The Mars Relay Network (MRN) Participation Guide

Version 1.0 (Initial Release), 1 July 2025

Questions on the content of this document or to begin discussions on participation in the MRN should be submitted to imrcwg_admin@jpl.nasa.gov.

The contents of this document were assembled at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration (80NM0018D0004), with the participation of the European Space Agency.

Specific contributors to this document include Roy Gladden, Eve Pereira, Brandon Sauer, and Neil Chamberlain from the Jet Propulsion Laboratory, California Institute of Technology; and Silvia Sangiorgi, Peter Schmitz, and Olivier Reboud from the European Space Agency.

Table of Contents

1 Introduction

The current Mars Relay Network (MRN) represents a highly successful international collaboration between the National Aeronautics and Space Administration (NASA) and the European Space Agency (ESA) [1]. Its architecture is the result of an evolution of capabilities that were first operationally exercised when the twin Mars Exploration Rovers (MER), Spirit and Opportunity, arrived at Mars in early 2004 [2] [3]. Originally considered a technology demonstration, the provision of telecommunications relay to landed spacecraft via orbiters has become a fundamental infrastructure to facilitate the exploration of the surface of Mars.

Today, there are five orbiters from NASA and ESA providing relay services to NASA's Curiosity and Perseverance rovers [4]. These orbiters use a variety of ground stations on Earth to transfer over 99% of all data to and from the rovers on the surface of Mars (see Figure 1.1).

This guide is intended to provide context and direction for new space agencies and other organizations that are interested in participating in the MRN, contributing spacecraft and/or other elements that can enhance the ongoing exploration of Mars. All such organizations are invited to review this guide and consider how they may appropriately and meaningfully participate.

Additional material is provided in a separate technical support package (TSP) to provide more detail and context regarding the architecture and operation of the MRN, from both flight and ground perspectives. This TSP will be referenced at meaningful locations throughout this guide. Other references are included throughout as representative supplemental material, not as a comprehensive bibliography.

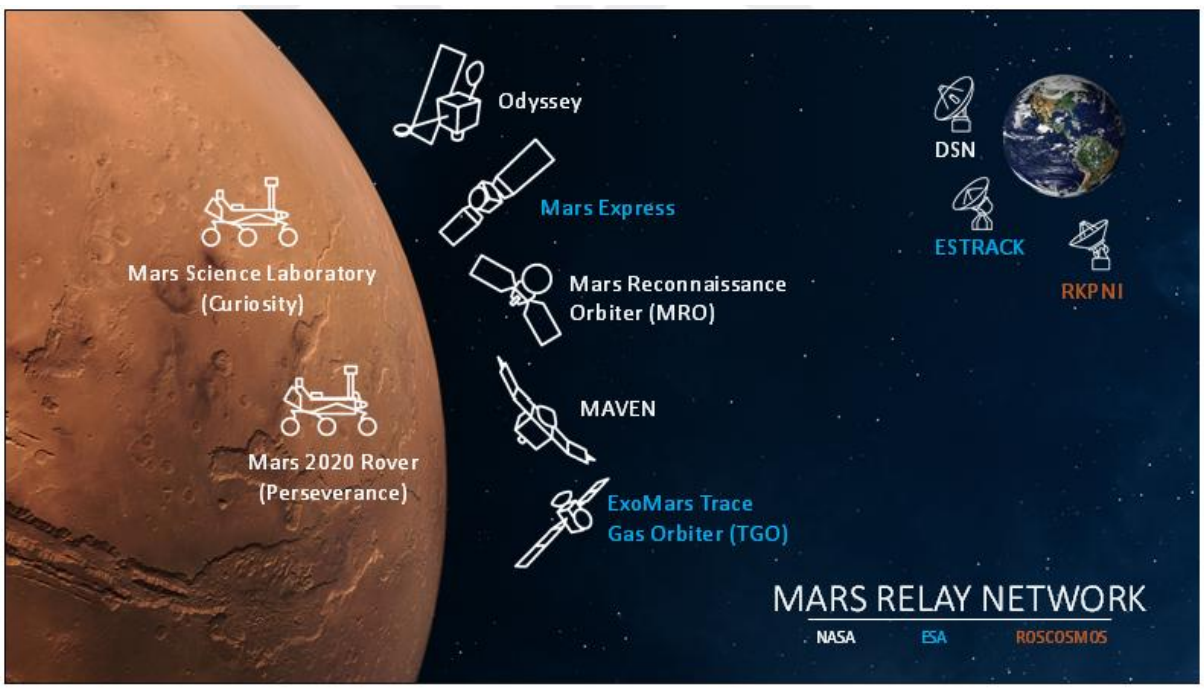

Figure 1.1. Elements of the Mars Relay Network in 2024.

1.1 Rationale for the Relay Network

There are several benefits that orbital relay provides to existing landed spacecraft. On the surface of Mars, landers usually lack line-of-sight to the Earth and the Sun for about 12 hours per day due to planetary rotation. In addition, the long and variable distance between Earth and Mars, coupled with the power requirements needed to communicate over that long distance, would make it necessary for a lander to carry large and power-hungry communications systems, or else suffer from low data rates. In comparison, Mars orbiters tend to have near-continuous access to solar energy and experience generally short or non-existent Earth occultations, depending on their orbit.

From the lander viewpoint, the shorter range to an orbiter, being a small fraction of the range to Earth, makes it possible to achieve high-rate data transfers using low-mass, low-power, and physically smaller radios, relying on the orbiters to carry the burden of transferring the data over the long distance to Earth.

In addition, it has proven vital to observe how vehicles perform during critical events, such as a Mars lander's entry into the Martian atmosphere, descent to the surface, and eventual landing; a series of events referred to as EDL [5]. In particular, lander antennae are tested when communicating with Earth during descent, primarily because of high Doppler caused by drag, uncontrolled vehicle motion, and power limitations. Here again, the shorter range to an orbiter makes it possible to achieve higher rates during this very challenging mission phase.

Radiometric observations (e.g., range and Doppler) made on these proximity links also enable precision position and navigation services in the Mars reference frame [6].

To date, relay services have been provided via orbiters whose primary purpose is to acquire science data, with relay capabilities implemented as a secondary function. Ground operations for the orbiters are impacted with greater complexity due to the need to interface with the lander systems, and the data throughput to and from the orbiters themselves are reduced somewhat due to the need to transfer relay data. Similarly, ground operations for the landers are also impacted by the introduction of new interfaces that allow them to access and take advantage of the relay services provided by the orbiters.

Despite these ramifications, NASA and ESA have agreed that including this relay function on science orbiters enhances the overall objectives of both space agencies to explore Mars. The collaborative relationship formed between the two agencies has resulted in an environment where higher-level exploration objectives can be pursued, cross-agency contributions to new missions are typical, and science activities across mission boundaries provide new insights to enhance our understanding of Mars.

New missions joining the Mars Relay Network stand to gain significant operational advantages, including enhanced data return, improved communication reliability, and access to contingency relay support. However, these benefits are contingent on completing all necessary steps to fully integrate the mission into the MRN. This includes establishing compatible communication interfaces, implementing validated operational procedures, and ensuring coordinated scheduling with existing network assets. For example, the capability to accommodate late-notice relay requests or provide robust contingency support can only be realized once the new mission's systems and processes have been properly aligned and tested within the MRN framework. Therefore, early and thorough integration efforts are critical to unlocking the full value of participation in the network.

1.2 Applicability of This Guide

This guide is intended to be applicable to any organization that expects to materially participate in the MRN, either by deploying a spacecraft (usually an orbiter) to provide relay services at Mars or by sending a spacecraft[1] that requires the use of those services (either as landers or as orbiters). It provides guidance for how one may join the network from administrative, technical, and operations standpoints.

This guide includes the following sections:

- **Section 2:** Information regarding the process whereby organizations may administratively join the network when sending new spacecraft to Mars that will operate as part of the MRN.
- **Section 3:** Information uniquely suited for those who wish to send a new spacecraft to Mars to provide relay services to existing and future spacecraft. This section describes the nature of the critical interfaces in the proximity environment (i.e., at Mars) and on Earth, information useful to join the network from a technical standpoint.
- **Section 4:** Information uniquely suited for those who wish to send new landed or orbiter spacecraft to Mars with the intent of utilizing existing relay services. As in Section 3, this section discusses the critical interfaces in both the proximity environment and on Earth, information useful to join the network from a technical standpoint.
- **Section 5:** Information for all participants in the MRN relating to ground systems and interfaces for transferring data between organizations (on Earth), and for planning and coordinating the relay activities described in Section 6, information useful to join the network from an operations standpoint.
- **Section 6:** Information for all participants in the MRN relating to ground planning processes that are in use today, information useful to join the network from an operations standpoint.
- **Section 7:** A summary of all roles and responsibilities expected of new MRN participants.
- **Section 8:** A summary of additional documentation that may be provided to new MRN participants.

All sections are intended to be complementary and refer to each other to avoid a duplication of information.

Since the MRN is an operational network, many of the nuances of its operations have already been identified and realized. New participants are therefore asked to interface with existing systems, processes, and services in known ways – especially leveraging established communications protocols – wherever possible. Necessary changes and augmentations to the existing network are expected to be carefully

[1] It is recognized that future users of the MRN may include human explorers, which may experience the same technical challenges of securing communications services as described in this guide.

coordinated through existing forums (see Section 1.5) to maintain interoperability and to enable efficient implementation.

Some explanations are included throughout this guide for why the existing network is architected the way it is and operates the way it does. Some information is also included to suggest capabilities that are not present in the existing network but may be desired for next-generation relay services, capabilities that may be first steps towards the implementation of a solar system internetwork [7]. These discussions are included in text boxes.

1.3 Terminology

Throughout this guide, various terms are used to describe the relay network and its functionality, as described here:

- Relay: The movement of data between two systems via the utilization of an intermediary system. In this context, this includes the movement of data between Earth and a spacecraft at Mars (either on or near the surface, or in orbit) through other spacecraft at Mars.
- Relay service provider (or simply "provider"): An entity that functions to transfer data from one entity to another. In some contexts, this may be considered a network "node", but here this term describes an orbiter at Mars that provides relay services and equally the organization that operates it. In the current MRN, this includes the relay orbiters at Mars, such as NASA's Mars Reconnaissance Orbiter (MRO) and ESA's ExoMars Trace Gas Orbiter (TGO).
- Relay service user (or simply "user"): An entity that is a source for data that relies on relay services to transfer data from one place to another. In some contexts, this may be considered a network "endpoint", but here this term describes either an orbiter or a landed spacecraft at Mars that utilizes relay services and equally the organization that operates it. In the current MRN, this includes NASA's rovers on Mars, Curiosity and Perseverance.
- Relay services: A collective term for the sum of all services that are provided by the operators of a relay service provider spacecraft and its operators on behalf of a relay service user spacecraft and its operators.
- Relay network: A generic term used here to describe the sum of all entities that participate in the transfer of data between a relay service user spacecraft at Mars and its operators on Earth.
- Direct-to-Earth (DTE) link: A radio interface that transmits data from a spacecraft at Mars directly to Earth-based ground stations.
- Direct-from-Earth (DFE) link: A radio interface that transmits data from an Earth-based ground station directly to a spacecraft at Mars.
- Direct-with-Earth (DWE) link: A term that includes both a DTE link and a DFE link, the bi-directional, direct transmission of data between Earth-based ground stations and spacecraft at Mars.

- Proximity link: A radio interface that is used to transmit data from one spacecraft to another at Mars.

- Return-link: A term that describes the movement of a relay service user's data from Mars to Earth via a relay service provider to the user's ground operators. In the context of the proximity link, this refers to the transmission of data from a relay service user spacecraft to a relay service provider spacecraft.

- Forward-link: A term that describes the movement of a relay service user's data from Earth to Mars via a relay service provider to the intended user spacecraft. In the context of the proximity link, this refers to the transmission of data from a relay service provider spacecraft to a relay service user spacecraft.

- View period (or "overflight"): A period during which a relay service provider spacecraft is geometrically in view of a relay service user spacecraft. This represents an opportunity for a relay session to occur.

- Relay session: A period during which a relay service user spacecraft and a relay service provider spacecraft intend to transfer data between them, either uni-directionally or bi-directionally. A relay session occurs during a view period, though a relay session may be scheduled to occur or extend outside of a view period for various reasons.

1.4 Simplified Topology of the Mars Relay Network

As shown in Figure 1.1, the Mars Relay Network (MRN) consists of elements operated by several organizations. Each element has a connection with most other elements via radio or ground system interfaces. This web of interfaces may be reduced to a simplified topology, as shown in Figure 1.2.

Today, the *Relay Service User Operators* are the operators of the Mars rovers, which themselves are the *Relay Service User Spacecraft*. The *Relay Service Provider Operators* are the operators of the Mars orbiters in the MRN, which themselves are the *Relay Service Provider Spacecraft*. The *Deep Space Tracking Network(s)* are those deep space antennas and their operators on Earth that facilitate communications over the long and variable distance between Earth and Mars. The *proximity link* and *Direct-with-Earth (DWE) Link* are the radio frequency (RF) links between the relevant spacecraft.

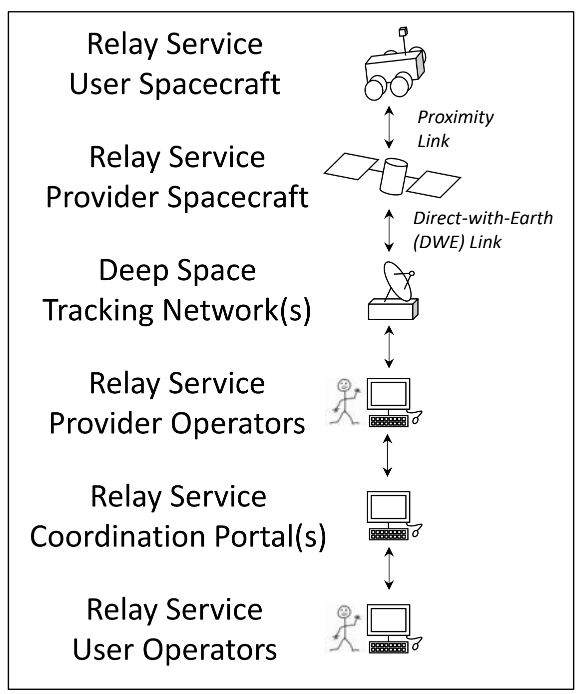

Figure 1.2. Simplified Topology of the Mars Relay Network

The *Relay Service Coordination Portal(s)* are services that facilitate data exchanges between the operators. As will be shown in Sections 5 and 6, they serve a key role in enabling the cross-organizational operations of the network and provide a means to equitably allocate the scarce resources of the network, both in bandwidth and in time.

The top half of the diagram in Figure 1.2 (from *Deep Space Tracking Network(s)* and up) represents radio interfaces. The proximity link is specifically discussed in Sections 3.2 and 4.2. The bottom half of the diagram (from *Deep Space Tracking Network(s)* and down) represents ground system interfaces between the various operators, including the *Relay Service Coordination Portal(s)*.

It is the responsibility of the relay service providers to secure tracking support by *Deep Space Tracking Network(s)*, being an intermediate system in the flow of relay data through the network and an integral part of the provision of relay services. Guidance for securing such support is not provided in this guide.

1.5 Organizational Context of the Mars Relay Network

The MRN is presently operated as a joint venture between NASA and ESA under the Terms of Reference (ToR) for the International Mars Relay Coordination Working Group (IMRCWG) [8]. Though NASA and ESA independently retain decision-making authority over the operations of their own spacecraft at Mars, the IMRCWG is the forum in which cross-agency issues related to the relay network are resolved for existing and near-term Mars spacecraft. This includes managing non-conformances by new participants as they join the network.

Under direction from NASA Headquarters, the Mars Exploration Program (MEP) Office, hosted at NASA's Jet Propulsion Laboratory (JPL), takes a lead role for NASA within the IMRCWG to coordinate the activities of the MRN. The MEP acts as the focal point for all NASA-related activities within the MRN. ESA's Mission Operations Department at the European Space Operation Centre (ESOC) serves as the focal point for all ESA-related activities within the MRN, working in close collaboration with ESA's Human Spaceflight and Robotic Exploration Directorate's Mars and Beyond Projects Group.

The proximity links at Mars are implemented using protocols defined by the Consultative Committee for Space Data Systems (CCSDS), and the frequencies utilized at Mars are specified by the Space Frequency Coordination Group (SFCG).

NASA's MEP co-sponsors with NASA's Multimission Ground Systems and Services (MGSS) Program the Mars Relay Operations Service (MaROS), acting as a Relay Service Coordination Portal as shown in Figure 1.2 to provide standardized ground interfaces in support of the planning and coordination of the relay services provisioned at Mars, as discussed in Sections 5 and 6. ESA's ESOC additionally operates the European Relay Coordination Office (ERCO), which manages ESA-internal relay activities and interfaces with MaROS in support of the planning and coordination of relay services.

The IMRCWG collaborates with the International Mars Exploration Working Group (IMEWG) and the Interagency Operations Advisory Group (IOAG) to coordinate recommendations related to the future of the MRN to participating organizations.

2 Joining the Network

The general process to join the MRN includes administrative, technical, and operations considerations, as briefly outlined here. New participants should perform each of the following, as appropriate:

1. Establish familiarity with the MRN by reviewing this guide.
2. Affirm interest in participating in the MRN (either as a provider of new relay services and/or as a user of existing relay services) and confirm achievability of participation with the IMRCWG. Such interest might be expressed through a formal letter submitted by the new participant to the IMRCWG.
3. For non-NASA and non-ESA government entities, join the IMRCWG by signing the Terms of Reference.[2]
4. Develop all necessary agreements (as described below), typically within the context of the IMRCWG.
5. Develop relay-related flight and ground interfaces with the participation of the IMRCWG.
6. Construct new spacecraft with flight interfaces consistent with the interfaces developed in Step 5, and perform pre-launch testing of all flight interfaces (see Sections 3.4 and/or 4.4).
7. Implement ground systems, ground interfaces, and operations systems consistent with the interfaces developed in Step 5; and perform all appropriate testing (see Sections 5.5 and 6.6).
8. Secure ground tracking services, as needed, consistent with expectations for participation in the MRN. Guidance for securing these services is not provided in this guide, though the IMRCWG can provide advice on and advocate for such services.
9. Deliver new participating spacecraft to Mars. Guidance for launch and/or delivery services to Mars is not provided in this guide, though the IMRCWG can provide advice on and advocate for such services.
10. Perform any *in situ* commissioning and in-flight testing and characterization, as needed.
11. Operate the new spacecraft at Mars with its relay activities coordinated as a part of the MRN, as intended.

Regarding Steps 2 and 3, communication with the IMRCWG can be initiated by sending an email to imrcwg_admin@jpl.nasa.gov. This email address can also be used if there are any questions regarding the contents of this guide.

Regarding Step 4, the agreements that need to be written vary depending on the nature of the new participant's organizational relationship with the members of the IMRCWG. For example, non-U.S. entities that are not already affiliated with ESA may require the development of a Technical Assistance Agreement

[2] At this writing, it is not yet clear how non-government organizations should engage with the IRMCWG. However, it is expected that non-government organizations within the U.S. or within member states of ESA would first engage with NASA or ESA, as appropriate, before engaging with the IMRCWG.

(TAA) with Caltech to comply with U.S. International Traffic in Arms Regulations (ITAR). Once appropriate agreements are in place, additional technical documentation can be provided that describe in detail the systems, services, operations approaches, and interfaces used by the MRN, as described in Section 8.

Other agreements that may need to be written include:

- Memorandums of Understanding (MOUs) or Memorandums of Agreement (MOAs), typically written at the level of participating space agencies to establish a basic agreement for relay activities.
- Service Agreements, which typically outline the expectations for what relay services will be provided or used. These may be written between participants in the MRN or with adjacent organizations, such as those that operate deep space tracking networks.
- Interface Definition Documents (IDDs), which outline the nature of relay services that can be provisioned and what technical interfaces are required to secure them.
- Other documentation to describe activities around special events, operations processes, operating procedures, unique data exchanges, etc., as needed.

Regarding Steps 6 and 7, testing of pre-launch flight elements, related ground systems and interfaces, and operations processes are described in Sections 3.4, 4.4, 5.5, and 6.6. All cross-organizational tests are identified in Service Agreements, with the details documented separately as conditions merit. IDDs typically guide the nature of needed flight and ground interface testing. Planning process testing is generally guided by the operations processes themselves (outlined in Section 6.4). New participants typically determine the schedule and rigor of all testing.

Adaptation of this process to join the MRN may be coordinated with the IMRCWG, if necessary.

3 Relay Service Guide Unique to Relay Service Providers

3.1 Applicability of This Section

This section is applicable to any organization that expects to materially participate in the MRN by sending a spacecraft to Mars to provide relay services. It provides guidance for how one may participate in the network from a technical standpoint, and specifically describes the following:

- The proximity link to be implemented (Section 3.2).
- The relay services to be provided (Section 3.3).
- Relay service interface testing (Section 3.4).

3.2 Proximity Link Implementation

Currently, the MRN implements the proximity link at Mars using the CCSDS Proximity-1 Protocol [9] [10] [11] [12] on UHF frequencies. New providers are expected to be compatible with this protocol and frequencies to provide relay services to existing and forthcoming relay service user spacecraft, and it is acknowledged that new relay service providers may enhance these capabilities with other protocols and on additional frequencies.

The Proximity-1 specification permits enterprise-specific tailoring, and therefore it is necessary to work with the IMRCWG to understand the nuances of any new implementation to ensure interoperability. This is needful regarding both the implemented behavior of the protocol and the frequencies to be supported. See TSP Section 2 for a description of the current implementation of the Proximity-1 protocol.

> *Future Consideration: Next-Generation Protocols and Frequencies*
>
> Though the current MRN presently relies on the Proximity-1 protocol to facilitate interoperability, new protocols are expected to be available in the coming years. Most notably, an update to the Proximity-1 specification is forthcoming that will enable the use of higher data rates and will specify other operating frequencies in the UHF band. The Unified Space Link Protocol (USLP) [13] is similarly expected to be useful.
>
> However, the continued use of the Proximity-1 protocol and designated frequencies, as presently implemented, is expected to be needed in support of existing and known forthcoming spacecraft, and to facilitate new spacecraft that may be unable to take advantage of more capable protocols. To this end, new relay service providers are encouraged to be capable of providing relay services using this "legacy" protocol and related frequencies, and to implement or stand ready to implement additional protocols and frequencies in future years.

In this context, the future augmentation of the MRN with new capabilities is expected to be influenced by other organizations in coordination with the IMRCWG. Notably, the IOAG's Mars and Beyond Communications Working Group has issued a set of recommendations [14] and NASA and ESA have jointly issued the LunaNet Interoperability Specification [15], which are likely to be applicable.

3.3 Relay Service Descriptive Information

This section provides guidance to new relay service providers regarding how relay services should be documented for new users and implemented in both the flight and ground systems.

3.3.1 Interface Definition Document (IDD)

An organization that operates a relay service provider spacecraft at Mars must describe the available relay services, including the following:

- A general description of the proximity link radio(s) that will provide relay services.
- A description of the antenna(s), associated antenna pattern(s), and radio frequency (RF) gain information of the relay service provider spacecraft.
- A description of any self-generated electromagnetic interference (EMI) that may affect relay performance.
- A description of the supported protocol modes and any other supported enterprise-specific modes (see TSP Section 3 for examples).
- A description of the DWE communications infrastructure, DWE RF frequencies (including optical, if applicable), expected DWE data rates as a function of Mars-Earth distance, and the expected ground station availability.
- A description of how forward-link data and return-link data will be handled onboard the relay service provider spacecraft and in the ground system, including data flows and the identification of when, where, and how data is expected to move through the system end-to-end (which may involve MaROS, see Section 5).
- A description of expected end-to-end data latencies, described as a function of the DWE system, the onboard data handling system, the ground system, and any other potential sources of delay when transferring data through the system end-to-end.
- A description of parameters that are required to schedule individual relay sessions.
- A description of telemetry and other quality of service data that will be made available to monitor the health and performance of the proximity link on a per-session basis (see also Section 6.4.5).

IDDs are typically written as a one-to-many document whereby the described interface is between one relay service provider and many relay service users. IDDs are necessarily comprehensive documents that describe all possible modes of operation that a relay service provider offers, and it should be expected that relay service users may utilize only a portion of the available modes of operation. To this end, the IDD acts as the principal guide for what testing should be performed (in terms of exercising use cases) to verify radio compatibility (see section 3.4.1) and to bound activities during operations. An outline of the contents of a generic IDD, including a basic description of common parameters required to schedule individual relay sessions, is provided in TSP Section 3.

Note that the information described above and as documented in an IDD should not be confused with regular and ongoing data exchanges that are necessary during operations, as described in Section 6.

3.3.2 Relay Service Physical- and Data-Link Layers

In terms of the specific implementation of hardware and software to provide relay services, new relay service providers in the MRN are strongly encouraged to consider and trade the full system of components that incorporates all aspects of the physical- and data-link layers upon which data transport will depend. These components can be viewed from the perspective of an RF link budget, the major considerations of which are summarized here.

Radio hardware and software:

- Modes of operation, such as duplexity (communicating simultaneously in both the return- and forward-link directions during a relay session), multiplexing (supporting more than one user simultaneously), quality of service (especially expressing the reliability of successfully transferring data during a relay session), and availability of data throughput optimizations (such as an adaptive data rate algorithm, see TSP Section 2).
- Supported frequencies, data rates, modulation schemes, and encoding schemes.
- Power amplifier output power and low noise amplifier.
- In-field programmability.
- Frequency agility.

Antennas:

- Polarization.
- Beam pattern and gain.

Propagation Channel:

- Ephemeris, slant range, antenna pointing, and overflight duration.
- Terrain scattering and multipath effects.
- Noise and interference.

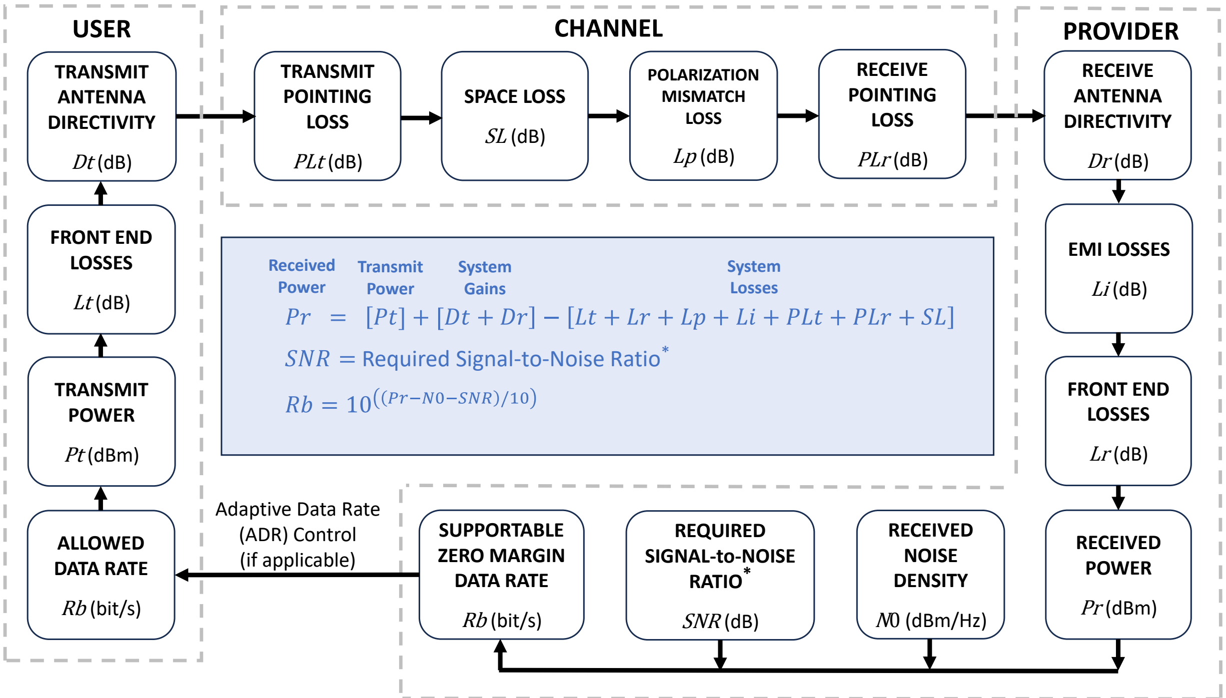

** The required signal-to-noise ratio depends on the modulation scheme, the required bit error rate (BER), and the channel coding scheme (e.g., SNR ≳ 5 dB for binary phase shift keying, BER = 10^{-6}, and convolutional coding).*

Figure 3.1. Example return-link signal flow and link budget calculation.

An example return-link signal flow with its link budget calculation is illustrated in Figure 3.1.

As a reference point for such trades, we provide a summary of the relay configurations as presently implemented by JPL in the MRN in TSP Section 4, with representative descriptions of the hardware and software implementations.

3.3.3 Relay Service Operational Considerations

MaROS, as will be described in detail in Section 5, serves as the centralized hub for scheduling inter-agency and NASA-internal relay sessions and for reporting the performance of executed relay sessions (the beginning and the end of the end-to-end relay process in the flight context). It is the preferred means of exchanging relay planning and coordination data across all participants in the MRN. It is notable that ESA's implementation of the ERCO is a ready example of how an agency may implement a standalone relay system to coordinate activities between spacecraft managed by a single agency. In this case, ESA activities that affect or interact with the balance of the relay network are reported to MaROS to enable cross-agency coordination activities and to facilitate situational awareness. The functionality of ERCO and how it interoperates with MaROS is described in TSP Section 9.

MaROS can act as a portal for transferring forward-link data between the operators of a relay service user and the operators of a relay service provider. Similarly, MaROS can also serve as a portal for delivering return-link data.

The selection of useful relay sessions is the responsibility of the relay service users, and there are three important operational considerations that influence these selections:

1. The availability of the relay service provider spacecraft at desired times.

   Ideally a relay service provider spacecraft is available to provide relay services during all view periods, though it is understood that operational realities may make this impractical. It is the task of the relay service provider operators to communicate periods when the provision of relay services might be constrained.

2. The amount of data that may be transferred during a relay session, especially in the return-link direction for Mars science missions.

   The task of calculating how much data can be transferred during a relay session is typically left to the relay service user operators. This is most often achieved by considering all aspects of the *in situ* physical and data-link layers, including the data rates achievable during the relay session, the antenna patterns of the radios, the geometric relationship between the two radios during a relay session, etc. Relay service user operators may choose to implement their own prediction tools, or they may utilize the Relay Telecom Predictor (RTP), available via MaROS (see Section 5.3).

3. The latency (or delay) associated with transferring data to and from the relay service user operators and the relay service user spacecraft.

   This may be important to facilitate the operations cadence of the relay service users, to understand when forward-link data must be inserted into the network, to scale the amount of data that can be collected during operations, etc. MaROS can calculate these latencies, as described in Section 5.3, or the operator of a relay service provider spacecraft may estimate them independently.

Relay service provider operators are responsible to provide all the information needed to meaningfully schedule relay sessions given these operational considerations. Some of this information should be made available in an IDD, with the rest made available during operational interactions as described in Section 6 via the methodologies described in Section 5. A summary of the operational responsibilities of relay service providers is included in Section 7.

It is understood that there may be peculiarities to how relay services may be provided because of the way they have been implemented by a relay service provider. For example, most NASA and ESA orbiters support the ability to provide relay services to two relay users during similar view periods in a time-shared arrangement. In this scheme, a relay session with one user would occur and then be followed by a second relay session with a different user after a short interruption[3]. It is the responsibility of the relay service provider operators to ensure that these peculiarities are well-communicated, such as in an IDD.

[3] This capability enables a certain flexibility when providing relay services to co-located relay service user spacecraft, but should not be confused with a true Time-Division Multiple Access (TDMA) scheme [16], though it shares some common features.

Another operational consideration might be when the transmission of relay data to Earth-based ground tracking stations is interrupted by problems with the ground stations or with the relay orbiter itself, resulting in the delivery of partial data products and necessitating the retransmission of missing data. In cases such as these, effective communication between the operators is needed, and this communication is facilitated using both software and processes as outlined in Sections 5 and 6.

3.4 Relay Service Interface Testing

The critical interfaces for a relay service user are the flight system interfaces of the proximity link and the ground system interfaces for coordinating relay services. Ground system interfaces are largely implemented via MaROS, and the testing of these interfaces are described in Sections 5.5 and 6.6. This includes those related to the transfer of data through applicable ground systems and the operations interfaces related to the scheduling of relay services (and the timelines for doing so).

Flight system interfaces in the proximity environment (at Mars) must be tested to confirm that the radios can functionally communicate with each other across the desired operating modes. This testing is focused on ensuring that data can be successfully passed through the network without corruption and in a timely manner. To this end, radio compatibility testing (see Section 3.4.1) and end-to-end system testing (see Section 3.4.2) are performed to verify these interfaces.

3.4.1 Radio Compatibility Testing

To verify flight interfaces on the ground prior to flight operations, new relay service providers should make available a testbed that emulates the intended flight system. Typically, such a testbed should:

- Be capable of receiving flight-like sequences and forward-link data via a process that emulates the DFE link.
- Set up and run the relay session, as connected with a similar relay service user testbed, to demonstrate the exchange of forward- and return-link data.
- Perform the transfer of collected relay data, any

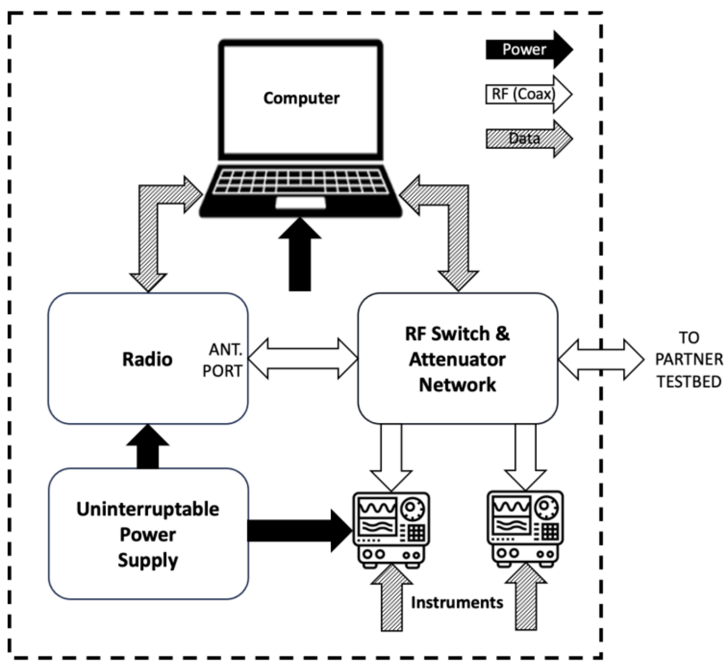

Figure 3.2. Generic testbed layout.

ancillary data, and any engineering telemetry to the operators of the relay service user spacecraft.

Historically, the relay service provider's testbed and relay service user's testbed have been connected by a coaxial cable, and not through open-air antennas, and the RF channel is commonly modeled by varying attenuators in-line with the coaxial cable link. Optionally, a channel simulator can be added to emulate Doppler shift and fast fading. See Figure 3.2 for an example of a generic testbed layout.

This testing typically requires significant advanced coordination and preparation. This includes the documentation of test plans and procedures (including test matrices and success criteria) and addressing the logistics of transporting necessary test equipment to appropriate test venues (which may involve navigating import and export laws if international travel is required). A summary of the expected contents of a test plan is included as TSP Section 5.

Historically, a typical test campaign takes about one work week to perform between one relay service user and one relay service provider. Additional time is often needed for setting up and checking out the test equipment, analyzing the data, troubleshooting discovered issues, performing any necessary go-backs in response to discovered issues, and concluding the test campaign. The duration of the test campaign should reflect its complexity and scope, and it is recommended that healthy schedule margins be included to ensure that at least minimal test goals can be accomplished while dealing with issues as they are discovered.

> *Future Consideration: A Portable Relay Testbed*
>
> NASA's MEP is considering the development of a portable relay testbed (PRT) to facilitate compatibility testing with future network participants. This PRT is envisioned to have similar functionality to relay test equipment developed at JPL and elsewhere while being sufficiently compact to transport with supporting personnel. Presently conceived as supporting the Proximity-1 protocol on UHF frequencies, the PRT architecture would be intended to be adaptable so that it may ultimately facilitate the testing of anticipated future communication protocols and frequencies as they arise.
>
> This PRT is not yet available.

3.4.2 End-to-End System Testing

End-to-end (E2E) system tests demonstrate the flow of relay data from request to receipt. This process necessarily begins and ends with the relay service user, the endpoints in the relay network, and involves the participation of both the relay service provider and the relay service user. Typically, a subset of the compatibility testing (per Section 3.4.1) is exercised as part of this testing. Key elements of E2E testing include (with responsibilities indicated in parentheses):

- The selection of the compatibility test cases that will be used for E2E testing (user and provider).
- The selection of relay sessions to execute, nominally through MaROS (user).
- The acknowledgement of selected relay sessions, nominally through MaROS (provider).

- The verification of all necessary commanding to implement the selected relay sessions (user and provider).
- The transfer of data across a simulated DFE link of the relay service provider commanding from the relay service provider mission operations center (MOC) to the relay service provider testbed, which may or may not include test articles that emulate ground stations (provider).
- The transfer of any forward-link data from the MOC of the relay service user to the MOC of the relay service provider, nominally through MaROS (user).
- The transfer of data across a simulated DFE link of any forward-link data from the MOC of the relay service provider into the relay service provider testbed, which may or may not include test articles that emulate ground stations (provider).
- The execution of the compatibility test cases in the interfaced relay service provider and relay service user testbeds (user and provider).
- The transfer of data across a simulated DTE link of collected relay data with any ancillary data and/or engineering telemetry from the relay service provider testbed to the MOC of the relay service provider, which may or may not include test articles that emulate ground stations (provider).
- The transfer of collected return-link data with any ancillary data from the MOC of the relay service provider to the MOC of the relay service user (provider).
- The conversion of the return-link data with any ancillary data into usable data (user).
- The reporting of engineering telemetry that reports the performance of executed relay sessions to a centralized repository, nominally MaROS (user and provider).

In flight operations, this described sequence of events happens on an established timeline, with many of the events occurring back-to-back in an automated fashion (e.g., relay data is transferred to the relay service provider MOC as soon as the relay session is completed, assuming the availability of a deep space antenna). During these tests, it is often acceptable to perform E2E testing in a piecewise fashion without adhering to a flight-like timeline. This can be the case, for example, if the E2E flight interfaces are not operational at the time of testing, necessitating the use of some other acceptable surrogate pathway. It may also be appropriate to test multiple relay sessions in rapid succession to expedite the test campaign, a scenario which would likely never occur during flight operations. It is recommended that each test case be carefully considered to determine if there is a need to test it in as flight-like a manner and timeline as possible, exercising the full E2E system.

The mechanisms for performing data interchanges via MaROS are described in Section 5. The processes for selecting and acknowledging selected relay sessions, for exchanging forward- and return-link data in the ground system, and for reporting post-session engineering telemetry, as outlined above, are described in detail in Section 6.

4 Relay Service Guide Unique to Relay Service Users

4.1 Applicability of This Section

This section is applicable to any organization that expects to materially participate in the MRN by sending a spacecraft to Mars that will utilize relay services. Such service users are typically Mars landers or rovers but may also include orbiters. This section provides guidance for how one may participate in the network from a technical standpoint, and specifically describes the following:

- The proximity link to be implemented (Section 4.2).
- The relay services that are available for use (Section 4.3).
- Relay service interface testing (Section 4.4).

4.2 Proximity Link Implementation

Currently, the MRN implements the proximity link at Mars using the CCSDS Proximity-1 Protocol [9] [10] [11] [12] on UHF frequencies. New users are expected to be compatible with this protocol and frequencies to interface with existing relay service provider spacecraft, though it is acknowledged that new relay service providers may enhance these capabilities with other protocols and additional frequencies.

The Proximity-1 specification permits enterprise-specific tailoring, and therefore it is necessary to work with the IMRCWG to understand the nuances of the available implementations. This is needful regarding both the implemented behavior of the protocol and the frequencies that are supported. See TSP Section 2 for a description of the current implementation of the Proximity-1 protocol.

4.3 Relay Service Descriptive Information

This section provides guidance to new relay service users regarding how relay services are documented and how these services may be accessed in both the flight and ground systems.

4.3.1 Relay Service Utilization

An organization that operates a relay service user spacecraft at Mars must understand how they will access the available relay services, which are typically documented consistent with Section 3.3.1. To access these services, a new relay service user must also understand the following about their own systems:

- The proximity link radio(s) that will utilize relay services.
- The antenna(s), associated antenna pattern(s), and RF gain information of the relay service user spacecraft.

- The impact of any self-generated electromagnetic interference (EMI) on link performance.
- The supported protocol modes that are desired to be utilized (see TSP Sections 2 and 3 for examples), and the means whereby relay sessions are scheduled (see Sections 5 and 6).
- The means whereby forward-link data and return-link data will be transferred to and received from the relay service providers, nominally via MaROS (see Section 5).
- A description of the telemetry and other quality of service data that will be made available to monitor the health and performance of the proximity link on a per-session basis (see also Section 6.4.5).

Note that this information describes the nature of the relay service user's system that will access the provided relay services. Unlike similar data needed from relay service providers (as in Section 3.3.1), this data does not necessarily need to be documented and exposed to the relay service providers. This information should not be confused with regular and ongoing data exchanges that are necessary during operations as described in Section 6.

4.3.2 Relay Service Physical- and Data-Link Layers

In terms of the specific implementation of hardware and software to access relay services, new relay service users in the MRN are strongly encouraged to consider and trade the full system of components that incorporates all aspects of the physical- and data-link layers upon which data transport will depend. Fundamentally, the major considerations for a relay service user and a relay service provider are the same; refer to Section 3.3.2.

4.3.3 Relay Service Operational Considerations

MaROS, as will be described in detail in Section 5, serves as the centralized hub for scheduling inter-agency and NASA-internal relay sessions and for reporting the performance of executed relay sessions (the beginning and the end of the end-to-end relay process in the flight context). It is the preferred means of exchanging relay planning and coordination data across all participants in the MRN. It is notable that ESA's implementation of the ERCO is a ready example of how an agency may implement a standalone relay system to coordinate activities between spacecraft managed by a single agency. In this case, ESA activities that affect or interact with the balance of the relay network are reported to MaROS to enable cross-agency coordination activities and to facilitate situational awareness. The functionality of ERCO and how it interoperates with MaROS is described in TSP Section 9.

MaROS can act as a portal for transferring forward-link data between the operators of a relay service user and the operators of a relay service provider. Similarly, MaROS can also serve as a portal for delivering return-link data.

The selection of useful relay sessions is the responsibility of the relay service users, and there are three important operational aspects that influence these selections, as described in Section 3.3.3.

Relay service user operators are responsible to use all information made available to them by the relay service providers when scheduling relay sessions. Some of this information may be made available by the relay service providers in IDDs (as in Section 3.3.1), with the rest made available during operational interactions as described in Section 6 via the methodologies described in Section 5. A summary of the operational responsibilities of relay service users is included in Section 7.

It should be understood that there may be peculiarities to how relay services are provided because of the way they have been implemented by the relay service providers. The relay service providers are responsible for communicating these, typically in an IDD, though unique operational situations may also need to be addressed. Communication between users and providers is facilitated using both software and processes as outlined in Sections 5 and 6.

4.4 Relay Service Interface Testing

The critical interfaces for a relay service user are the flight system interfaces of the proximity link and the ground system interfaces for coordinating relay services. Ground system interfaces are largely implemented via MaROS, and the testing of these interfaces are described in Sections 5.5 and 6.6. This includes those related to the transfer of data through applicable ground systems and the operations interfaces related to the scheduling of relay services (and the timelines for doing so).

Flight system interfaces in the proximity environment (at Mars) must be tested to confirm that the radios can functionally communicate with each other across the desired operating modes. This testing is focused on ensuring that data can be successfully passed through the network without corruption and in a timely manner. To this end, radio compatibility testing and end-to-end system testing are performed to verify these interfaces, as described in Sections 3.4.1 and above3.4.2.

As with relay service providers, relay service users are expected to make available a testbed that emulates the intended flight system along with applicable ground systems to support the tests.

5 Ground System Interfaces for All Network Participants

5.1 Applicability of This Section

This section is applicable to all organizations that expect to materially participate in the MRN by sending a spacecraft to Mars either to provide relay services and/or to utilize available relay services. It provides guidance for how one may participate in the network from an operational standpoint, and specifically describes the existing ground services, as follows:

- Overview of relay ground services used for planning and coordinating relay activities (Section 5.2).
- The native capabilities of MaROS (Section 5.3).
- The interfaces used to access MaROS (Section 5.4).
- Relay service interface testing with MaROS and related systems (Section 5.5).

5.2 Relay Ground Systems Overview

Inter-agency and NASA-internal relay planning and coordination activities of participants in the MRN are enabled by the Mars Relay Operations Service (MaROS), a live ground service provided by JPL that has been in continuous use since 2010 [17]. MaROS is at the center of several related systems that operate in an "ecosystem", as shown in Figure 5.1 [18]. MaROS is principally a server and database that provides a singular interface to the various participants in the MRN to facilitate the transfer of relay planning and coordination data. A highly configurable system, MaROS supports the wide variation in the nature and types of relay services provided by the relay service providers. The methodology for gaining access to MaROS and related systems is described in TSP Section 6.

Future Consideration: Delay/Disruption Tolerant Networking

In anticipation of the adoption of Delay/Disruption Tolerant Network (DTN) technologies in the Mars Relay Network, NASA's MEP is currently working with leaders in DTN development to define MaROS's role in a DTN-enabled network. Work continues to prototype the generation of contact graph routing (CGR) plans using data that is already hosted in MaROS, and using these to simulate DTN transactions with existing MRN nodes. In the future, MaROS is expected to fulfill the key role of aggregating contact opportunities across the network and enabling the generation and distribution of CGR plans to DTN-enabled nodes.

5.3 MaROS Native Capabilities

Using data provided by the participants in the MRN, MaROS and its related systems (refer to Figure 5.1 and TSP Section 6) provide the following native functions:

- The automated calculation of view periods (by the Metrics Predict Executive, or MPX) using the latest (and sometimes frequently updated) spacecraft ephemerides provided by the spacecraft navigation teams (via the Deep Space Network's (DSN's) Service Preparation Subsystem, or SPS). These view periods represent relay opportunities.
- The provision of features used to negotiate relay service requests between the relay service users and the relay service providers.
- The automated calculation of latency predictions. For the calculation of return-link latencies, this is enabled by the modeling of the onboard data storage of relay service provider spacecraft and includes the detection of potential overflows.
- The automated detection of potential scheduling conflicts and the validation of relay session parameters for each relay opportunity.

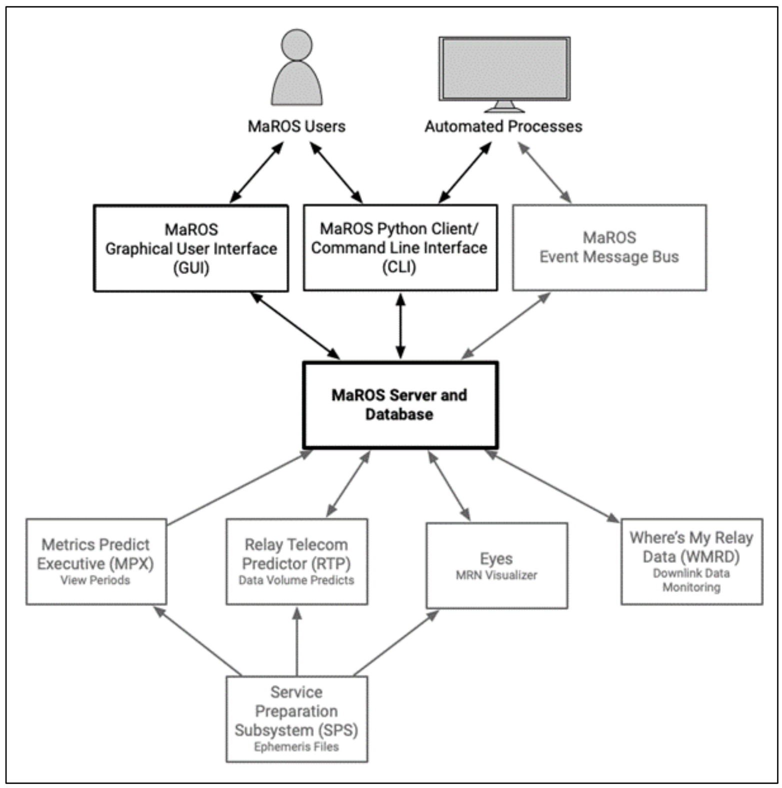

Figure 5.1: MaROS and its ecosystem.

- The provision of features to enable the transfer of files between spacecraft operators in both the forward-link and return-link directions.
- The computation of data volume predictions for each relay session (via the Relay Telecom Predictor, or RTP).
- The collection of as-flown "accountability" metrics from all participants in the MRN for each relay session to enable the reconciliation of planned/predicted performance against actuals.
- The aggregation of monitoring and tracking data for return-link data products (via the "Where's My Relay Data" service, or WMRD) to allow users to analyze the performance of the delivery of

return-link data in real-time, enabling more timely explanations for discrepancies against predicted latency times.

- The provision of a three-dimensional visual representation of the MRN to enhance situational awareness (via the "Eyes on the MRN" service, publicly available at https://eyes.nasa.gov/mrn).

5.4 MaROS Interfaces

Spacecraft operators can interact with the MaROS server via two primary interfaces: a graphical user interface (GUI) accessible through a web browser, and a command-line interface (CLI) in the form of a Python-based client that allows software interactions with MaROS, as shown in Figure 5.1. Though users can submit most relay data directly via the GUI, most interactions with the service are operationally performed via the upload and download of data files, either via the GUI or more commonly via the CLI. Both interfaces can be used to publish relay coordination data to MaROS when formatted as one of a set of extensible markup language (XML) files. Detailed software interface specifications (SISs) for these XML-based file formats are supplied to new participants in the MRN (see Section 8) and a brief introduction to these file formats is included in TSP Section 7.

Future Consideration: CCSDS Ground Interface Standards

The MaROS-supported XML-based file formats have been in continuous operational use by the ground operators of the MRN since 2010. Adopted by all NASA and ESA Mars missions that participate in the MRN, these file formats represent *de facto* multi-mission file interface standards. When MaROS first went online, there were no released CCSDS standards that could manage the necessary data interchanges. CCSDS is only now bringing to maturity a set of Cross-Support Service Management standards [19] that might be exercised in this domain. It is expected that MaROS will eventually conform to or otherwise be compatible with these specifications in the future.

Though users of the MaROS GUI can perform most of the same data exchanges that are available via the CLI, the GUI is typically used to make one-time changes to planning data, such as those that might be needed during the Strategic and Tactical Planning Processes and for Forward-Link Commanding (see Section 6.4). However, the GUI provides visuals that are not supported via the CLI, such as an overflight planning timeline, overflight search capabilities, customizable charting features in support of post-pass reporting, and other data visualizations.

An additional interface with MaROS has been recently introduced, the Event Message Bus (EMB), as shown in Figure 5.1. As events occur in MaROS (i.e., a new ephemeris file is published, a new relay service request is received, etc.), the EMB can be used to notify external systems of these changes to enable automated reactions.

Future Consideration: The Event Message Bus

The MaROS Event Message Bus (EMB) was introduced to enable event-based automation in the relay service planning process. In practice, such automation has not yet been widely adopted by the existing participants in the MRN, despite the service standing ready to do so. As of this writing, access to the EMB for users outside of JPL has yet to be established, but future support is planned. As the EMB continues to mature it is expected to play a valuable role in facilitating new forms of network automation.

5.5 Relay Service Interface Testing

For new participants in the MRN, interface testing with MaROS is performed in a standalone software environment, the MaROS "testbed" (MaROS-TB). MaROS-TB shares the same codebase as the operational version of MaROS (MaROS-Ops), but with a separate database and user permission set. This segregation between environments enables flight-like testing to occur simultaneously with ongoing operations, preserving the integrity of the operational workflows. MPX, RTP, and WMRD similarly have testbed variants that interface with MaROS-TB.

Interface testing between new MRN participants and MaROS occurs as needed by the participants. After access is granted to MaROS, most interface testing is focused on ensuring that data files are properly formatted and contain the proper content. Additional testing may be focused on implementing automated interactions with the service. As mentioned in Sections 3.4 and 4.4, MaROS also may play a role in radio compatibility testing and end-to-end system testing. Other testing may be focused on practicing operational processes, procedures, and data exchanges with relay partners, as described in Section 6.6.

6 Ground Planning Processes for All Network Participants

6.1 Applicability of This Section

This section is applicable to all organizations that expect to materially participate in the MRN by sending a spacecraft to Mars either to provide relay services and/or to utilize available relay services. It provides guidance for how one may participate in the network from an operational standpoint, and specifically describes the relay planning processes, as follows:

- Overview of relay service processes used for planning and coordination (Section 6.2).
- Overview of cross-organization meetings and other communications necessitated by relay planning and coordination activities (Section 6.3).
- Overview of the five relay coordination processes (Section 6.4).
- Overview of how emergency situations within the relay network are handled (Section 6.5).
- Relay service process testing (Section 6.6).

6.2 Relay Service Process Overview

MaROS acts as the primary service to support relay planning and coordination activities and stands ready to support the automation of these activities. However, at present, these activities are enabled by processes that remain human-involved, with the operators of each spacecraft playing key roles. Relay planning and coordination activities are therefore shepherded by the IMRCWG, which collaboratively performs the following functions:

- Ensures the proper development of all the agreements mentioned in Section 2, including Service Agreements, IDDs (as in Sections 3.3.1), test plans and procedures (as in Sections 3.4, 4.4, 5.5, and 6.6), etc.
- Facilitates communication between organizations regarding the numerous aspects of relay operations via the sponsoring of working meetings and the provision of email lists (see Section 6.3), the hosting of cross-spacecraft test and training activities, etc.
- Develops and maintains a detailed, cross-organization operations schedule, referred to as the "Mars Mid-Range Schedule".

The Mars Mid-Range Schedule describes the consolidated timeline for all scheduled relay-related data exchanges between the participants in the MRN, providing the relay community with a clear timeline for operations due dates. These scheduled data exchanges reflect the relay coordination processes, as described in Section 6.4. All these processes are focused on ensuring the smooth and successful implementation of relay activities at Mars.

6.3 Meetings and Other Communications

The ongoing success of the MRN can largely be attributed to the community of operators from the various spacecraft participating in the MRN. This community cooperates to seek solutions to challenges confronting all involved. In-person or virtual meetings and teleconferences are a key aspect of this ongoing collaboration. Several standing meetings are held:

- The IMRCWG meets in monthly teleconferences and twice-yearly in person for a Technical Interchange Meeting (TIM). Participants in the working group include management and operations personnel from all current and future MRN participants. The working group functions to manage high-level discussions of ongoing operations and future plans, including coordinating flight and ground test campaigns that cross agency boundaries. Though not a decision-making body, the IMRCWG directly guides or influences the development of agreements at all levels (as in Section 2).
- The Short-Range Relay Coordination (SRRC) Meeting is an international meeting that is held every other Tuesday and involves all current relay participants. The purpose of this meeting is to finalize the details of upcoming relay sessions and to approve the relay plans for the upcoming weeks. The meeting also serves as a forum to discuss any operational issues, to coordinate any Mars Mid-Range Schedule changes (representing process schedule changes to account for holidays, solar conjunctions, etc.), or to raise any concerns regarding relay operations. Notably, this meeting represents the end of the Strategic Planning Process, as described in Section 6.4.1.
- Non-regularly scheduled meetings are called as needed. Examples of such meetings are those useful to prepare for high-level events (such as a new Mars landing), status meetings during emergency situations (see Section 6.5), small-group meetings to discuss flight and ground system anomalies, or any other topics that need to be discussed outside of the nominally scheduled meetings.

Daily communications between all the MRN participants are facilitated using email, instant messengers, and by phone. In particular, email distribution lists are critical for handling short-term changes to the relay plan (as per the Tactical Planning Process, as described in Section 6.4.2), reporting delays in the return of relay data, and to provide notification in the event of a spacecraft emergency.

6.4 Relay Coordination Processes

Today, the successful implementation of relay activities at Mars involves five primary relay coordination processes, as shown in Figure 6.1. All these processes are supported by MaROS. Participants in these processes include management, mission planners, navigation team representatives, sequence system engineers, ground data system engineers, telecommunications engineers, science representatives, and others as necessary.

The IMRCWG acts as an observational body to ensure that all relay agreements between the MRN participants are being met and guides the flow of these processes. However, the IMRCWG as a body does not actively participate in the relay planning processes themselves, instead enabling the processes to function at the mission level.

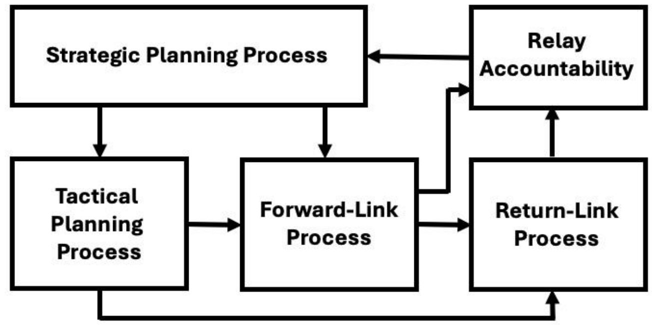

***Figure 6.1:** Summary View of the Relay Coordination Processes.*

6.4.1 The Strategic Planning Process

The Strategic Planning Process involves the definition of long-duration planning periods and the identification of overflights during which a relay session may occur. MaROS supports the negotiation, selection, and implementation of relay services between the lander and orbiter teams as part of this process.

To facilitate consolidated planning across all MRN orbiters for all relay service users, relay activities are planned in 2-week increments, called "planning periods". The Strategic Planning Process is active between 4 months to 1 week before the first relay session of a planning period, depending on the MRN orbiter involved, and officially completes at the bi-weekly SRRC meeting.

The Strategic Planning Process enables the relay service user operators to establish a baseline plan for the planning periods defined by the relay service provider operators, and follows the steps outlined in TSP Section 8. The IMRCWG facilitates the smooth execution of the Strategic Planning Process by providing the various forums (i.e., the meetings as in Section 6.3) for use by the relay community, while NASA's MEP provides the needed tools (e.g., MaROS, email lists, etc.).

Future Consideration: Demand Access Networking

Today's MRN is operated on two-week planning cycles, as described above. This cadence is an historic artifact of the nature of the MRN, which consists of orbiters whose primary purpose is their science investigations, with the provision of relay services being supported as an additional function.

Given the long and variable distance between Earth and Mars, these orbiters were each developed with operations paradigms that enable the spacecraft to self-operate for periods as long as 8 weeks. The current approach of performing consolidated planning across all the orbiters in two-week increments represents a compromise solution that enables the Mars lander teams to plan relay services across all the orbiters for a given two-week period, with roughly the same start and end dates for each orbiter. This required adaptation by both the relay orbiter operators, who naturally desire to command their spacecraft with long-duration sequences, and with the lander (especially rover) operators, who naturally desire to command their spacecraft every day to respond to changing on-the-ground conditions.

In the future, it is envisioned that relay orbiters may be implemented with "demand access" capabilities. Under this paradigm, relay orbiters may no longer need to be commanded with long-duration sequences to implement the desired relay services, but instead may be capable of accepting and supporting unscheduled relay service requests from relay service users. The operators of today's MRN orbiters support this notion for "emergency" relay sessions with varying levels of responsiveness.

However, as long as the available relay services remain over-constrained, it is expected that the human-involved planning approach as described here will remain necessary to ensure the fair and equitable distribution of the scarce resources of the network.

6.4.2 The Tactical Planning Process

The Tactical Planning Process involves managing any changes to relay plans that were established during the Strategic Planning Process, usually after command products have already been generated by the participants to implement those plans. This process is active beginning at the completion of the SRRC Meeting and continues through real time. MaROS supports the negotiation and implementation of short-term changes, typically in response to unexpected events within the network. The ability to respond to these requests is restricted by the relay service provider operators and varies widely by orbiter. Examples of allowed changes may include:

- Adjustments to the timing of a relay session.
- Adjustments to relay session parameters, such as any pre-selected data rates or encoding.
- The addition of a new relay session due to a spacecraft emergency on a relay service user spacecraft.
- The cancellation of a previously scheduled relay session due to an operational change or a spacecraft emergency.

In practice, tactical changes require unplanned effort by all affected parties and are therefore discouraged.

The Tactical Planning Process generally follows the steps outlined in TSP Section 8. In the case of changes initiated by the relay service provider operators (typically in response to a spacecraft emergency), there tends to be no recourse for the relay service user operators, and notification of changes is provided by the relay service provider operators as a courtesy and to facilitate situational awareness.

All tactical data exchanges are via MaROS. MaROS calculates a forward-link latency called the "last nominal uplink time" (LNUT) for every relay session (or is otherwise provided this information by the relay service provider operators), which represents a deadline for requesting, negotiating, and securing a tactical change.

As with the Strategic Planning Process, the IMRCWG facilitates the smooth execution of the Tactical Planning Process by providing the various forums (i.e., the meetings as in Section 6.3) for use by the relay community, while NASA's MEP provides the needed tools (e.g. MaROS, email lists, etc.).

Future Consideration: Automatic Radiation from Deep Space Tracking Networks

In today's MRN, nearly all DFE commanding to the spacecraft at Mars include either human-involved or human-assisted processes. Historically, the inclusion of humans in the DFE chain has been shown to increase the risk of command errors even while it provides a safety net to protect against them [20]. Regardless, the inclusion of humans in this process introduces additional latency when transferring data through the relay network. The removal of humans from this process is expected to be a necessary step prior to the implementation of a truly autonomous SSI.

6.4.3 The Forward-Link Process

The Forward-Link Process involves the transfer of relay service user data and other similar data products via the MRN in the forward-link direction. Forward-link data exchanges can occur via MaROS, which can act as an intermediary to transfer this data from the relay service user's MOC to the relay service provider's MOC for packaging and transmission to the relay service provider spacecraft, in anticipation of its eventual delivery to the relay service user spacecraft (refer to Figure 1.2 to review the simplified network topology). MaROS may be bypassed for intra-agency forward-link data transfers, as in the example of ESA's ERCO where forward-link data intended for an ESA lander is transferred via an ESA orbiter, as described in TSP Section 9.

Like the Tactical Planning Process, the identified LNUT for a relay session represents the last opportunity for forward-link data to be inserted into the network prior to the designated relay session during which the forward-link data is intended to be transferred.

This process represents the primary means for relay service users to submit forward-link data to their relay service user spacecraft via a relay service provider. For MaROS-supported forward-link data transfers, this process generally follows the steps outlined in TSP Section 8.

Once forward-link data has been submitted into the network, whether via MaROS or otherwise, there is assumed to be no means to intercept the data in transit to the relay service user spacecraft. The relay service provider operators assume responsibility for correctly receiving this data (confirming data integrity at every transfer between systems) and transferring it in a timely way via deep space tracking networks and the relay service provider spacecraft to the relay service user spacecraft.

It is possible for forward-link data to not be completely delivered to the relay service user spacecraft on the first attempt, though efforts by all are made to ensure a successful delivery. When this happens, the current paradigm is to have the relay service user operators re-submit the forward-link data for transfer during another relay session, thereby starting this process over again.

NASA's MEP facilitates the smooth execution of the Forward-Link Process by providing the various tools (e.g. MaROS, email lists, etc.) for use by the relay community.

Future Consideration: Forward-Link Data Handling

In today's MRN, most forward-link data transfers are designated by the Mars rover operators to be transmitted to the rovers during specific relay sessions. This ensures that the forward-link data, which usually contains rover command data, are completely received by the rover prior to when these commands are expected to execute.

Occasionally, the rover teams need to transmit very large command products to the rovers. These large products usually include flight software updates that are significantly larger than the usual command products. These often cannot be transmitted to the relay orbiter from the deep space tracking networks during a single DFE communications session and also typically cannot be transmitted to the rovers during a single relay session. However, they tend to not need to be received by the rover by any specific time, which grants some flexibility.

In this scenario, the rover operators fragment these very large command products into several smaller products, sized specifically to be transferrable during likely transmission windows from the deep space tracking networks and during likely relay sessions. The rover operators then transfer these fragments via MaROS to the orbiter operators with instructions to transfer the forward-link data on a best-efforts basis over as many relay sessions as necessary, until all the data is successfully transmitted to the rovers and can there be reassembled into the original large command product.

In all cases the orbiter operators are unaware of the contents of the forward-link data. They receive the data via MaROS and manage each data product as a distinct data product that must be transferred completely, in order, and without corruption – delivering on time to the rover exactly what was received via MaROS. The forward-link data may contain many different data types that are handled in a variety of ways onboard the rover, and the orbiter operators remain unaware of this.

In the future, there are several technologies that may make these forward-link transfers simpler. The use of the Bundle Protocol (BP) [21], for example, could be used to automatically fragment forward-link data products. This protocol includes support for the transfer of useful meta-data such as data type and priority, supports encryption, and could enable custody transfer schemes. The use of an overlying protocol such as the Licklider Transmission Protocol (LTP) [22] or the CCSDS File Delivery Protocol (CFDP) [23], which has already been used at Mars in a partial implementation, can be used to facilitate the reliable transfer of these bundles in a complete manner.

The use of these protocols would likely reduce the ability of the rover operators to control when and how their data is transferred to their rover (which may not be desired by the rover operators), but it would manage the over-sized forward-link data problem natively.

6.4.4 The Return-Link Process

The Return-Link Process involves the transfer of relay service user data and other similar data products via the MRN in the return-link direction. MaROS may act as an intermediary to reliably transfer this data from the relay service provider's MOC to the relay service user's MOC after it has been received on Earth. The objective of this process is to ensure the complete transfer of return-link data to the relay service user operators as fast as possible.

This process generally follows the steps outlined in TSP Section 8, though an additional pathway is also available, as described in the "Topology Exception" box below.

The relay service provider operators assume responsibility for executing the designated relay session (as planned in the Strategic Planning Process and as modified, if applicable, via the Tactical Planning Process) to receive return-link data from the relay service user spacecraft, transfer it to the relay service provider MOC via deep space tracking networks, and subsequently deliver it to the relay service user MOC.

MaROS models and estimates the end-to-end time-of-transfer (i.e., "return-link latencies") (or is otherwise provided this information by the relay service provider operators) to inform the relay service user operators when they should expect their return-link data to be available in their MOC.

It is possible for return-link data from a given relay session to not be completely delivered to the relay service user operators on the first attempt, though efforts by all are made to ensure a successful delivery. When this happens, the current paradigm is to have the relay service provider operators recover as much of the data as possible by retransmitting the data from the relay service provider spacecraft, if appropriate, or otherwise retrieving or reprocessing the data from various parts of the ground system, including from deep space tracking network(s), as available.

The relay service provider is only responsible for delivering the data it has actually received during a relay session, which may represent partial data products as transmitted by the relay service user spacecraft. Should the relay service user operators need to complete a partial product, or if data is permanently lost by the relay service provider in transit, the relay service user operators must re-send the return-link data from the relay service user spacecraft during another relay session, if desired, thereby starting this process over again.

NASA's MEP facilitates the smooth execution of the Return-Link Process by providing the various tools (e.g. MaROS, email lists, etc.) for use by the relay community.

Topology Exception: NASA's DSN as an Intermediary in the Return-Link Process

The description included above regarding how return-link data is transferred represents a methodology that conforms to the topology shown in Figure 1.2. However, in today's MRN, the actual pathway for delivering return-link data is quite different.

For NASA's relay orbiters, the DSN parses the orbiters' data streams as they are being received, recognizes the presence of relay data, and extracts the return-link data from them. The DSN then directly transfers any return-link data to the rover MOCs. For ESA's relay orbiters, regardless of which deep space tracking network is used, the entirety of the orbiters' data is transferred to ESOC where it is de-packaged and then delivered to a server from which the DSN extracts the return-link data and delivers it directly to the rover MOCs.

This DSN-provided service persists the early architecture developed for the MRN in the early 2000s. It is notable that this approach is topologically consistent with other architectures proposed for future capabilities, specifically the use of DTN where BP would be used network-wide and would allow for the direct routing of data. This approach does require the DSN to be able to parse the contents of an orbiter's data stream, which may be undesirable for new relay service provider operators; and also mandates an additional interface between the DSN and the relay service user operators to receive their return-link data.

For new participants in the MRN, either approach may be pursued, namely this DSN-supported service or the MaROS-based approach as described above. The MaROS-based approach would be especially useful for cases where the DSN is not included in the end-to-end data transfer chain, or for cases where the transmission from the relay service provider orbiters is encrypted.

Future Consideration: Return-Link Data Handling

In today's MRN, return-link data is received by the relay orbiters from the Mars rovers as discrete frames of configurable size, consistent with the Proximity-1 protocol. These frames usually represent portions of a larger data product as sourced from the Mars rovers. During relay sessions, the Proximity-1 protocol is typically configured to ensure the reliable transfer of this data using a "go-back-N" scheme to confirm the contiguous transfer of data. As the exact amount of data that may be transmitted during any given relay session is uncertain, the Mars rovers are configured to keep transmitting data (with a maintained backlog or with fill data) until the relay session terminates, either per the defined relay session schedule or due to other interruptions, such as terrain occlusions. This typically results in the very last data product in the transmission being received incompletely.

The rover operators therefore have implemented systems to track what data products are generated by the rover and confirm when those data products are completely received in the rover MOC. If a partial data product is received, the rover operators then must decide to retransmit the remainder of the data product (or, more typically, the product in its entirety) from the rover during a future relay session ("try again"), allow the data to be overwritten onboard the rover ("timing out"), or actively delete the product from the rover memory ("giving up"). In any case, the recovery of a partial data product that was never completely received by a relay orbiter is the responsibility of the rover operators.

In the future, there are several technologies that may make these return-link transfers simpler. Again, the use of BP, for example, as discussed in the "Forward-Link Data Handling" box above, could be used to automatically fragment return-link data products. This protocol includes support for the transfer of useful meta-data such as data type and priority, supports encryption, and could enable custody transfer schemes. The use of an overlying protocol for transfer reliability, such as LTP or CFDP, can be used to facilitate the reliable transfer of these bundles in a complete manner.

The use of these protocols would likely reduce the ability of the rover operators to control when and how their data is transferred from their rover (which may not be desired by the rover operators), but it would manage the return-link fragmentation problem natively.

6.4.5 The Relay Accountability Process

The Relay Accountability Process involves the monitoring of the network's overall performance. MRN participants report radio and other network performance data to MaROS, which acts as a centralized repository for this data. This reporting is typically done for every relay session and from various locations in the end-to-end data return path as return-link data transits through the network. This data is used to assess, monitor, and trend the overall performance of the network, enabling the MRN participants to address fundamental interconnectivity and interoperability concerns between the various spacecraft in the network.

This process generally follows the steps outlined in TSP Section 8.

The accountability data may include information such as the total number of bits transferred during a relay session, signal strengths experienced, and success indicators for whether all forward-link data were transferred to the relay service user. It may also include reports on when return-link data is transferred through each node in the network until it is received by the relay service user operators. Timely provision of this data can be helpful to provide situational awareness of where data is in the network, such as per the WMRD service (see TSP Section 6).

Both partners in a relay session are expected to generate this type of data and should therefore report what they can. Should any relay data be lost or corrupted during relay sessions or if relay sessions perform unexpectedly, a comparison of both relay participants' reports can facilitate the identification of where in the network problems might have occurred and suggest corrective actions.

It is possible for the engineering data related to a relay session to be received by the relevant spacecraft operators even before any return-link data from that relay session comes available.

Publication of this accountability data is expected to occur as soon as reasonably possible, but typically no later than the next working day after the data is received by the spacecraft operators. It is acceptable for publications of partial data at early times to be followed by more complete reporting when the full set of accountability data comes available.

This data is made available to all participants in the network via MaROS, both by viewing it in the MaROS GUI (including charting capabilities, see TSP Section 6) or by extracting it from the MaROS database via the CLI for offline review, processing, and charting.

NASA's MEP facilitates the smooth execution of the Relay Accountability Process by providing the various tools (e.g., MaROS, email lists, etc.) for use by the relay community. The IMRCWG actively monitors the health of the network using this data and may pose questions about unanticipated network behaviors (see Section 6.5).

6.5 Emergency Support and Other Unanticipated Network Behaviors

A spacecraft experiences an "emergency" when the health or safety of the spacecraft is at risk. Emergency relay support may be requested during the following two scenarios:

1. A relay service user spacecraft experiences an emergency and necessitates changes to the nominally scheduled relay plan.
2. A relay service provider spacecraft experiences an emergency, and the operators of the relay service user spacecraft determine they require additional support from other relay service providers to compensate for lost relay sessions.

In practice, both scenarios are very rarely exercised. Today, the Mars rovers schedule a variety of relay sessions across the relay orbiters, making them less sensitive to a failure in any one of them, and are generally unable to respond to emergencies with sufficient agility to make the scheduling of new relay sessions useful.

Because these emergencies often occur in near real-time, the responses might need to be handled differently from nominal changes conducted via the standard Tactical Planning Process. To facilitate interchanges between participants, an MRN-wide email distribution list and contact information for the management of each relay participant is made available to all participants.

At times, the performance of the network may not be as expected. Underperformance might be caused by a spacecraft issue, EMI, terrain occlusions, reflectivity, etc. Any MRN participant may initiate an investigation into these situations and ask for investigative support from affected parties.

6.6 Relay Service Process Testing

The relay coordination processes described here each have their own timelines for execution. New participants may wish to practice them as their capabilities mature and as they approach active operations. The successful testing of these processes typically relies upon the successful testing of the flight and ground interfaces (see Sections 3.4, 4.4, and 5.5) and therefore has historically been the last testing performed prior to a new spacecraft actively participating in the MRN. Ultimate responsibility remains with the new participant to ensure they can perform all necessary relay coordination and planning functions as described here.

In practice, a new MRN participant typically requests a series of "operational readiness tests" or "operations simulations" during which these processes are practiced by operations personnel using the pre-launch flight (or flight-like) software and hardware, as needed. The objective of these tests is to confirm that new participants can properly engage within the MRN from an operations viewpoint, with a thorough understanding of their operations processes (including timelines), procedures, and data exchange mechanisms.

The nature of cross-spacecraft support for these and all other tests are typically outlined in the Service Agreements (see Section 2). The IMRCWG provides support to execute these tests, as needed.

7 Summary of MRN Participant Responsibilities

New MRN participants assume a variety of responsibilities, consistent with the nature of the MRN as an inherently collaborative venture. Participants, including both relay service providers and relay service users, assume these responsibilities for all their ground and flight elements. The summary presented here of these responsibilities is drawn from the earlier contents of this guide.

Administrative responsibilities that apply equally to relay service providers and to relay service users include:

- Engage with the IMRCWG to develop all needed agreements and other documentation.
- Remain in compliance with all documented agreements.
- Perform testing of all flight and ground interfaces prior to direct participation in the operational MRN.
- Confirm an ability to engage with all the relay planning processes prior to direct participation in the operational MRN.

Operations responsibilities that apply specifically to relay service providers include:

- Operate the relay service provider spacecraft and related ground systems.
- Perform all spacecraft navigation, especially for making position information available for DWE communications, for relay planning purposes, and for collision avoidance monitoring.
- Perform all spacecraft attitude control activities and specifically make spacecraft attitude information available for relay planning purposes.
- Define periods of time when relay services are unavailable for use, as necessary.
- Secure deep space tracking time and specifically make this and related information (such as expected DWE data rates) available for relay planning and operations.
- Accept forward-link data from the operators of a relay service user spacecraft and commit to transferring this data through the network as received and without corruption, within the predicted latency.
- Deliver return-link data from a relay service user spacecraft to the operators of the spacecraft as received and without corruption, within the predicted latency.

Operations responsibilities that apply specifically to relay service users include:

- Operate the relay service user spacecraft and related ground systems.
- Perform all spacecraft navigation, especially for making position information available for relay planning purposes and for collision avoidance monitoring, as applicable.
- Perform all spacecraft attitude control activities, as applicable, and specifically make spacecraft attitude information available for relay planning purposes.

- Provide forward-link data to the operators of a relay service provider spacecraft on a timeline consistent with published deadlines.
- Receive return-link data from a relay service provider as it is made available.

Operations responsibilities that apply equally to relay service providers and to relay service users include:

- Participate in good faith in all planning and coordination activities to schedule relay services.
- Implement scheduled relay services as planned and actively communicate deviations from the planned activities to affected parties.
- Monitor the health and performance of applicable flight and ground systems and report performance data related to the provision of relay services in a timely manner.
- Be responsive and proactive during nominal operations and during unexpected situations.

The successful operation of the network relies upon participants meeting these responsibilities [24]. The IMRCWG and its membership may aid in the performance of these responsibilities through a variety of means.

8 Summary of Additional Resources Provided to New MRN Participants

After establishing a formal agreement to join the MRN, additional documentation may be provided to new MRN participants. These documents include more detailed information about the nature and operation of the MRN and are especially important during operations.

- Access to existing agreements and guidance on newly required documentation, as outlined in Section 2, as appropriate.
- The Mars Relay Network Handbook: This document describes the MRN relay service capabilities available for use by future relay service user spacecraft. It presents key specifications of each of the MRN orbiters in a common format, at a level suitable to support user mission planning and analysis. The primary focus of this handbook is on the proximity link for Mars relay services.
- The MRN Planning Approach Document: This document describes in greater detail the relay coordination processes outlined in Section 6.4, spanning all network-level activities regardless of institutional ownership of the respective spacecraft.
- The Mars Mid-Range Schedule: This document outlines the comprehensive schedule listing milestones and due dates for the relay coordination processes. This is updated regularly in support of ongoing operations.
- MaROS Software Interface Specifications (SISs): These document the XML data structure for each file interface type (see TSP Section 7).
- User guides for MaROS and related systems, describing in detail how users can gain access to these systems and how each of these systems are used in the relay operations workflow. In the case of the MaROS user guide, it also includes descriptions of how forward- and return-link latencies are calculated, the nature of identifiable relay planning conflicts, etc.
- Invitations to MRN coordination meetings, especially IMRCWG and SRRC meetings, as appropriate, and access to other communications mechanisms (see Section 6.3).

References

[1] Wenkert, D., e. a. (May 2016). Enabling International Data Relay at Mars. AIAA SpaceOps 2016, Daejeon, Korea.

[2] Edwards, C. (October 2004). A Martian Telecommunications Network: UHF Relay Support of the Mars Exploration Rovers by the Mars Global Surveyor, Mars Odyssey, and Mars Express Orbiters. International Astronautic Congress, IAC-04-M.5.07, IAC-04-M.5.07, Vancouver, British Columbia, Canada.

[3] Gladden, R., e. a. (May 2025, forthcoming). The History of the Use of Orbiter Relay at Mars, 1972 to 2024. SpaceOps 2025. Montréal, Québec, Canada.

[4] Gladden, R. (March 2022). Preparing the Mars Relay Network for the Arrival of the Perseverance Rover at Mars. IEEE Aerospace Conference, Big Sky, Montana, USA.

[5] Krasner, S. M., e. a. (September 2021). Reconstruction of Entry, Descent, and Landing Communications for the InSight Mars Lander. Journal of Spacecraft and Rockets, DOI: 10.2514/1.A34892.

[6] Elliott, H., e. a. (March 2024). Spacecraft Position Determination with the Electra UHF Transceiver: A Comparison of One-Way and Two-Way Doppler Observations. IEEE Aerospace Conference, Big Sky, Montana, USA.

[7] CCSDS (July 2014). Solar System Internetwork (SSI) Architecture. CCSDS 730.1-G-1.

[8] NASA/ESA (November 2017). International Mars Relay Coordination Working Group Terms of Reference. Signed jointly by NASA and ESA.

[9] CCSDS (July 2013). Proximity-1 Space Link Protocol – Physical Layer. CCSDS 211.1-B-4.

[10] CCSDS (October 2019). Proximity-1 Space Link Protocol – Coding and Synchronization Sublayer. CCSDS 211.2-B-3.

[11] CCSDS (July 2020). Proximity-1 Space Link Protocol – Data Link Layer. CCSDS 211.0-B-6.

[12] CCSDS (December 2021). Proximity-1 Space Link Protocol – Rationale, Architecture, and Scenarios. CCSDS 230.1-G-2.

[13] CCSDS (June 2024). Unified Space Data Link Protocol. CCSDS 732.1-B-3.

[14] The Future Mars Communications Architecture, a report of the Interagency Operations Advisory Group Mars and Beyond Communications Architecture Working Group, 22 February 2022.

[15] LunaNet Interoperability Specification Document (LNIS), Version 4, issued jointly by NASA and ESA on 12 September 2022. Version 5 is known to be forthcoming.

[16] Grami, A. (2016). Time Division Multiple Access. In Introduction to Digital Communications (pp. 457-491). Academic Press.

[17] Gladden, R.E., e. a. (April 2010). Mars Relay Operations Service: Rationale and Approach. SpaceOps 2010, Huntsville, Alabama, USA.

[18] Sauer, B. (May 2025, forthcoming). Mars Relay Operations Service (MaROS): Adapting to the Needs of the Mars Relay Network. SpaceOps 2025. Montréal, Québec, Canada.

[19] CCSDS (February 2021). Cross Support Service Management – Common Data Entities. CCSDS 902.12-M-1.

[20] Meshkat, L., e. a. (January 2013). Managing the Risk of Command File Errors. 51st AIAA Aerospace Sciences Meeting, Grapevine, Texas, USA.

[21] CCSDS (September 2015). Bundle Protocol Specification. CCSDS 734.2-B-1.

[22] CCSDS (May 2015). Licklider Transmission Protocol (LTP) for CCSDS. CCSDS 734.1-B-1.

[23] CCSDS (July 2020). File Delivery Protocol (CFDP). CCSDS 727.0-B-5.

[24] Gladden, R. (March 2024). Lessons Learned from the Mars Relay Network: Considerations for Future Relay Networks. IEEE Aerospace Conference, Big Sky, Montana, USA.

Acronyms

BER – bit error rate
BP – Bundle Protocol, a CCSDS standard
Caltech – The California Institute of Technology
CCSDS – Consultative Committee for Space Data Systems
CFDP – CCSDS File Delivery Protocol, a CCSDS standard
CGR – contact graph routing, a DTN-related technology
CLI – command-line interface
dB – decibels
DFE – direct-from-Earth
DSN – NASA's Deep Space Network
DTE – direct-to-Earth
DTN – delay (or disruption) tolerant networking
DWE – direct-with-Earth
E2E – end-to-end (system tests)
EDL – entry, descent, and landing
EMB – event message bus, a MaROS-provided capability
EMI – electromagnetic interference
ESA – European Space Agency
ERCO – European Relay Coordination Office
ESOC – ESA's European Space Operation Centre
ESTRACK – ESA's European Space Tracking network
FDMA – frequency division multiple access
GUI – graphical user interface
IDD – Interface Definition Document
IMEWG – International Mars Exploration Working Group
IMRCWG – International Mars Relay Coordination Working Group
IOAG – Interagency Operations Advisory Group
ITAR – International Traffic in Arms Regulations
JPL – Jet Propulsion Laboratory
LNUT – last-nominal uplink time, a calculated forward-link latency
LTP – Licklider Transmission Protocol, a CCSDS standard
MaROS – Mars Relay Operations Service
MaROS-Ops – the MaROS operations software environment
MaROS-TB – the MaROS testbed software environment
MAVEN – NASA's Mars Atmosphere and Volatile Evolution (orbiter)
MEP – NASA's Mars Exploration Program
MGSS – NASA's Multimission Ground Systems and Services (Program)
MOC – mission operations center
MRN – Mars Relay Network
MRO – NASA's Mars Reconnaissance Orbiter
MOA – Memorandum of Agreement
MOU – Memorandum of Understanding
MPX – Metrics Predict Executive, an SPS service
NASA – National Aeronautics and Space Administration (U.S.)
PRT – portable relay testbed
RF – radio frequency
RKPNI – Russian Complex of Receiving of Science Data
ROSCOSMOS – Russian State Corporation for Space Activities
RSP – relay service provider
RSU – relay service user
RTP – Relay Telecom Predictor, a MaROS service
SFCG – Space Frequency Coordination Group
SIS – software interface specifications
SNR – signal-to-noise ratio
SPS – Service Preparation Subsystem, a DSN-provided service
SRRC – Short-Range Relay Coordination (meeting)
SSI – solar system internetwork
TAA – Technical Assistance Agreement
TDMA – time-division multiple access
TGO – ESA's ExoMars Trace Gas Orbiter
TIM – technical interchange meeting
ToR – terms of reference
TSP – technical support package
U.S. – The United States of America
WMRD – Where's My Relay Data, a MaROS service
XML – extensible markup language