data may need to accumulate in the forward direction well in advance of scheduled contacts between the LNSP and user mission. The LNSP may need to reserve onboard storage space separately for different users or applications that it is performing on each relay spacecraft.
- Agreement on quality of service treatment, e.g. if there should be some priority scheme, e.g. determining transfer order amongst forward data flows towards the user spacecraft.
- Potentially, indications of limitations in the mission's supported packet or bundle sizes, etc., or expression of other mission-unique capability limits.

The LunaNet architecture is intended to enable multiple LNSPs to be simultaneously operating, and they may even support one another in providing reach-back through shared trunk links with Earth, for instance, in cases where one LNSP has a relay that provides local access at the moon, and another LNSP connects a second relay to the first and provides DWE services to it. In this case, LNSPs need to have means of coordinating their operations and

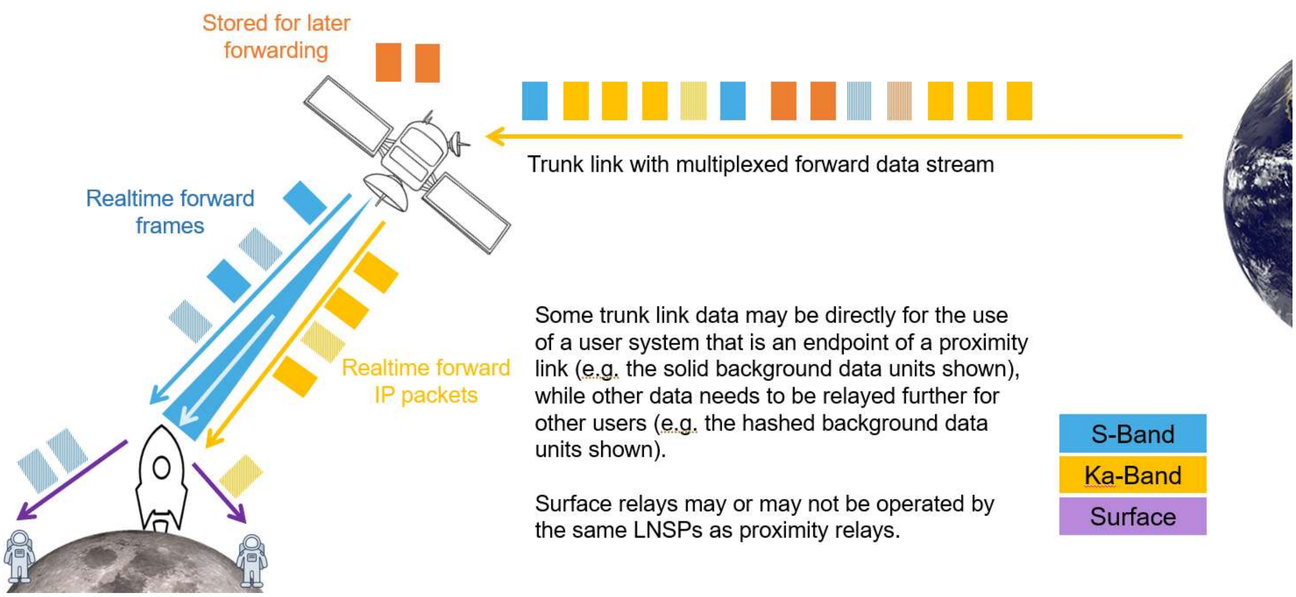

Figure 3 Multiplexing Including Multihop Routed Flows for Surface Links

managing how they will support one another, including any network routing decisions that may be required.

There are lessons that can be applied from experience with the Mars Relay Network (MRN) [16], even though the way that it technically operates has several significant technical differences (e.g. there is no network layer routing unifying the network, as envisioned for LunaNet). Some particular aspects that need to be addressed still for LunaNet operations include:

- Like the MRN "office" coordination, significant coordination may be needed between relevant LunaNet groups (users, LNSPs, etc.) that are sharing the early network and its initially limited resources. Since organizations have different agendas and priorities, a centralized body may need to be created to facilitate this.

- Similar to the Mars Relay Operations Service (MaROS) capabilities, tooling may need to be developed for LunaNet, to help support planning, situational awareness, and other functions, especially as configurations of routing and lower layers need to be orchestrated both within and between different agencies and companies.

- Similar to the MRN (and the Internet), LunaNet can embrace the "organic architecture" concept, where operators (LNSPs) and mission users each have significant autonomy, and some aspects of the network may come about over time without being explicitly architected from the top-down.

Since LunaNet offers network-layer interfaces, however, there are some new aspects that LNSPs and users must be aware of, and design for.

1. LNSPs must each have designs to manage and orchestrate the routing configurations within their networks with their users and likely with other LNSPs that they and their users do business with. For lunar use of IP, and for DTN, there are no analogous systems to the Internet's Border Gateway Protocol (BGP), and its concepts of Autonomous Systems , the Default Free Zone, etc. LNSPs will need to adapt or develop new paradigms for how to exchange and maintain routing information with one another (and it must be secure and reliable, since disruption of lunar assets may have high value as a target for terrorists, nation-state adversaries, and other threat actors). Early-on, routing data coordination and joint planning might be handled largely though terrestrial systems that manage each LNSP's space segment, and later on transfer to more autonomous in-space interactions as the network scales up.

2. LunaNet is intended to be Internet-like, but not to provide direct connectivity to the Internet. However, the network protocols (IP and DTN) are not limited to the strict scope of a particular mission, LNSP, or LunaNet, but *could* be more widely routable. Some nodes might act as gateways to other networks, and this could happen either purposefully by design, accidentally by misconfiguration, or maliciously as the result of an attack, backdoor, etc. Systems using LunaNet should not assume that it is a closed network or

walled garden, even though it may practically be so early in its lifetime.

Security of the lunar relay system is of paramount importance and can be separated into LNSP and end-user responsibilities. This is analogous to the way the terrestrial Internet works today with Internet service providers and subscribers. Similar to cloud FedRAMP services or Software as a Service (SaaS) applications, there is an implied shared security responsibility for LunaNet. LNSP ground systems and spacecraft will be secured, including telemetry command and control encryption and anti-interference mechanisms required for risk reduction. Mission end users are ultimately responsible for security of their data transiting both terrestrial and space networks and should employ strong encryption (both in-transit and at-rest) and appropriate authentication for assets and software operating in cislunar space.

A set of assessment criteria similar to or tailored for lunar use from the Space Force Infrastructure Asset Pre-Approval (IA-Pre) or other NASA assessments may be a longer-term need, if a rich ecosystem of providers and connectivity emerges at the moon. Unique aspects that individual LNSPs would need to consider in order to balance different types of risks might include the practices for hardening and managing the security configuration of their relay payloads, separation of service data ("user plane") and relay operations ("control plane") data within their network, and vulnerability tracking and mitigation approaches for their in-space assets.

5. ONBOARD PROCESSING DESIGN SPACE

The LCRNS team has studied the potential space of onboard processing design options that LNSPs might select within their designs in order to meet the LCRNS requirements. This section provides an overview of some of the results.

The LCRNS requirements include some aspects that drive the overall scale and performance needs of an onboard processing system, including Ka-band data rates up to 50 Mbps, simultaneous operation of multiple S-band, Ka-band, and PNT services, as well as DWE links, and ample onboard data storage. The team studied the capabilities of onboard processing hardware that would be suitable for the lunar orbit environment, as well as for integration into a payload that could fit into a typical package for launch and transfer to lunar orbit.

Given the needs for multiplexing of multiple user data flows, as well as command and telemetry for the relay spacecraft constellation itself, one of the key aspects studied included potential methods to perform trunking[2] on DWE links, and the related implications of protocol processing onboard.

Figure 2 illustrates the basic needs for a trunking protocol mechanism, that can support mixture of forward data on a DWE link. On receiving each unit of data, the relay must be able to recognize if it is intended to be transmitted (for instance) as a frame on an S-band link (blue line), or if it corresponds to messages for the AFS broadcast (blue triangle), or IP packets that must be encapsulated and put into frames for a Ka-band user link (yellow line), or if it may be DTN data that should be stored for later transmission (colored orange).

The situation becomes even more complex when multi-hop relaying is considered, rather than just a single DWE trunk and proximity link. Two cases of this are:

- There may be surface wireless users that do not have sufficient power to close proximity links to orbital relays, or do not support the right protocols for these links (e.g. they may only support WiFi or 3GPP protocols), as shown in Figure 3. Base stations on the surface may relay packets or DTN bundles between these surface nodes and proximity relay spacecraft, adding an additional hop. These base stations are further beneficial because they allow a greater number of surface user nodes to be served through fan out/in on the surface, while keeping the proximity relay capacity for simultaneous users small.
- There is also potential for an LNSP to have intersatellite links (ISLs) within the constellation, assuming some relays are reachable from Earth only via other relays over ISLs at certain points in time. In this case, the DWE trunk links may carry data that needs to be identifiable for routing to any potential ISL links, in addition to the other possibilities.

There is a plethora of different protocols that could be used to implement the needed multiplexing functionality on trunk links. The specific protocols that an LNSP chooses will be their own selection, as they are not visible outside the LNSP's system and LCRNS requirements do not specify these design details. The selection should agree well with each LNSP's particular onboard hardware/software switching and processing design, management capabilities, and methods of operating the ground system and flight payload.

[2] The term "trunking" is adapted here from common use in terrestrial networks, where trunk links aggregate data to and from multiple access links. Protocol capabilities keep the data from the different access links logically separated while in-transit on the trunk links and while being processed by a router or switch.

	Solution Approach	Main Advantages	Main Disadvantages
Data Link Layer	1.CCSDS virtual channel IDs	CCSDS standard	Management complexity
	2. CCSDS frame encap/slicing	CCSDS standard	Newer standard; management complexity
	3. ITU GFP-F	ITU standard, commonality w/ optical	New implementation for RF links
	4. VLAN tagging	IEEE standard	New implementation
Network Layer	5. IP/UDP tunnel encapsulation	Flexibility, generality	Increased overhead
	6. IP/IPsec tunnel encapsulation	Flexibility, security	Higher overhead and processing
	7. MPLS encapsulation	Flexibility	Not inter-provider
	8. DTN BP encapsulation	Unifies all forwarding under DTN	Highest overhead and complexity

Table 1 Trunking Protocol Options Considered

Despite this freedom for an LNSP to select its own trunking methods, the use of common protocols may enable greater efficiencies and interoperability, resulting in more overall robustness for the network and services, e.g. since it could facilitate easier cross-support between LNSPs and/or potential for layered transport networks over a myriad of systems in the lunar domain. Especially since the number of LNSPs will be small initially, while the specific trunking design is internal to each LNSP, and not evident to users, commonality might be encouraged to promote manageability, scalability, extensibility, and reuse of technology and hardware/software components between providers.

The key needs for any trunking protocol will be to:

1. Distinguish between multiple types of enclosed user service data requiring different treatment or processing by the relay. User data being relayed may be CCSDS frame-based, IP-based, or DTN BP-based.

2. Data may be intended for either (1) immediate real-time relay passthrough, (2) reception and processing by the current relay, or (3) forwarding for later processing by some other relay “downstream” over an ISL either in real-time or in the case of DTN data, at some future point in time.

3. Data intended for processing by a receiving relay may be a mixture of local payload command traffic, or user service data. Data passing through may be user service data, or telemetry and service monitoring data from other relay satellites.

These multiple streams of data need to be merged and split by relays within the required timeframes given the latency budgets in the LCRNS requirements and at the data rates and unit-per-second (in units of frames, packets, or bundles) that are implied by the maximum S and Ka-band data rates, along with the worst case (i.e. smallest) possible frame, packet, and bundle sizes.

Trunking protocols might be implemented in either onboard software (e.g. running on general purpose processing cores) or within specific hardware designs accelerating the particular protocol options. Combinations of hardware/software functionality are possible along the full spectrum in between the pure hardware and pure software approaches as well. The rows in Table 1 summarize several different approaches leveraging either data link or network layer protocols or mechanisms in order to support trunking. Physical layer approaches have also been considered, but in our studies have been found undesirable due to multiple disadvantages and limitations. The main advantage of a physical layer approach, such as frequency division multiplexing is that it can be very simple to implement on the terrestrial side. However, disadvantages include:

- It may be difficult to size each channel bandwidth appropriately in advance, especially given that there are limited DWE link spectrum resources that must be worked within.
- Physical layer multiplexing requires setup and management of many different channel and modem resources over time. This could be complex for onboard processing, where multiple parallel modems for different channels are uncommon for the radios used in DWE communication.
- Frequency division is not flexible for varying proximity link bandwidths, and variable data flow rates. Over short time scales, the balance of real-time and store-and-forward data sharing a trunk may be dynamic, for instance.
- The use of ISLs magnifies the challenges, as a relay constellation may grow to include several relay nodes and multiple ISL hops for data routing.

	Solution	Scalability	Security	Overhead	Market Availability	Complexity
Data Link Layer	CCSDS Virtual Channel IDs	< 6 bits	LNSPs may need to limit user VCID usage and filter on user VCIDs	Lowest overhead	Commonplace for CCSDS radios	Well-understood and widely-used paradigm, but may be difficult to manage in multilateral setting
Data Link Layer	CCSDS Frame Encapsulation/ Slicing	6 bits	N/A	Low overhead	Not yet widely used	Relies on radio having newer CCSDS slicing support
Data Link Layer	ITU GFP-F	N/A; must be combined with another approach	N/A	Dependent on the other options GFP-F is paired with in a complete solution	Not common for CCSDS space links	Simple configuration
Data Link Layer	VLAN Tagging	12 bits (and can be nested)	N/A	Medium or low overhead depending on whether Ethernet is already used on trunk.	Not common for NASA missions nor part of LNIS	Simple configuration
Network Mechanisms	IP/UDP Encapsulation	16 bits	N/A	Low header overhead and medium computational overhead	Commonly used terrestrially, and could layer to traverse IP LunaNet services	Simple configuration
Network Mechanisms	IP/IPsec ESP Encapsulation	32 bits	IP packets could be encrypted	Relatively high header overhead, and high computation overhead	Common terrestrially, but may not be part of many flight stacks	Needs more comprehensive key management plan, etc.
Network Mechanisms	MPLS Encapsulation	20 bits	N/A	Lower overhead than other network approaches	Commonly used terrestrially and by satellite operators, but may not be part of flight stacks	Requires some management not common for NASA missions
Network Mechanisms	DTN BP Encapsulation	Not practically bounded	BPsec could be used	High header overhead and relatively high computation overhead	Multiple software packages are available, but not mature commercial systems	Greatest configuration burden to be managed

Table 2 Stoplight Analysis of Trunking Methods Considered

In contrast, we found that the onboard processing needed for data link and network layer solutions (summarized in Table 1) is reasonable within the capabilities of modern flight hardware, at the data rates required for LCRNS.

Some of the viable approaches are based on specific capabilities of CCSDS space link standards. For instance, CCSDS Virtual Channel IDs can be leveraged to separate parallel flows. However, there are limited VCIDs available, and there would be management complexity in dynamically assigning their usage for particular types of flows on different trunk links within the network. CCSDS has also standardized means of frame encapsulation and slicing that could be employed for multiplexing on a trunk link, however it is a newer capability and could have similar management complexity to a VCID-based approach. An encapsulation of frames such as the ITU GFP-F for optical links could be used, however, it would need to adapted for RF links, and it would also be new hardware/software (depending on the implementation approach) specific to the lunar space use case, likely unable to leverage COTS. Similarly, VLAN tagging as used in Ethernet is well known from IEEE standards, however, usage on a lunar trunk link would be an entirely new adaptation and have similar management issues as use of CCSDS VCIDs.

Network layer options in general were found to be much more flexible, and to potentially scale better in complex and dynamic network situations. Some options tend towards more of a software-oriented implementation rather than simple hardware or FPGA-based designs, however, and so may have an impact on the onboard processing in terms of efficiency, power consumption, throughput, and latency.

Among network layer approaches, MPLS as often used terrestrially, including in satellite constellations, offers a

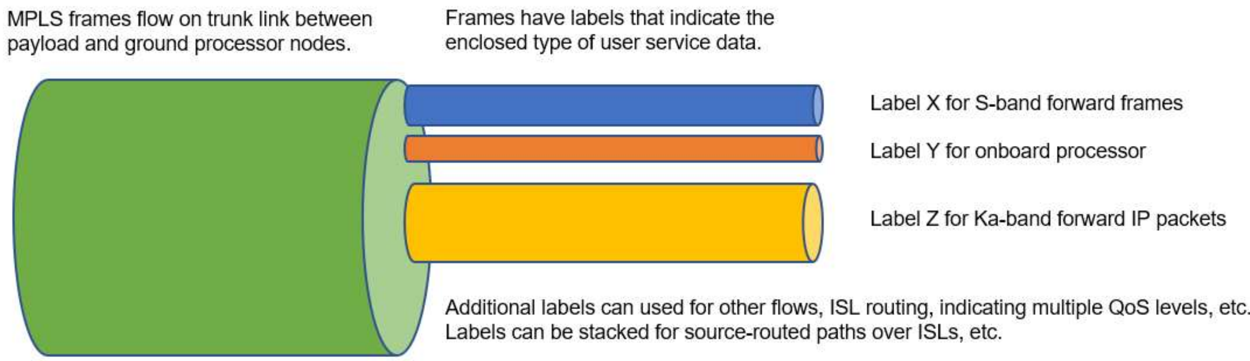

Figure 4 Example Multiplexing Using MPLS Encapsulation

very simple, efficient, yet powerful option that an LNSP might be able to leverage in conjunction with tools and management protocols re-usable from terrestrial networking (e.g. MPLS management data standards, MPLS ping, MPLS traceroute, etc.). For each of the options considered, we analyzed exactly how the protocol fields might work in order to support the needed degree of multiplexing. Figure 4 depicts a summary of how this would work for MPLS, for instance, with a (green) trunk link carrying MPLS frames with multiple different labels, indicating based on the label what services or other onboard processing each MPLS frame should be forwarded to by the onboard processing.

IPsec was also found to be an interesting option, as it may pair a flexible and powerful approach to security with the ability so support multiplexing of different flows on separate ESP Security Associations (SAs). The onboard processing for this may either be a concern or a strength, depending on whether there is the ability to leverage onboard hardware to accelerate and make the cryptography efficient, and whether it would be duplicative of other onboard crypto hardware (e.g. bulk encryption).

Leveraging DTN BP in order to carry all traffic was also an interesting option considered, as it could unify all of the routing and processing within the network, if other data types are encapsulated into bundles at the edge of the network. Since BP traffic already needs to be handled at the edges, the support could be extended naturally throughout the LNSP network. The header overhead and onboard processing capabilities are relatively higher than other approaches, however, and an implementation would require more complex software than other approaches considered.

Table 2 provides a summary "stoplight" analysis on the key properties we found for each of the options considered in terms of its scalability (number of bits to identify specific services or service levels, etc.), built-in security properties, overhead in terms of additional protocol headers needed, market availability, and complexity in terms of onboard processing burden (including onboard management capabilities needed).

6. ENABLING ARTEMIS OBJECTIVES

The increased need for onboard processing of networking protocols is driven by the Artemis mission plans and Moon to Mars objectives. Specifically, one of the top-level Moon to Mars sub-objectives for Communications, Position, Navigation, and Timing (CPNT) is for the architecture to scale to support long-term science, exploration, and industrial needs [4]. Without networking and the implied onboard processing capabilities, it would not be feasible to meet this scaling objective.

Artemis plans include orbiters, landers, multiple rovers (pressurized and unpressurized), a base camp, and astronaut suits, all with the need to function as nodes on the network with different data flows over time. If relay capabilities were limited to only bent-pipe operations, either at the RF or at a data link frame level, then the Artemis Mission Control (MCC-H) would need to arrange relay services on a point-to-point basis, and handle all of the networking required themselves, "over the top" of the relay services. This is viable at first for a small network, but is labor-intensive, hard to fully automate, and error-prone as the network grows in increased complexity, dynamism, and presence of diverse commercial and international partners.

Relays offering networked services with sophisticated onboard processing and storage, operating at the network layer, will instead enable simplified planning for the network, as well as more efficient routing paths, especially in cases where point-to-multipoint communication patterns might be desirable (e.g. voice communications, messaging, etc).

The nature of the lunar south pole mission plans, landing sites, terrain, and contact opportunities between the mission elements, relays, and Earth also drive the relay capabilities for onboard storage. Simulation studies by our team have shown a quantitative benefit to increase total user data return (e.g. for science) by using all available user-relay proximity contact opportunities, and onboard storage along with and intelligent scheduling of data transmissions between onboard storage and RF links. This is especially beneficial when DWE opportunities are limited.

7. Conclusions & Future Work

The LunaNet architecture and LNIS will soon start to be instantiated around the moon by LNSPs, including through NASA's LCRNS project.

In designing the data services definitions and interfaces for the LNIS, we strove to balance the specific limitations and concerns of space networking (low SWaP, specialized protocols, security considerations, etc.) with the goal of being able to offer an open Internet-like experience for users. Onboard processing needs for LunaNet are considerably more complex than what has been common to-date in space communications systems. Lessons are being incorporated from past space networking experiences, and operational systems, such as the other NSN elements, the Mars Relay Network, and earlier DTN and IP demonstration experiments and operational capabilities on the ISS and elsewhere.

The specific needs and challenges related to trunk link protocols (and ISL/crosslink protocols) as well as LNSP onboard processing needs were considered in detail during LCRNS formulation. Our study found that several diverse approaches to implementing trunk link functionality are technically viable, but have different operational, scalability, and security advantages and disadvantages.

As LNSPs begin to deploy systems and offer lunar data services, the effectiveness of different onboard processing designs and specific hardware/software combinations can be assessed in more detail. At first, LNSPs may field point designs that are tailored for the particular requirements, such as those for LCRNS, in terms of the data rates, storage, and user services that they support. Over time, a more rich services landscape may develop, as more capable flight computing systems become available, and more capabilities can be software-defined, or onboard processing can be leveraged to build more over-the-top services such as data caching, edge computing, etc. The LunaNet data services defined in LNIS should provide a solid basis for building such services and meeting mission needs.

Because of the complex and multi-layered nature of the LNIS definitions, a flexible software-defined testset is necessary to test and support lab validation of compatibility for LNSP systems. The LCRNS project is working to develop a testset capable of supporting the set of LNIS protocols and configuration options [17]. The testset is designed to support new needs for "full stack" compatibility testing that cover not just the RF and physical layer interfaces that are traditionally tested for space network compatibility, but also covering the onboard processing functionality including testing of frame relaying, IP, and DTN services. The software-defined nature of the test set allows it to support multi-layer tests, and an expansive number of permutations of physical, link layer, and network layer service configurations that are viable for user missions.

Ultimately, these enhanced networking capabilities provided by LunaNet compliant systems with onboard processing will allow for NASA's Artemis systems to be more connected and utilize enhanced capabilities to support the objectives of NASA's future Lunar exploration plans.

Acknowledgements

The authors are part of the LCRNS project at NASA Goddard, and the material in this paper includes contributions from the project Data Services team, including our colleagues Chris Mishaga, Shawn McGraw, Tiffany Navas, and Timothy Walker.

References

[1] D. J. Israel et al., "LunaNet: a Flexible and Extensible Lunar Exploration Communications and Navigation Infrastructure," *2020 IEEE Aerospace Conference*, Big Sky, MT, USA, 2020, pp. 1-14, doi: 10.1109/AERO47225.2020.9172509.

[2] LunaNet Interoperability Specification Document, version 4, LN-IS V004, September 12, 2022. https://www3.nasa.gov/sites/default/files/atoms/files/lunanet_interoperability_specification_version_4.pdf

[3] NASA, "Moon to Mars Architecture Definition Document", ESDMD-001 Revision A, NASA/TP–20230017458, March 2024.

[4] M. Zemba, K. Vaden, R. Reinhart, C. Gramling, and G. Heckler, "NASA's Lunar Communications and Navigation Architecture", 2023 Moon to Mars Architecture Concept Review, white paper, 2024. https://www.nasa.gov/wp-content/uploads/2024/01/lunar-communications-and-navigation-architecture.pdf?emrc=f1a91a

[5] D. J. Israel and D. Zillig, "CANDOS Experiment Overview", *Proceedings of the AIAA SpaceOps 2002 Conference*, October 2002. doi: 10.2514/6.2002-T3-26.

[6] D. Israel, R. Parise, K. Hogie, and E. Criscuolo, “Demonstration of Internet Technologies for Space Communications”, *Proceedings of the AIAA SpaceOps 2002 Conference*, October 2002. doi: 10.2514/6.2002-T5-18.

[7] D. J. Israel, “Space Network IP Services (SNIS): An Architecture for Supporting Low Earth Orbiting IP Satellite Missions”, *2005 IEEE International Conference on Networking, Sensing and Control*, Tucson, AZ, March 2005.

[8] D. J. Israel, A. Hooke, K. Freeman, and J. Rush, “The NASA Space Communications Data Networking Architecture”, *Proceedings of the AIAA SpaceOps 2006 Conference*, June 2006.

[9] L.Wood, W. Ivancic, W. Eddy, D. Stewart, J. Northam, and C. Jackson, “Investigating operation of the Internet in orbit: Five years of collaboration around CLEO”, IEEE Communications Society Satellite and Space Communications Technical Committee newsletter, vol. 18 no. 2, pp. 10-11, November 2008.

[11] D. Brooks, W. Eddy, G. Clark and S. Johnson, “In-Space Networking On NASA’s SCaN Testbed”, Proceedings of the 34th AIAA International Communications Satellite Systems Conference, Cleveland, Ohio, October 2016. doi: 10.2514/6.2016-5754.

[12] “Internet in Space: Bringing Wifi to Earth Orbit, the Moon, and Mars.” Accessed: September 23, 2024. https://broadbandnow.com/report/internet-wifi-service-in-space.

[13] A. Schlesinger, B. M. Willman, L. Pitts, S. R. Davidson and W. A. Pohlchuck, "Delay/Disruption Tolerant Networking for the International Space Station (ISS)," 2017 IEEE Aerospace Conference, Big Sky, MT, USA, 2017, pp. 1-14, doi: 10.1109/AERO.2017.7943857.

[14] “Five Years of Wi-Fi Outside the International Space Station”, https://ntrs.nasa.gov/citations/20210015337.

[15] L. Daigle, “Permissionless Innovation – Openness, not Anarchy”, Internet Society blog post, April 2014. https://www.internetsociety.org/blog/2014/04/permissionless-innovation-openness-not-anarchy/

[16] R. Gladden, E. Pereira and B. Sauer, “Lessons Learned from the Mars Relay Network: Considerations for Future Relay Networks”, 2024 IEEE Aerospace Conference, March 2024. doi: 10.1109/AERO58975.2024.10521332.

[17] E. Grayver, D. Lee, E. McDonald, J. Verville, J. Dubois, “Testset for Cis-Lunar Communications and Navigation”, *2025 IEEE Aerospace Conference*, to appear, 2025.

BIOGRAPHY

Jonathan Verville *received his B.S. in electrical engineering from Michigan Technological University. With over 21 years at NASA Goddard Space Flight Center, he has supported ground and spaceflight developments across 9 projects, including TDRS, LCROSS, and MMS, focusing on Communications and Data Systems. In his current role as the Data Systems lead for the Lunar Communication Relay and Navigation Systems project, Jonathan leads a team responsible for ensuring the successful formulation, implementation, and verification of end-to-end data services to support the Artemis missions.*

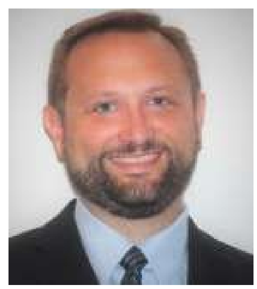

Wesley Eddy *serves as Chief Technologist at MTI Systems, Inc., where he supports government and industry customers on projects in advanced networking, Internet protocols, satellite routing, and space communications technology projects. He has contributed to several IETF working groups and standards, including serving as Transport Area Director, and chair of multiple Working Groups. He holds M.S. and B.S. degrees in Computer Science from Ohio University.*