IOAG.T.RC.001.V1

Interagency Operations Advisory Group
Space Internetworking Strategy Group

Operations Concept for a Solar System Internetwork (SSI)

15 October 2010

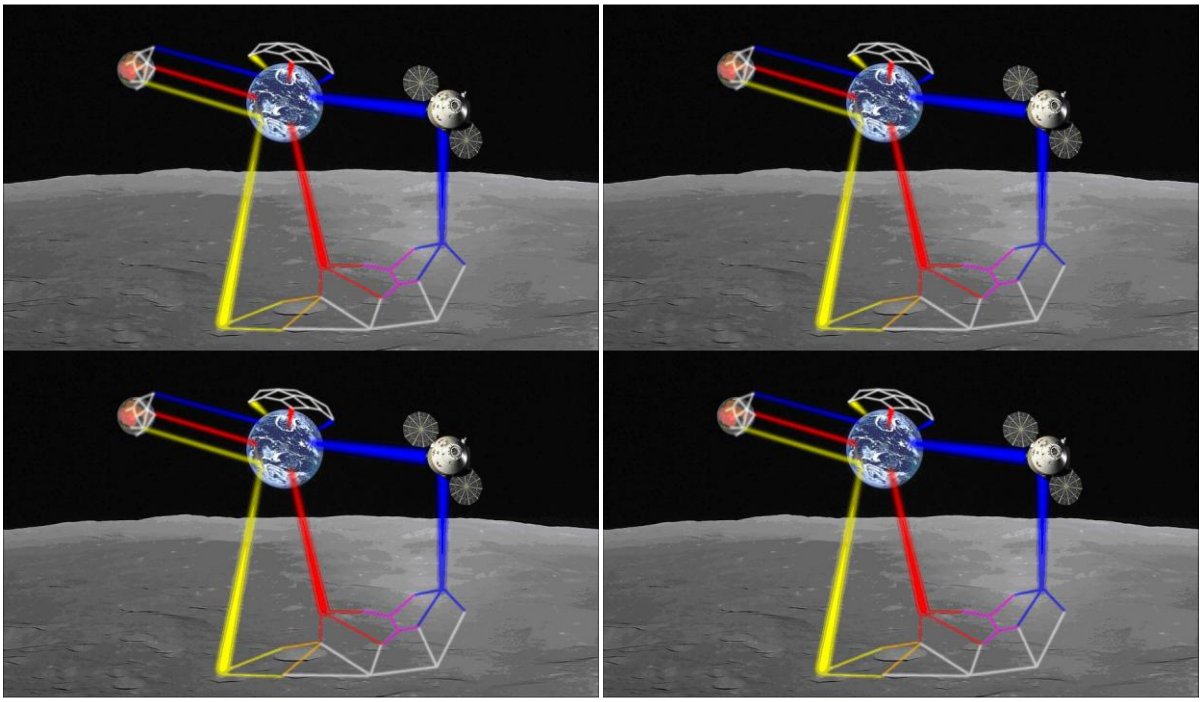

Cover art background: SELENE HDTV image provided by JAXA/NHK

Membership of the Interagency Operations Advisory Group (IOAG) Space Internetworking Strategy Group (SISG)

Co-chairmen:

Wolfgang Hell	ESA ESOC
John Rush	NASA HQ

Members:

Francois Allard	ESA ESTEC
Erik Barkley	NASA JPL
Lena Braatz**	NASA/Booz Allen Hamilton
Fred Brosi	NASA/GST
Scott Burleigh	NASA JPL
Gian Paolo Calzolari	ESA ESOC
Vint Cerf	NASA/Google
Matthew Cosby	UK Space Agency/QinetiQ
Michel Denis*	ESA ESOC
Robert Durst	NASA/Mitre
Chad Edwards*	NASA JPL
Roy Gladden	NASA JPL
Adrian Hooke	NASA HQ
Dave Israel	NASA GSFC
Narita Kaneaki	JAXA
Greg Kazz	NASA JPL
Mike Kearney	NASA MSFC
Paolo Maldari	ESA ESOC (ret.)
Jane Marquart	NASA GSFC
Gilles Moury	CNES
John Pietras	NASA/GST
Martin Pilgram	DLR
Luca Salotti	ASI
James Schier	NASA HQ
Michael Schmidt	ESA ESOC
Klaus-Juergen Schulz	ESA ESOC
Keith Scott	NASA/Mitre
Peter Shames	NASA JPL
Jason Soloff	NASA JSC
Wallace Tai	NASA JPL
Chris Taylor	ESA ESTEC
Takahiro Yamada	JAXA

* Co-author
** Editor

Table of Contents

Figures

Executive Summary

To date, most space communication scenarios have involved simple point-to-point links between a spacecraft and Earth. Today's link-layer protocols and carefully planned and scheduled link operations have been adequate to meet the needs of past missions. However, future space exploration concepts will introduce much more complex topologies, with data transfer over multiple hops via relay spacecraft and other intermediate nodes. Such scenarios call for the implementation of a true Solar System Internetwork (SSI), with the introduction of a network layer in the space communications protocol stack to provide reliable routing and forwarding of data across multiple links.

Like the terrestrial Internet, the SSI will offer users a well-defined, standardized platform upon which to build a wide variety of applications by accessing end-to-end network services. Also like the terrestrial Internet, once aggregate bandwidth needs are satisfied, the user has no need to plan and manage the details of the end-to-end data path; intermediate SSI provider nodes will forward user data to their end destination based on the specified destination network address and knowledge of the overall network topology. However, the SSI will need to address the unique aspects of space communication. For scenarios in which light times are short and a permanent end-to-end communication path exists, conventional Internet Protocol (IP)-based protocols may suffice; more generally, however, the SSI will need to utilize the Disruption Tolerant Networking (DTN) protocol suite, which can be used in any scenario, including those with longer light times or where link connectivity may be intermittent.

The SSI will provide a solid foundation for implementing a wide range of user applications, such as file transfer, messaging, voice and video streaming, network timing services, and radio metric tracking. The SSI network layer will provide the end-to-end data transfer services supporting these applications. Consequently, the user application itself will only be exposed at the endpoints of the network service; all intermediate nodes will only need to operate at the network layer and below. (This situation is in contrast to today's nascent planetary relay scenarios, in which intermediate relay nodes implement their store-and-forward functionality in ad hoc, customized, application-layer software.)

Exploiting information about the overall network topology, the SSI will enable more efficient use of communication link resources by allowing user data to flow over multiple network paths as they become available, and by seamlessly integrating the flow of data from multiple users across the network. Quality of service parameters associated with each network service request will support prioritized allocation of network resources when necessary, and knowledge of alternative network paths will allow rapid, automated recovery from isolated network anomalies.

A set of SSI network management activities will support the operations of the SSI network layer. Key elements include: a) management of the SSI network address space, b) determination of the SSI contact plan (i.e., the temporal connectivity of the overall network), c) dissemination of the SSI contact plan to all participating SSI nodes, d) monitoring of SSI network performance, and e) identification and resolution of any detected network anomaly.

The SSI is envisioned to be a confederation of independent, cooperating space agencies. Individual agencies will act as SSI service providers for their user nodes, based on the capabilities of their provider nodes. In addition, cross-support services will be negotiated via peering agreements, defining how one agency can access another agency’s SSI network elements.

Based on the anticipated trend toward more complex space networking scenarios, and on the clear benefits of improved bandwidth utilization, simplified user operations, standardized network interfaces, and more straightforward interagency cross support, the IOAG’s SISG recommends a timely evolution toward a fully operational SSI.

1 Introduction

1.1 Background

The vast majority of space communication scenarios to date have involved relatively simple single-hop links between a spacecraft and its ground-based mission operations center. For such scenarios, straightforward link-layer functionality is sufficient to support end-to-end data transfer and a relatively mature set of link-layer and cross-support standards have been developed under the aegis of the Consultative Committee for Space Data Systems (CCSDS) and are routinely used in flight operations. With these simple point-to-point topologies, there is little need for a network layer in the overall communication protocol stack.

However, future space mission scenarios are envisioned with much richer topologies involving multiple spacecraft and with data flowing across multiple hops and over multiple paths to achieve end-to-end data transfers. Such scenarios will demand the implementation of a functional SSI network layer to support routing of data over this more complex topology. Like the terrestrial Internet, such an architecture will support simple network interfaces for users, including provision of end-to-end network data transfer services; however, when required it will also address the intermittent connectivity characteristic of space missions and the long light-time delays inherent in interplanetary links.

In fact, relay communication is already being used today operationally at Mars to great advantage. The Mars Exploration Rovers, Spirit and Opportunity, have returned 98% of their data via ultrahigh frequency (UHF) relays through Odyssey, Mars Global Surveyor, Mars Express (MEX), and the Mars Reconnaissance Orbiter (MRO)—far more than could have been returned on the rovers' X-band direct-to-Earth links. The 2007 Phoenix (PHX) Lander mission dispensed completely with direct-to-Earth X-band links on its landed spacecraft, utilizing only UHF communications via Odyssey, MRO, and MEX to reduce mass, power, and cost. However, these data relay operations, while very effective, are at this point still highly idiosyncratic, with mission-unique, store-and-forward, application-layer functions implemented directly above the link layer. Routing of data is based on private methods and ground-based sequencing, without any functional network layer. Ultimately, this approach drives implementation and operations cost and complexity and cannot scale efficiently to more complex network scenarios.

The IOAG anticipates a wide range of future space exploration concepts that will exploit network communications, motivating a rapid transition to a true Solar System Internetwork. Examples include:

- **Mars exploration**: The coming decade of Mars exploration will involve a range of large in situ rovers, ranging in size from Mars Exploration Rover (MER)-class up to Mars Science Laboratory (MSL)-class, with high-rate payload suites including stereoscopic panoramic imaging, hyperspectral imaging, microscopic imaging, and High Definition video. Spatially distributed networks of small, stationary landers may collaborate to perform geological and meteorological investigations. In the 2020s, a collection of spacecraft, including one or more landers and rovers, a Mars Ascent Vehicle, and an Earth return vehicle, may cooperate to return the first Martian samples to Earth. Supporting these various users will be an orbiting relay infrastructure consisting of multiple relay orbiters flown by multiple space agencies, all adopting a common end-to-end networking approach. Some of these orbiters may be hybrid science/telecommunications orbiters, gathering remote sensing science observations while also offering SSI network services to other users; at some point this hybrid network will likely be augmented with dedicated relay orbiters, with payloads and orbits optimized for their telecommunications function.

- **Earth science**: Future Low Earth Orbit (LEO) scenarios may include high-rate science downlinks utilizing store-and-forward networking for rate buffering, both onboard a spacecraft and at Earth-based stations. Use of network protocols for rate buffering at ground stations will enable science file transfers and other mission applications to communicate between the mission and end user, while making the underlying network and ground station infrastructure transparent to user data transfers. By taking opportunistic advantage of available data transfer possibilities, data driven routing of the SSI may also enable low-rate science and emergency alert distribution between spacecraft, science operation centers, Mission Operations Centers (MOCs), and other locations. Use of protocols interoperable with terrestrial internetworking will enable sensor web applications that include space-based and non-space-based sensors. Multi-sensor missions with payloads on board satellites, planes, balloons, or on the ground will provide end users with connectivity to their payloads, while multi-satellite systems will conduct distributed science experiments from constellation or formation-flying scenarios using inter-satellite links and/or relay spacecraft services to route science and engineering data.

- **Human exploration**: Future human exploration scenarios will similarly require connectivity between multiple in situ elements, as well as back to Earth. These scenarios will also require higher data rates in all directions and highly reliable low-rate voice and data services. In the case of Lunar or other near-Earth exploration, near real-time interactions will be required, along with store-and-forward data transfers for those times when end-to-end connectivity is either not available or when all bandwidth is consumed by real-time services.

In this end state, we envision a rich network topology with data-driven routing controlling the flow of information between various network nodes. In addition to typical transfers of information between a spacecraft and its Earth control center, the SSI architecture will support direct exchange of information among spacecraft without Earth in the loop. For instance, two collocated landers/rovers may exchange information to support collaborative exploration, or a lander and orbiter may interact to coordinate in situ and orbital observations. Large data products may flow over multiple, intermittently connected paths, with network-layer protocols ensuring reliable end-to-end data transport. The availability of SSI services will minimize the effort and cost borne by users and relay service providers to plan and execute data transfer across the network. At the same time, however, the highly dynamic nature of space exploration will continue to demand user visibility into the quality of service measures (e.g., latency, total available bandwidth) predicted for end-to-end data flow.

Together, these future operational scenarios provide a compelling case for development of an SSI with standardized network services and interfaces for all users and service providers.

1.2 Progress

Following studies on future communications architectures by several agencies, the IOAG discussed the need to evolve interoperability into space at its IOAG#11 meeting, held in Cebreros, Spain in June 2007. At that meeting, the IOAG passed a resolution to form the SISG to provide recommendations to the IOAG concerning a strategy for internetworking in space. Co-chaired by the National Aeronautics and Space Administration (NASA) and the European Space Agency (ESA), the SISG also includes representatives from the Japan Aerospace Exploration Agency (JAXA), the German Aerospace Center (DLR), the French space agency (CNES), and the Italian space agency (ASI). The SISG's initial task was to provide the IOAG a report by spring 2009 that:

- Identified concepts of internetworking operations and associated users scenarios
- Analyzed profiles of existing CCSDS standards necessary to implement future space internet working
- Provided a roadmap for how agencies will evolve to support new capabilities

The SISG completed its initial report[i] in July 2008, and the findings were presented and discussed at IOAG#12, held the following September at DLR-Oberpfaffenhofen, and also at the Interoperability Plenary (IOP-2) held in December 2008 in Geneva, Switzerland.

The IOP-2 endorsed the SISG report, concluding the first phase of the SISG's work, and asked the SISG to start a second phase of work to define a proposal for an Internetworking Architecture and associated Operations Concept. A communiqué recording the results of IOP-2[ii] states:

> *The IOAG's Space Internetworking Strategy Group (SISG) should formalize a draft Solar System Internetwork (SSI) Operations Concept and candidate architectural definition in time for IOAG-13 and should prepare a mature architectural proposal for review and endorsement at the third Inter-Operability Plenary meeting (IOP-3). At that time, the IOAG is requested to present an enhanced service catalogue for endorsement. The IOP Agencies should ensure representation from their programs and projects to work with SISG to identify potential missions which may benefit from adoption of the SSI-related standards, leading to a gradual build up of in-space and ground-based space internetworking infrastructure.*

1.3 Purpose of this Operations Concept

This document is the SISG's response to the request by the IOP-2 to formalize an SSI Operations Concept (i.e., operations based on end-to-end services transferred between end users via routed communications). To formulate the concepts described in this document, the SISG conducted several studies to examine and resolve technical issues related to SSI implementation[iii]. This Operations Concept is fully consistent with the SISG's SSI technical issue study results and provides the top-level definition of SSI operations, including references to the elements and services that will be defined further in a separate SSI architecture document.

The focus of the SSI is interoperability of space system elements built and operated by different space agencies. This Operations Concept provides a high-level description of the SSI elements, including their types, locations, and interactions. It describes the activities involved in planning and executing SSI services and monitoring and reporting of SSI functions.

This Operations Concept proposes a framework for:

- Identification of required SSI functions and activity scenarios (which may lead to the definition of new requirements and impact the overall SSI architecture)
- Identification of the network elements and functionalities required/available to support missions in the 2015-2025+ timeframe

1.4 Document Overview

This document consists of several sections plus appendices:

- Section 1 provides background information
- Section 2 states the purpose of the SSI, defines the SSI elements and functions, and explains the evolution to the SSI end state
- Section 3 contains SSI principles
- Section 4 describes the SSI network management processes
- Section 5 describes views from the perspectives of the user and provider
- Section 6 presents the functional deployment aspects of the SSI in a simple case study
- Section 7 describes SSI network operations scenarios
- Appendix A provides an acronym list
- Appendix B lists referenced documents

2 General Description of the SSI

2.1 SSI Scope and Purpose

The purpose of the SSI is the provision of internetworked data communications services across the Solar System in support of space mission users. The SSI will employ a confederated and interoperable infrastructure of nodes owned and operated by multiple space agencies to achieve a level of service that individual agencies would otherwise be unlikely to achieve.

Fundamental to the end goal of the SSI concept is that all participants in the confederation will expose standard and agreed cross-support services at the network layer of the International Standards Organization (ISO)/Open Systems Interconnection (OSI) Reference Model. Similar cross-supported services for network monitoring and accounting will also be essential.

The scope of the SSI will consider constraints such as communications links and visibility, complex interactions between user and provider to deliver interoperability, large delays in the data transfer paths, and the need to provide appropriate qualities of service for data, reporting, capability, and management of widely distributed assets.

2.2 SSI Composition

The SSI provides communications services among nodes in an interplanetary network. An SSI node is defined as any network entity that can serve as a source or destination of information at the network layer. All SSI nodes may exchange information with each other at the network layer, but end-to-end performance will depend upon the physical location of and connections among the nodes and the individual node capabilities (throughput, bandwidth, and the length and frequency of contact opportunities).

All SSI nodes will use the DTN Bundle Protocols (BP) suite and/or the internet protocol (IP) suite as the basis for network-layer connectivity. SSI nodes can use the BP suite to communicate in any scenario (including those in which all communicating nodes are continuously connected with low-delay, and those in which they are intermittently connected and/or include long-delay paths). Use of the internet protocol of the IP suite, however, is restricted to scenarios in which the communicating nodes are well connected (i.e., the network path between them is continuously connected and relatively low-delay). If an application on an SSI node must communicate between IP and BP domains, application-layer gateways or bridging/tunneling can be used. Use of application-layer gateways is not recommended because they often have to be application-specific (one gateway or pair of gateways per application) and they may introduce unintended interdependencies between environments. Bridging/tunneling between IP and BP is simpler to implement, but may not work for all applications. It is worth noting that the two networking suites (BP and IP) can be run concurrently over the same sets of underlying links, and that the BP suite can also run “on top of” the IP suite, using IP transport protocols as “links” between BP nodes. Running BP over IP allows easy BP connectivity across the terrestrial Internet (e.g., to connect a BP node at a ground station to BP node at a MOC).

The SSI is composed of two types of nodes:

- **User nodes** reside at the periphery of the SSI and have the ability to access SSI services, but do not have the capability to provide network-layer forwarding functionality (i.e., store-and-forward capability, potential support for multiple simultaneous links, etc.). User nodes can use SSI services, but cannot provide SSI services for other nodes. A terrestrial Internet analogy for an SSI user node is the computer belonging to an individual Internet user, who is accessing services offered by a local Internet Service Provider.

- **Provider nodes** have the ability to offer network-layer routing and forwarding capabilities. Such nodes can perform routing over one or more links, have store-and-forward capability, and may support multiple simultaneous links to enable simultaneous support to multiple user nodes and/or real-time forwarding functionality. These nodes act as intermediate relay nodes for end-to-end network services. Provider nodes offer standardized service interfaces to provide network functionality to SSI users. They can connect to any other peer network service provider nodes using standard peering protocols to exchange user data, as well as routing and other management information. When it is acting as the terminus of an SSI service, a provider node functions like a user node. Examples of a provider node acting as a user node include a hybrid science/telecommunications orbiter that is transmitting its own science telemetry, or a dedicated telecommunications orbiter that is sending or receiving data related to its own spacecraft operations.

Just as in the terrestrial Internet, user nodes must implement an agreed minimum set of standard protocols, and the SSI provider nodes must offer the agreed set of standard services for routing, staging, and forwarding data. There will be procedures for adding a new user node to (or deleting it from) the SSI, for requesting services, and for inserting a new SSI provider node into (or removing it from) the network.

2.3 Definition of the SSI Nodes

Note that for purposes of this document, Earth is treated separately from the other planets, and the term “planet” is extended to mean any other Solar System body (moon, asteroid, comet, etc.) that is being investigated.

Figure 2-1 shows the types of SSI nodes, each of which is described in detail in subsequent subsections. Lines between the SSI nodes indicate potential communications paths; actual paths will depend upon the location of the nodes and individual node capabilities.

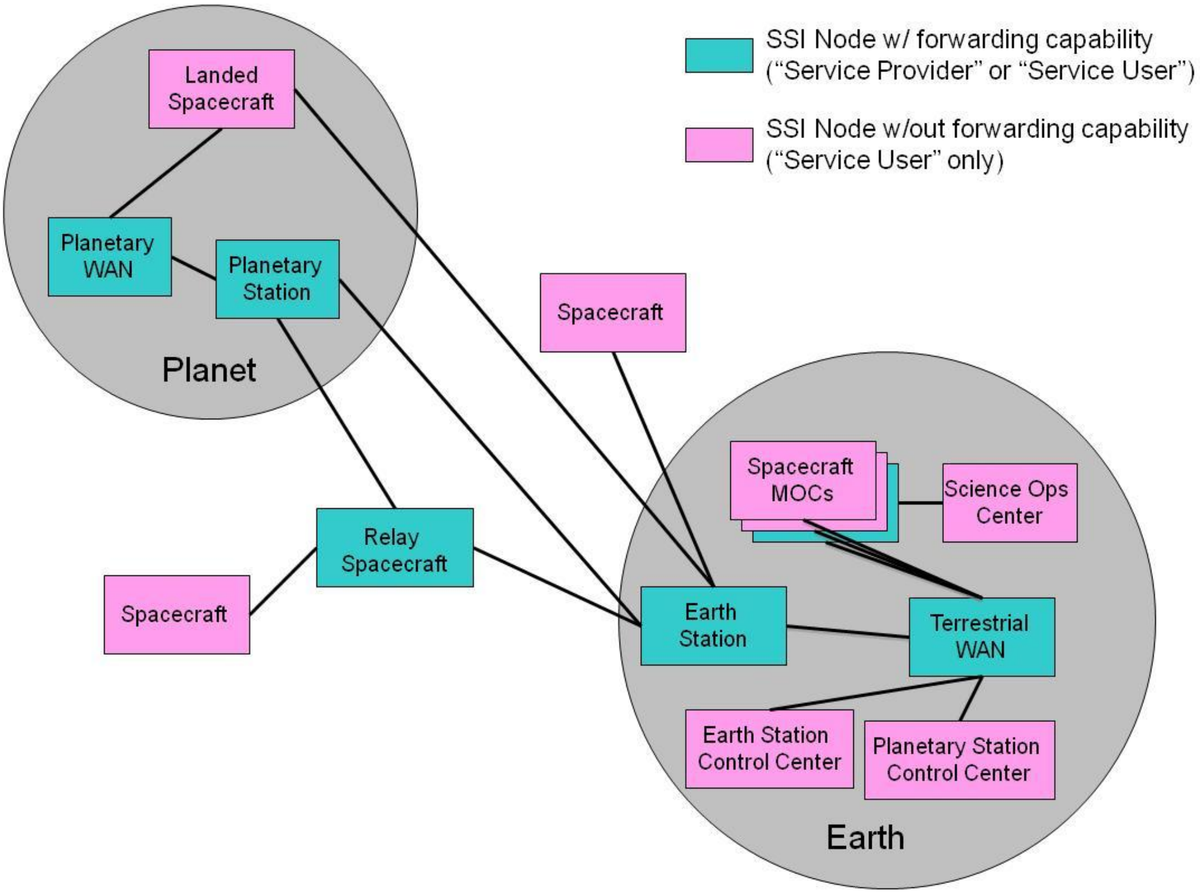

Figure 2-1: SSI Nodes

2.3.1 SSI Provider Nodes

SSI provider nodes are a set of interconnected elements that work together to provide end-to-end communications services for mission users (user nodes). It is the fundamental precept of the SSI that all provider nodes be equipped to use the internet protocol of the IP suite and/or DTN Bundle Protocols (BP) as the basis for network-layer connectivity, and that they route data in the form of IP datagrams and/or BP bundles. The following subsections describe types of SSI provider nodes.

2.3.1.1 Earth Station (e.g., Earth-space link terminals or ground stations)

An Earth Station provides network-layer connectivity over a space link to route traffic between a Terrestrial Wide Area Network (WAN) and space or planetary SSI nodes. In the context of the SSI, the Earth Station establishes and maintains link-layer and network connections with other provider nodes (e.g., Relay Spacecraft, Planetary Stations) and directly with user nodes (either user spacecraft or landed assets).

2.3.1.2 Relay Spacecraft

The Relay Spacecraft performs routing and storage functions over multiple space links (e.g., Advanced Orbital Systems [AOS]/Telemetry [TM]/Telecommand [TC], Proximity-1). In the context of the SSI, a Relay Spacecraft may establish space links with other provider nodes (e.g., Earth Stations, other Relay Spacecraft, Planetary Stations) and directly with user nodes (either user spacecraft or landed assets).

Some Relay Spacecraft will be Hybrid Relay Orbiters whose main mission may be the collection of science data, but which have the capability to serve as SSI provider nodes, as well. Depending on resource availability, season, geometry, flexibility of pointing and planning, Hybrid Relay Orbiters may be able to perform science-related activities and support the SSI provider node role. Often, however, the ability of Hybrid Relay Orbiters to support relay operations may be constrained due to science-driven orbits, limited resources, science-driven pointing constraints, body fixed or omni antennas for proximity operations, etc. Once their science operations requirements are satisfied, some Hybrid Relay Orbiters may be granted an extended mission, during which time they may serve as dedicated SSI provider nodes.

2.3.1.3 Planetary Station (e.g., planet-Space Link Terminal)

A Planetary Station is essentially the same as an Earth Station. It provides the space link over which traffic is routed between a Planetary WAN and space or terrestrial SSI nodes. In the context of the SSI, the Planetary Station establishes and maintains link-layer and network connectivity with other provider nodes (e.g., Relay Spacecraft, Earth Stations) and directly with user nodes (either user spacecraft or landed assets). An example of a Planetary Station would be a fixed terminal on the lunar surface at the site of a human outpost, providing communication services to proximate users, as well as dedicated high-rate links to lunar relay spacecraft and/or Earth Stations.

2.3.1.4 Terrestrial Wide Area Networks

A Terrestrial WAN is a network of nodes that provides network-layer connectivity among terrestrial SSI user nodes and Earth Stations to support the exchange of data.

2.3.1.5 Planetary Wide Area Networks

A Planetary WAN is essentially the same as a terrestrial WAN. It provides network-layer connectivity among planetary SSI user nodes and Planetary Stations to support the exchange of data.

2.3.2 SSI User Nodes

User nodes are the entities served by the SSI. As in a terrestrial network, any SSI user node may communicate with any other SSI node by routing data to its unique end-point address. One or more of the SSI provider nodes will enable such communications. Some user nodes perform functions necessary to the operation of the SSI (e.g., coordinating or operating SSI provider nodes), but are considered user nodes since they do not have the capability to provide full network-layer routing and forwarding functionality. Individual SSI user nodes will access the SSI using the internet protocol of the IP suite and/or the bundle protocol of the DTN suite.

The following subsections describe types of SSI user nodes.

2.3.2.1 Spacecraft

Spacecraft user nodes may be located either in space or landed on a planet (rover or lander). A Hybrid Relay Orbiter spacecraft acts as a user node when it transmits its own science telemetry or sends or receives data related to its own spacecraft operations. (This document discusses the case of a single user node on the spacecraft. In general, one could envision multiple user nodes on a spacecraft corresponding to individual payload and instrument processors.) For the purposes of this SSI Operations Concept, the notion of “spacecraft” encompasses the full range of potential user platforms including launch vehicles, orbiters, landers, airborne vehicles, etc.

2.3.2.2 Spacecraft Mission Operations Center

The Spacecraft MOC is in charge of operating one or more spacecraft missions (spaceborne or landed). The Spacecraft MOC uses the SSI to communicate with its spacecraft. User Spacecraft MOCs operate one or more user spacecraft, and Provider Spacecraft MOCs operate one or more Relay Spacecraft. A Provider Spacecraft MOC ensures that the internetworking and relay functions of the satellite(s) it is operating are available and that those satellites perform according to the service commitments made. In the case of a Hybrid Relay Orbiter, the Provider Spacecraft MOC might be a separate physical entity or it might be part of the mission’s User Spacecraft MOC. In some cases, a Spacecraft MOC may act as a provider node for a separate science operations center or science investigator (see Section 2.3.2.5 on the Science Operations Center).

2.3.2.3 Earth Station Control Center

An Earth Station Control Center will be responsible for preparing, controlling, and monitoring one or more Earth Stations and the ground-based communications related to those stations. An Earth Station Control Center might not be collocated with the Earth Station(s) it operates. The Earth Station Control Center will coordinate use of such stations, often supplying the service management interfaces used to plan, schedule, and request Earth Station services. Agencies coordinate with each other to provide cross-support services through their respective Earth Stations.

2.3.2.4 Planetary Station Control Center

A Planetary Station Control Center will be responsible for preparing, controlling, and monitoring one or more Planetary Stations and the planet-based communications related to those stations. A Planetary Station Control Center might not be collocated with the Planetary Station(s) it operates. The Planetary Station Control Center will coordinate use of such stations, often supplying the service management interfaces used to plan, schedule, and request Planetary Station services. Agencies coordinate with each other to provide cross-support services through their respective Planetary Stations.

2.3.2.5 Science Operations Center

A Science Operations Center may be responsible for controlling and monitoring individual science instruments, as well as the reception and possible distribution of science data. A Science Operations Center may not be collocated with the MOC that operates the spacecraft housing the instruments. At the SSI conceptual level, the Science Operations Center is seen as a terminal entity issuing payload requests and recovering the final science data. Whether this Science Operations Center is a dedicated institution in one physical place, or instead exists as a network of science institutes and principal investigators interfacing with the MOC is not relevant to the SSI operation.

2.4 SSI User Applications

The SSI will provide intercommunications that support user applications residing on SSI nodes. Examples of SSI user applications include:

- File transfer
- Messaging
- Voice/video streaming
- Timing services
- Radio metric tracking

These user applications will utilize the underlying SSI network to transfer data, and it is recognized that operations such as timing and radio metric tracking will require specialized link and physical layer activities that are below the network layer, and are therefore outside the scope of this operations concept document.[1]

The SSI will also facilitate a set of so-called "last-hop" services for situations in which a target node needs to be accessed directly at the link or physical layer without use of any network-layer protocol. Examples of such scenarios include emergency command/telemetry, support for legacy assets without network-layer functionality, or support to very simple nodes without networking capability. In these last-hop service scenarios, a standardized SSI network-layer service will be initiated from a user source node to a Delivery Agent application on a penultimate node (i.e., a node one hop away from the target node) to request the last-hop service, providing any additional ancillary information needed by that application. That standardized Delivery Agent application will then initiate a separate physical or link-layer connection with the target node to complete the last-hop service.

2.5 SSI Supporting Functions

Successful SSI operations will be dependent on a number of supporting functions. SSI network management functions will involve the suite of activities that enable connectivity at the various OSI layers involved in supporting network operations. Network management will entail the processes leading to a confirmed schedule of internode connectivity, including Earth station, relay, and planetary station nodes, as applicable, along with associated link-layer configurations that determine their underlying bandwidth capabilities. This temporal connectivity information will be key to supporting network routing functions. Related to

---

[1] For example, in an SSI environment with multiple physical links potentially available between Earth and a given spacecraft, radio tracking observables can, in principle, be formulated on each individual physical link. Such observables offer new opportunities in terms of the richness of available geometric tracking information, but also potentially present new challenges in terms of understanding the accuracy and error contributions of the radio metric tracking observables on each individual link. In the context of this operations concept, we limit our scope to the notion of a user application on a given SSI node, which can formulate radio metric tracking observables on physical links involving that node, and can subsequently transfer those observations as data over the SSI network.

this, management of the SSI network address space will include unambiguous identification of SSI network nodes and designation of remote destinations for network services.

To assess SSI network performance, accounting data (e.g., data volumes, latencies, etc.) will be gathered from individual nodes and monitored to ensure nominal behavior. In the event of network anomalies, diagnostic data will be retrieved from relevant nodes to troubleshoot and resolve the situation. Individual agencies will be responsible for acquiring such data from their own assets; mechanisms and formats for exchanging such information between agencies need to be developed.

Management of user nodes will be conducted by the agency that owns and operates the user node. The user will guarantee the technical compatibility of the user node mission interfaces, configuration, operations concept, and procedures with the capabilities offered by the SSI during the mission lifetime.

Similarly, provider node management will be conducted by the agency that owns and operates the provider node to ensure that the provider node is operated to provide the agreed services without interfering with the other duties of the asset (e.g., the scientific operations of a Hybrid Relay Orbiter). Provider node management will guarantee that services are available, as per commitment to the SSI and to the user, while respecting the asset’s technical constraints.

If the SSI provider nodes necessary to provide a certain service consist of assets of the same Agency, coordination of service for the users may be performed by the MOCs for the provider assets, or externally by an agency-level coordinating function. For a complex relay scenario involving multiple assets of several agencies, this coordination activity will have to be organized and performed on a confederated basis, taking into account multiple project and agency constraints.

2.6 Evolution of the SSI

The introduction of internetworking in space implies a significant change to the mission operations concept with respect to the conventional operations of past missions, which were based on direct point-to-point communications between visible end users, or proprietary store-and-forward applications on individual relay spacecraft. The end state concept of space internetworking is based on intermediate nodes forwarding data between source and destination nodes, using standard interoperable protocols. Today’s conventional CCSDS communications provide only link-layer, end-point addresses; the destination is fixed in the single-hop, point-to-point space link. The conventional CCSDS application data units (space packets or files) cannot be used directly as network-layer data structures, but must instead be transferred within routable, network-layer protocol data units (PDUs)—either DTN bundles or IP packets. Other application data units, such as messages, may also be transferred using the network-based transfer protocols.

The context of the SSI space internetworking end state encompasses the DTN bundle protocol suite (suitable for use in any scenario, including those that are delayed or intermittently connected) and/or the internet protocol of the IP suite (suitable for use only in delay-unaware, continuously connected scenarios). To evolve from the present style of mission operations to the SSI end state, these new internetworking constructs need to be inserted into the ground and space assets that provide services, and into the user nodes that will use these services. This evolution may initially involve static routing and manual configuration of the space assets, but the end goal is for space internetworking to be supported by a higher degree of automation, which in principle may even include dynamic routing—and therefore dynamic reconfiguration—onboard. SSI evolution

- Constitutes a migration towards fully automated space internetwork (DTN- and IP-based), and will familiarize the user and provider missions with interoperable operations
- Responds to the needs of potential user missions planned for launch in the 2015-2020 timeframe
- Allows testing and validation of the Internetworking concept before it is fully adopted in its final configuration for future flying missions

This operations concept describes the SSI end state and, where appropriate, provides guidelines regarding the evolutionary phase toward that end state to ensure maximum coherence and continuity between the point-to-point mission operations of the past and the intended network of the future.

3 SSI Principles

The SISG has identified a set of management, planning, and execution principles that identify the roles of the different elements in the SSI and how they are intended to operate and evolve. These operational principles reflect both lessons learned from the terrestrial Internet, as well as experience with the unique aspects of space communications.

3.1 Management Principles

MA—1 SCOPE: End-to-end cross support shall be the goal of the SSI, where cross support embraces what is known today as ground cross support and space cross support.

MA—2 AGREEMENTS: Cross support is ruled by high-level documents agreed between the user agency and the provider agency(ies); typically a Memorandum of Understanding (MOU) and a service agreement. Lower level interface control documents (ICDs) are developed as necessary.

MA—3 ASSET RESPONSIBILITY: Each agency is responsible for the planning, control, and operations of its own assets, including payload configuration (e.g., for cross support).

MA—4 NETWORK RESPONSIBILITY: Each agency shall contribute interoperable elements to the overall functioning of the network in support of user services.

MA—5 COMMUNICATIONS PROTOCOLS: The operations concept shall, as far as possible, be independent of the communications protocols below the network layer to offer stable guidelines robust to the risks inherent to space missions (financial, delays, technical difficulties) and to ensure it is not bound too much to the specific implementations of such protocols.

MA—6 VALIDATION: An end-to-end validation shall occur between the definition/setup of the network and the start of its operational usage (e.g., via “demo passes”) or following major failure and recovery of network elements. Validation shall also be performed upon addition or removal of significant elements of the network.

MA—7 SECURITY: A trusted level of security/confidentiality must be agreed in advance and respected between the provider and the user. Security is required for the exchange of information and files leading to the establishment of the mission timeline, the configuration of sessions and services, and in the exchange of data to be transmitted to the user node via the provider nodes.

MA—8 ADDRESSING: Asset addressing must be constructed and managed at the network level using network-layer identifiers (IP address, DTN node name) and also existing link-layer elements (Spacecraft Identification [SCID], Multiplexer Access Point [MAP], etc.), along with the identifiers required to allow the intended routing and preclude non-intended communication. The SSI will require management and maintenance of the SSI address space.

MA—9 NETWORK SERVICES: The defining characteristic of the SSI end state operations concept is the presence of a functional network layer in the protocol stack. Application layer functionality is only present at the endpoints of an end-to-end service (e.g., CCSDS File

Delivery Protocol [CFDP], Asynchronous Message Services [AMS]); all processing at intermediate nodes is performed at the network layer.

MA—10 -INTEROPERABILITY: To provide a functional network layer in the protocol stack, all nodes that agree to provide SSI services will offer an agreed set of interoperable IP and/or DTN protocol services.

MA—11 ADDRESS SPACE MANAGEMENT: End state SSI operation requires interagency management of the network address space(s).

3.2 Planning Principles

PL—1 OVERALL PLANNING: The planning entities of the provider and user agencies must coordinate at long-term (typically geometry or flight dynamics), medium-term (typically resources or mission planning), and short-term (typically service request and delivery) levels. The development of international standards for exchange of planning information shall be promoted.

PL—2 NETWORK PLANNING: Network planning shall compile the routing possibilities and the loading level of affiliated provider assets. It shall answer user planning requests over the agreed time periods, and in case of conflict, propose feasible alternatives. It shall ensure that the finalized routing options are published to the users and providers for implementation.

PL—3 USER PLANNING: User planning shall produce and manage consolidated user requests in a manner globally compliant with the known capabilities of the network, as defined by the high-level agreements and applicable policies. Depending on the network planning feedback, user planning shall refine these requests with its end users at all envisaged planning cycles.

PL—4 TIMELINES: The network shall allow timely delivery, as required by the user, via managing the timing for delivery of the forward link product. Users will need to know the predicted epoch by which a given forward product will reach the destination node.

PL—5 PREDICTABILITY: It shall be possible to identify all provider components' latency and the resulting earliest/latest physical delivery times under normal conditions of SSI network operation. These physical limits may be variable within predefined parameters (distances between nodes, phases of the network deployment, etc.). Within these physical limits, the provider and user agencies shall agree which earliest/latest operational relay timings define the normal conditions of service; some flexibility may be introduced according to the criticality of the served mission's phase.

PL—6 RESOURCES: The SSI providers shall have visibility into the characteristics of the to-be-relayed data that have an impact on the resources of the provider nodes (like bandwidth required, priorities, data volume, pointing of relay spacecraft, consumption of fuel or other energy, manpower, etc.).

PL—7 RE-PLAN: The SSI shall allow re-planning or cancellation of a network service request prior to completion of the service or in event of an anomaly.

PL—8 ROUTINIZATION: After set-up/validation and outside critical or emergency phases, contact planning and relay services shall maximize compliance with the routine planning of each involved asset.

PL—9 CONTACT PLAN: Network planning and execution in the SSI end state hinge on a network contact plan, which establishes the temporal windows and communications capabilities (e.g., bandwidth) of individual node-to-node network links. A key element of SSI end state operations is the dissemination of all or a subset of the network contact plan to each individual network node. (The amount of information conveyed to each node is scalable; it could be limited to nearest-neighbor contact information or could entail full end-to-end network information. The more information a given node has, the better it can make routing decisions in terms of end-to-end service latency.)

PL—10 NETWORK PLANNING AND MANAGEMENT: In addition to standardized SSI network protocols, the SSI requires network planning and management functions to develop the network contact plan and execute network services.

PL—11 PEERING AGREEMENTS: Peering agreements will be used to implement interagency interfaces within the SSI. When planning mission communications, an individual user will arrange for service with its agency-level provider, who will, in turn, employ peering agreements to arrange end-to-end data flows that use different agencies' provider nodes.

3.3 Execution Principles

EX—1 SAFETY: The ultimate safety of the user/provider node is the responsibility of the user/provider agency. The user assets shall be robust to delivery problems from the network.

EX—2 USER CONTROL: The user MOC is responsible for the overall decision on when/whether to use the SSI network and what to do in case of a problem. The user MOC authorizes the sending of the forward products that shall transit via the SSI to the target asset and receives reporting accordingly.

EX—3 SSI PROVIDER NODE CONTROL: The MOC of the provider node is responsible for the management and scheduling of that SSI node, and determines what to do in case of a problem. The provider MOC manages the transfer of data and provides reporting accordingly.

EX—4 VERIFICATION: The SSI communications architecture and protocols shall allow the user MOC to verify proper execution of the data delivery operations, taking into account the delays resulting from physical constraints and operational latencies.

EX—5 MONITORING/REPORTING: The SSI providers shall provide the user MOC with feedback (at least at a high level to indicate what was done and what was not done) on the progress and success of the intermediate steps in the relaying process. If several paths are possible for the delivery, reporting on the path used is desired, as long as the delivery requirements (contents, timing) can be respected; when they cannot, the user MOC may expect first-level information on the location and impact of the problem.

EX—6 TRACEABILITY: It shall always be possible for the provider agency to accomplish the following for its own nodes: 1) successfully trace the items which have transited on a given node on Earth or in space in a certain period of time, 2) to report on delivery success or failure, and 3) to support redelivery of requested data for a certain (limited) time after the nominal transit date.

EX—7 ROBUSTNESS: Each node agreeing to participate as a service provider node in the network shall be robust to nominal operating variations, provided that the data stream respects the protocols in place, the data dimensions (volumes, transfer rates, transfer frequency), the service agreement, and the technical specification of the network and the node. Each agency and participating node has the right to abort sessions and discard data that are outside the agreed boundaries.

EX—8 TRANSPARENCY: Each node agreeing to participate in the network shall be indifferent and transparent to any contents of the transferred data units, provided that the data stream respects the protocols and the technical specification of the network. Each agency and participating node is expected to transfer data without examining the contents of the data.

EX—9 INTEGRITY: The SSI shall be capable of delivering complete, gap-free data products between any two nodes.

EX—10 USER EMERGENCY: The SSI shall allow for defining and using, under pre-agreed conditions, a path from the user MOC to the user node that is completely deterministic in geometry and timing (e.g., to recover from anomalies in the network and/or the user node).

EX—11 NETWORK EMERGENCY: An interagency alert system shall be supported. Cross operations (like interventions on user data in transit) shall remain exceptional, within a pre-agreed frame, and require close interaction between user and provider as for cases of emergency support.

EX—12 LAST-HOP SERVICE: The SSI shall facilitate a defined, cross-supported means for the user MOC to communicate with nodes that cannot implement SSI user node functions due to spacecraft emergency, or with legacy assets that are not capable of SSI functionality. SSI support of last-hop services will be achieved by establishing an SSI transfer from the user MOC to a standardized Delivery Agent application on the penultimate node, and that application will subsequently establish a link-layer or physical-layer connection with the user node.

EX—13 ROUTING FUNCTIONS: The SSI end state operations concept supports data flow over multiple possible network data paths. Forwarding of information is based on static or dynamic routing tables in the network-layer protocol with forwarding decisions based on information in the network-layer PDUs (as opposed to being driven by metadata or manually sequenced operations).

EX—14 DELAY-AWARE AND DELAY-UNAWARE OPERATIONS: The SSI end state will allow support for both delay-unaware operations (where short delay [<2 sec], continuous, end-to-end data paths are available) and delay-aware operations (where long transmission delays and multiple, temporally disjoint, piecewise network hops may be required to support end-to-end data flow).

EX—15 REPORTING: Individual network nodes will report status information in standardized formats to support network management and anomaly resolution.

4 Network Management Processes

4.1 SSI Coordination

SSI network management requires agreement by the management of each agency to adopt shared protocols, principles, norms, rules, decision-making procedures, and negotiating mechanisms governing the configuration and use of the SSI.

Nodes participating in the SSI may be owned and operated by different space agencies. It is expected that there will be no central management or ownership of the SSI and that participating agencies will follow established, documented, interagency agreements for configuring cross support. However, efficient long-term SSI operations will necessarily require increased automation of the mechanisms associated with negotiating and implementing these cross-support agreements to expedite service execution.

As shown in Figure 4-1, each agency will act as an "SSI Service Provider[2]" for its user nodes, offering SSI services based on the capabilities of its provider nodes. From a user perspective, this SSI Service Provider function is similar to the role of an Internet Service Provider (ISP) for terrestrial Internet users. Service agreements will document the SSI services that will be provided to the user based on its communications requirements across the life cycle of the mission.

In addition, agencies participating as providers in the SSI will reciprocally provide access to each other's nodes via a process known as peering. To be eligible, each participating SSI agency must offer peering services at the network layer of the OSI reference model.

To serve as a component of the SSI, an agency's provider nodes will offer a native network-layer routing service based on:

1. The bundle protocol of the DTN suite (suitable for use in any scenario), and/or
2. The internet protocol of the IP suite (suitable for use only in continuously connected, delay-unaware scenarios)

The peering points between different provider agencies' SSI nodes are controlled by interagency agreements.

---

[2] Based on this terrestrial ISP paradigm, the agency role of coordinating the provision of SSI services to user missions is sometimes referred to as a "Solar System Internetwork-Internet Service Provider," or SSI-ISP. A given agency may implement multiple SSI-ISP instantiations to accomplish activities in different solar system domains.

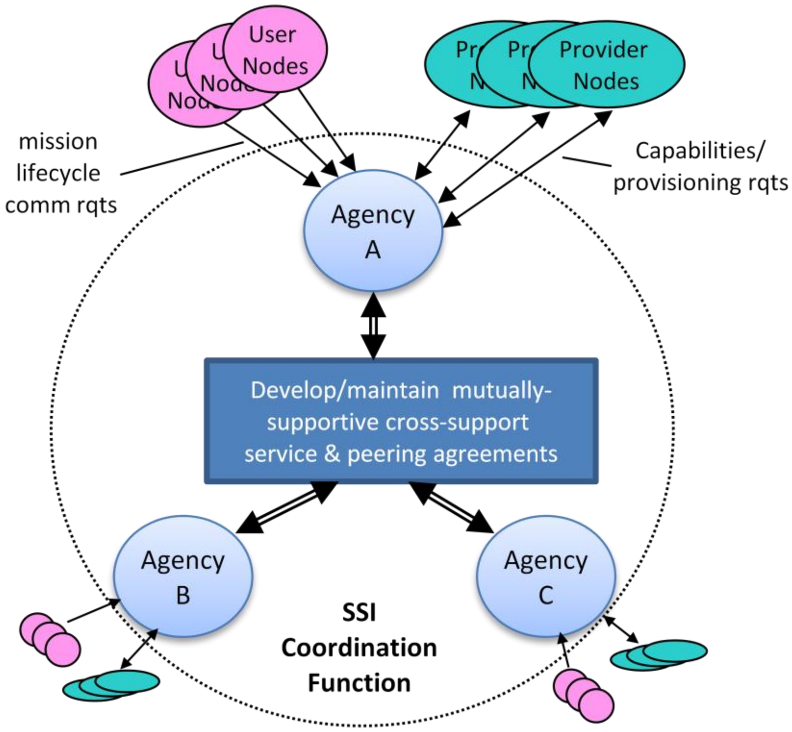

Figure 4-1: SSI Coordination

The confederation of participating agencies and the explicit set of user and provider nodes will evolve over time as new missions are launched and as operational missions are phased out or lost due to mission failure. Implementation of new provider nodes will be responsive to the needs of individual space missions or space mission campaigns involving multiple missions. In this way the overall configuration of the confederation and its participants will therefore change to match evolving international mission plans. On this relatively long-term time scale, service agreements and peering agreements will be updated to reflect changes in the capabilities of the evolving SSI. Participating agencies will document agreed primary mechanisms for offering and negotiating their participation in the SSI confederation, and for offering service commitments to their end users.

While terrestrial ISPs can typically absorb surges in demand or additional customers, SSI providers typically will not be able to do so without affecting other users of the network. SSI service agreements will therefore take into account the resource constraints of the SSI and the aggregate needs of SSI users.

All participating agencies must conform to a common naming and addressing scheme. Once defined, this scheme will be administered by the CCSDS's Space Addressing and Naming Authority (SANA), which will coordinate development of an internationally agreed mechanism for assigning and registering identifiers (names and numbers) necessary for configuration and operation of the SSI.

In addition, all participating agencies must implement a common security scheme at the network layer to manage access to the services and to provide confidentiality as needed for user data that is carried.

The SSI will be operated by different space agencies, requiring a level of integration based on architectural agreements that provide engineering coherence. These agreements will guide:

1. Procedures for assigning unambiguous identifiers (naming and addressing)
2. Methods that assure integrated, but private, data flows for multiple missions
3. Interoperable mechanisms for SSI network monitor and control
4. Collaborative scheduling procedures that maximize the utility of the confederated SSI, while respecting the terms of cross-support agreements among contributing agencies

In the evolutionary phase any “peering” agreements are expected to be documented on a mission-by-mission, interface-by-interface basis.

4.2 SSI Contact Planning Process

The need to schedule space links will not be eliminated by the evolution to an internetworked architecture. In fact, in the SSI architecture, management of the network contact plan becomes an essential aspect of operations. SSI contact planning involves determining mutual contact opportunities, producing an agreed-upon contact plan, disseminating the contact plan to participating SSI nodes, and monitoring and reporting on network utilization. Routing (determining the path to send the data) may be static (predetermined or managed) or it may entail contact graph routing, which involves understanding and exchanging the temporal aspects (contact times, one-way light time) and resources (e.g., storage and link capacity) for nodes participating in the SSI.

Since the provider node loading, service demands, communication opportunities, and communication geometry of the SSI are continually changing, contact planning will be a cyclical process, as illustrated in Figure 4-2.

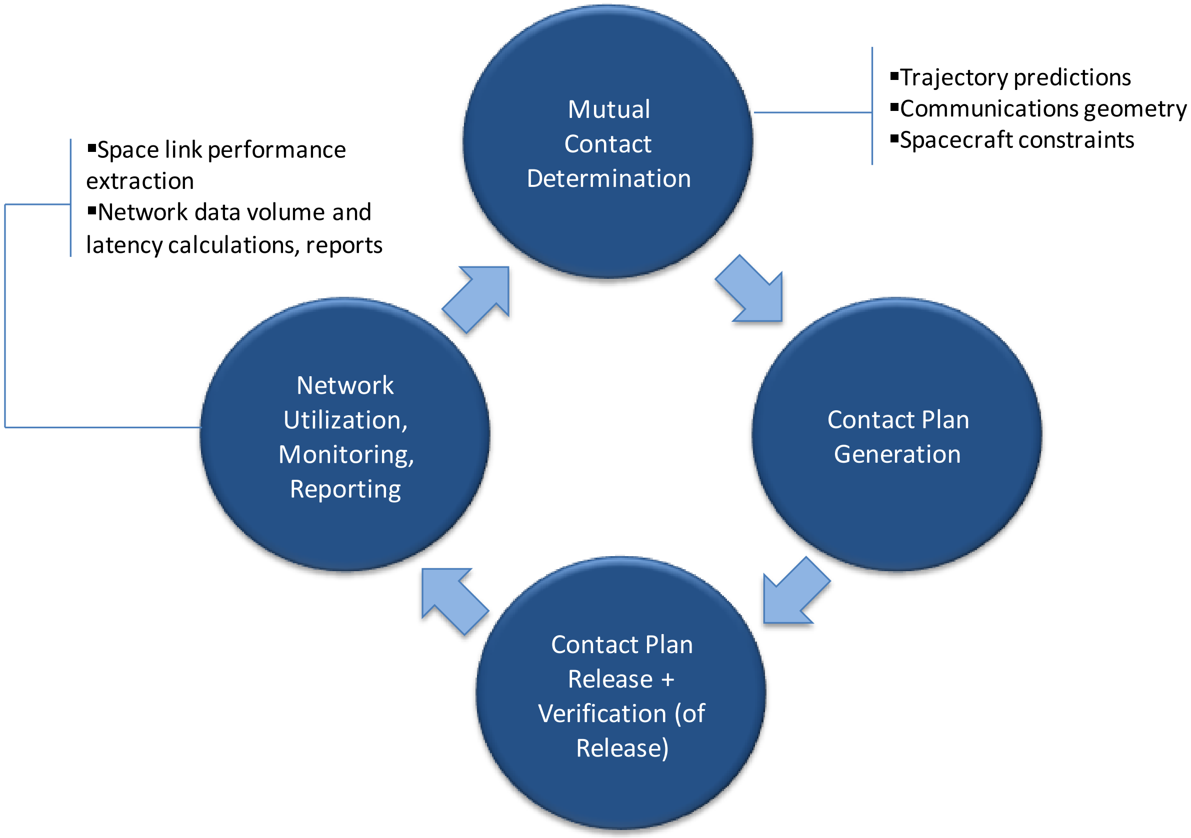

Figure 4-2: SSI Contact Planning Process Overview

During the evolutionary phase, the Provider MOC and User MOC will negotiate directly to conduct the contact planning process. Provider MOCs may have to negotiate with more than one User. Similarly, any given User MOC may require services from and need to negotiate with more than one Provider MOC. As more nodes are involved, scheduling and coordinating will become exponentially more difficult.

In the SSI environment, the traffic borne by a space link at any given time will no longer be constrained to a single spaceflight mission. Space links may be concurrently shared among missions. Contact Plans will be generated via a multiagency collaborative process, and will be used for real-time routing of IP packets and/or DTN bundles across a time-varying collection of links and nodes, but the establishment of those links and nodes is a prerequisite. Once these layer-2 links are established, all network data that can efficiently use the available link will be flowed toward its destination, subject to priorities and other policies.

4.2.1 Mutual Contact Determination

The first phase in configuring the SSI to support user missions is mutual contact determination. Trajectory prediction information, overlaid with spacecraft operational constraints that affect communications, will be used to determine mutual contact opportunities.

4.2.2 Contact Plan Generation

The second phase of the SSI contact planning process supports negotiation and selection of mutually agreed contact opportunities among the parties participating in a particular SSI configuration. The negotiations may be quite simple and straightforward for environments involving relatively short light time delays and/or a rich set of supporting nodes, but may be more involved for environments with long light time delays and/or a relatively sparse set of supporting nodes. In either case, a contact plan is produced as a result of the mutual agreements and is made available to all of the nodes participating in a particular SSI configuration.

A centralized scheduling process would likely yield a near-optimal solution (in a mathematical sense) to the space link scheduling problem. Using requests for connectivity among SSI nodes as input, the process would determine schedules for all space links in the SSI. Such a centralized process, however, would not preserve the relative autonomy of the SSI nodes belonging to different space agencies, and is therefore not practical for the foreseeable future.

Instead, the SSI will adopt a federated approach in which the scheduling of the various space links is distributed. An early version of such a shared-information scheduling process is currently being implemented in support of relay operations at Mars[iv]. To automate the process in a scalable fashion, more and different types of information will need to be available to and exchanged among the components of the SSI, for example:

- Additional/modified information in the schedule request submitted to SSI providers
- Additional planning information exchanged between users and providers
- Planning and scheduling information exchanged among SSI providers (e.g., for the purpose of load balancing on space links)
- Planning and scheduling information exchanged among SSI providers, and where required, the users (e.g., for the purpose of estimating aggregate loads on requested links and storage at intermediate nodes)
- Relative priorities of individual missions and mission phases

4.2.3 Contact Plan Release and Verification

After establishing the negotiated contact plan, those parties participating in a particular SSI network configuration take the necessary actions to implement it. For example, an Earth Station would compute its contact graph. A MOC would radiate the contact graph information to its participating spacecraft (potentially enabled by an earlier implemented contact plan), which would then perform a similar computation of the information into its contact graph. All parties would provide status about the contact plan implementation. This process is similar to the methods used to manage today's terrestrial Internet and update backbone routing information, except for the need to address intermittent connectivity and variable delays.

4.2.4 Network Utilization, Monitoring, and Reporting

With the contact plan disseminated, network operation in the SSI end state can proceed in a highly automated fashion, with user applications initiating network service requests as needed for inter-node data transfer, and with data flowing to destination node(s) based on the routing and forwarding capabilities of the network-layer components of the intermediate provider nodes.

The individual organizational entities responsible for the operation of each user or provider node will configure the telecommunication elements of that node to establish internode links in accordance with the contact plan.

The system will provide mechanisms for accountability to evaluate network and link-layer performance metrics relative to plans. Ongoing monitoring of the network will be required to confirm nominal SSI operations and/or to detect, isolate, and recover from network anomalies. At the network layer, this will involve each provider and user node generating statistics on transmitted and received bundles and on current bundle storage (for a DTN network scenario) or IP packet statistics (for an IP scenario). Standards for SSI network monitoring data and formats for reporting and disseminating such data have yet to be developed, but will be essential in facilitating the operation of a confederated multiagency network.

At a lower level, link-layer monitoring will provide information on the performance of individual links between pairs of nodes. For each link, both participating nodes will generate monitoring data, which will include link-layer statistics, e.g., Proximity-1 frame counts, frame error rates, and retransmission request counts for a relay link. Here again, standardized reporting formats for such data are needed to support efficient interagency operations.

5 SSI User and Provider Views

5.1 SSI User View

SSI user node operations will be focused on the interface to the SSI provider and the request and delivery of SSI network services between two spatially separated user nodes based on the internet protocol of the IP suite and/or the bundle protocol of the DTN suite. Users will be able to specify various quality of service parameters at the time of a network service request (e.g., data priority and/or required end-to-end latency). While the user will not manage the detailed node-by-node data flow across the network, the user will require status information on the end-to-end service.

From the point of view of an SSI user, the visible parts of the SSI are the network service request interfaces and the network service delivery interfaces.

End-to-end application services, such as reliable file and message delivery and other user application services, will be constructed in a layered fashion on top of standard networking services. These upper-layer services can take advantage of the end-to-end transfer, routing, forwarding, and reliability services provided by the network layer. In a similar fashion the network services are accomplished by the hop-by-hop delivery of data across the individual communication links along the end-to-end path.

Service delivery is initiated when the user sends any network-layer data, along with appropriate routing and addressing information, to the nearest routing/forwarding node. For IP traffic the data will be forwarded immediately, obeying the IP routing rules. For DTN/BP traffic the data may be routed immediately if a path is available, or they may be staged in intermediate storage until a suitable route is available.

5.2 SSI Provider View

SSI provider nodes are characterized by a functional network layer using the DTN BP suite (which is suitable for use in any scenario, including those that are intermittently connected, delay-aware, and delay-tolerant) and/or the internet protocol of the IP suite (which is suitable for use only in continuously connected, delay-unaware scenarios). SSI provider nodes will forward data units based on destination addresses, knowledge of potentially time-varying network topology and routing capabilities, and specific quality-of-service parameters associated with the user service request.

The operations view for the SSI service provider is centered on the execution of network service requests; data transfer; staging and forwarding, including fault detection, isolation, and recovery; and on the processes required to establish and disseminate the detailed network contact plan (i.e., the data base describing the potentially time-varying routing paths available across the network). The contact plan needs to reflect the bandwidth capability of each inter-node link, as well as any other relevant operational considerations, such as onboard storage limitations. Hybrid Relay Orbiter nodes must integrate performance of SSI network services into the broader mission timeline, and consider potential constraints imposed by the orbiter’s science activities.

From the SSI provider perspective, SSI user requests and contact plans are used in the SSI to plan the delivery of network services and to develop routing information that defines when and where there will be communication contacts to support data delivery. Routing information may be static (programmed into the SSI nodes “by management,” i.e., outside any standard network management protocols) or it may use contact plans that allow routes to be dynamically calculated based on priority of data and availability of links and node capacity. Regardless of the means used, SSI provider nodes will transfer routing information and other control and reporting information as part of a peer-to-peer network control process. These activities are done as “in-band” transfers between peer provider nodes in the SSI, and the methods and protocols are agreed to in advance by established peering agreements between agencies and missions.

These SSI peer elements may be on the Earth, on the surface of another planetary body, or in space, either free flying or in orbit around another planet or “hub” location, such as a Lagrange Point. Once the peering agreements are in place and communications links have been established, SSI traffic can flow as long as there are adequate link and provider node resources available. Each of the SSI provider nodes that is handling DTN/BP traffic must be able to stage data in intermediate storage until a suitable route is available, but it may route data immediately if a suitable path is available.

6 SSI Functional Deployment – A Case Study

This section describes the functional deployment of the SSI for a reference case in the end state. The descriptions are cast in the frame of a principle scenario and the activities and interactions of the set of network elements in that scenario. In the SSI end state, all nodes will have space internetworking functionality. Manual processes will be largely automated, and autonomous protocol execution will provide reliable data delivery. User requests for management of priorities will be provided together with contingency services for spacecraft in anomalous conditions.

6.1 Reference Scenario

Initial SSI deployments may be simple; for instance, the SSI reference scenario (depicted in Figure 6-1) is defined as:

- A spacecraft landed on another planetary body (e.g., Mars) and its MOC on Earth (Agency A)
- A Relay Spacecraft and its MOC on Earth (Agency A)
- A Relay Spacecraft and its MOC on Earth (Agency B)
- An Earth Station and its control center (Agency C)

For the purposes of this discussion, the reference scenario will consider communications from the landed spacecraft to/from its MOC only (communications between the Relay Spacecraft and their MOCs and the Earth Station and its control center are not discussed).

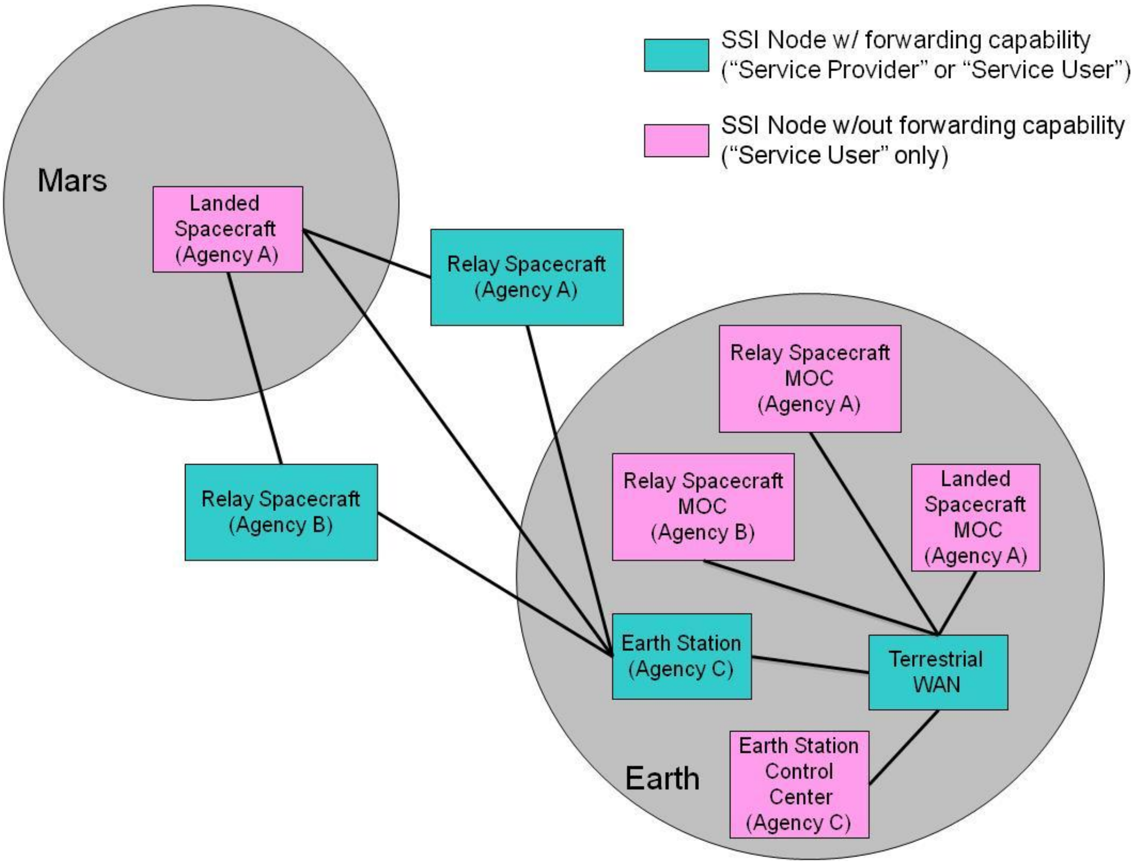

Figure 6-1: SSI Reference Scenario

6.2 Protocols

Figure 6-2 illustrates the general concept for the protocols involved in the SSI. The user applications exist only at the end nodes; only space internetworking (DTN or IP) is deployed at the intermediate nodes. Each of the intermediate nodes provides uniform and standardized functions for routing, storing, and forwarding data. In the reference scenario, the end nodes are the landed spacecraft and its MOC. Intermediate nodes include the Relay Spacecraft, the Earth station, and the terrestrial WAN.

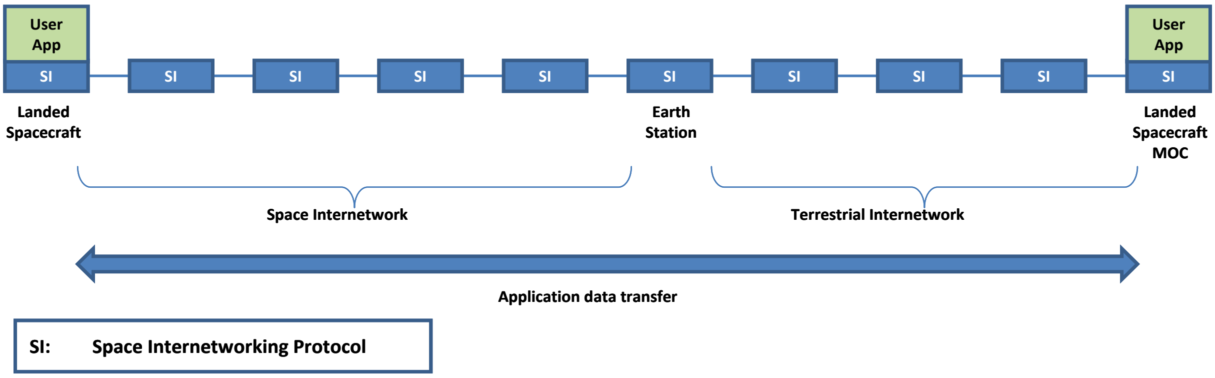

Figure 6-2: SSI Protocol Chain between End-user Applications

6.3 SSI User Applications

User applications include file, message, and other data transfer over DTN or IP. They are transparent to the SSI. The SSI will also support communications between users and standardized "last-hop" user applications.

6.4 Acquisition and Contact

The acquisition of the link between two SSI nodes (provider or user), and in particular between a landed spacecraft and a Relay Spacecraft, may be pre-programmed or use on-demand requests. The acquisition is carried out autonomously in accordance with the pre-negotiated contact plan, without near real-time intervention. Either node may send a hailing signal to indicate the start of the pass and initiate acquisition.

When establishing the link between the Relay Spacecraft and its MOC, the Relay Spacecraft MOC will use service management requests to initiate the space link, configure, manage, and monitor it. Once the space link is established, SSI traffic may flow directly from nodes that have staged data to the Relay Spacecraft.

The Relay Spacecraft will forward data from the landed spacecraft on the deep space link to an Earth Station. Due to orbital geometry or the capability of the Relay Spacecraft, the deep space link may not be contemporaneous with the link between the lander and Relay Spacecraft, in which case data will be stored on the Relay Spacecraft until the next available downlink opportunity. The network layer, typically based on DTN, will be used to address the intermittent link connectivity and the acknowledgement process in the presence of long, round-trip light times. (Today these are handled as a mixture of application and link-layer functions.)

A reverse process occurs to accomplish communications from the landed spacecraft MOC to the landed spacecraft.

End-to-end relay operations will be programmed to operate autonomously. A priori decisions of the agencies will ensure availability and coherency of the physical elements (e.g., frequencies, pointing, bandwidth, storage, etc.) and logical elements involved (e.g., authorisations, priorities). Relay Spacecraft nodes will be designed to accommodate the expected data and service requests of the lander(s) they are designed to support.

6.5 Node Capabilities

User nodes will initiate end-to-end data transfers. Data will be staged as DTN bundles for forwarding on the Relay Spacecraft, and will be distinguishable from data belonging to other relay sessions. Intermediate provider nodes will be able to store data until the next hop becomes accessible (possibly with a time out). Provider nodes will not reassemble user data (e.g., files) at each hop and may transfer user data in pieces, possibly via different relay paths. Only the end user node will reassemble the user data. Users may force a set of user data to be delivered fully assembled by sending the data in a single bundle.

Nodes will use DTN protocols to identify the next node, establish the communications with the next node, and transfer the data to that node. Nodes may employ static routing, or they may use Contact Graph Routing, which can adapt to available paths.

Nodes may verify data on receipt, may securely transfer data (with Bundle Security Protocol [BSP] or via user encryption), or may delete data if requested.

A data handling system is provided at each provider node for bundle data staging. Individual bundles will specify priorities, and provider nodes will forward bundles based on established policies. A user or provider (under certain defined conditions) will be able to delete a well defined set of user data (e.g., all bundles associated with the identified user data).

Link-layer services and protocols will support transfer of data over individual hops. The DTN Bundle Agents and upper-layer applications will manage reliable end-to-end delivery and signalling.

6.6 Emergency Services (Last Hop)

In support of emergency situations where the landed spacecraft no longer has network functionality, last-hop services will be employed to establish physical or link-layer command and telemetry service to/from the landed spacecraft and a standardized Delivery Agent application resident on the Relay Spacecraft.

To implement this emergency service, packets and other low-level frame/link transmissions must be supported as a point-to-point transmission across the “last-hop” proximity link. To provide this service, the User MOC of the mission must embed the required link-layer data structures, packets, or frames into files, along with the necessary link configuration and delivery information, and the files will be handled by the SSI as usual, until they are delivered to a standardized Delivery Agent application running on the Relay Spacecraft. This Delivery Agent accepts the data to be delivered and the instructions, and performs the necessary link configuration and data delivery services at the requested time.

Between the Relay Spacecraft and the landed spacecraft, Proximity-1 will support the direct transmission of space packets, but frames and other data structures (e.g., Bose-Chaudhuri-Hocquenghem [BCH] encoded TC frames) may also be transferred over reliable bit stream (User Defined Data [UDD]). The delivery instructions state how the link is to be configured, how the data are to be extracted and sent (packets, frames), when the data are to be sent, how often, and under what conditions this transmission is to be terminated. A similar return service is also implemented by an application on the Relay Spacecraft. As with the forward service, this application accepts a service request that instructs it how to configure the proximity link, what data to capture, and when to record the data. The resulting data set is placed in a file, along with a report of what was done and its success or failure, and this information is sent via SSI services back to the user. These return services may deliver essential telemetry, open loop sampled data from the Radio Frequency (RF) link itself, and timing or radiometric data from Proximity-1.

7 Network Operations Scenarios

SSI network operations involve a range of activities with multiple participants acting over multiple time frames. One set of activities takes place well in advance of the execution of network services, and is involved with the generation and dissemination of the network contact plan, using the processes described in Section 4. Figure 7-1 illustrates the contact planning process, with a focus on the various participants that must interact to establish a confederated contact plan that spans multiple user and provider nodes and may in general involve assets from multiple agencies. While individual robotic nodes may ultimately interact directly with the contact planning process, initially it is envisioned that a User MOC will provide contact planning inputs for both the User Spacecraft Node and the User Spacecraft MOC node. These inputs will include specific requests for contact opportunities for the Spacecraft and MOC nodes, as well as specific temporal constraints that would preclude contacts with other nodes. These inputs to the contact planning process will incorporate contributions from mission planning, navigation, flight system analysis, and science planning.

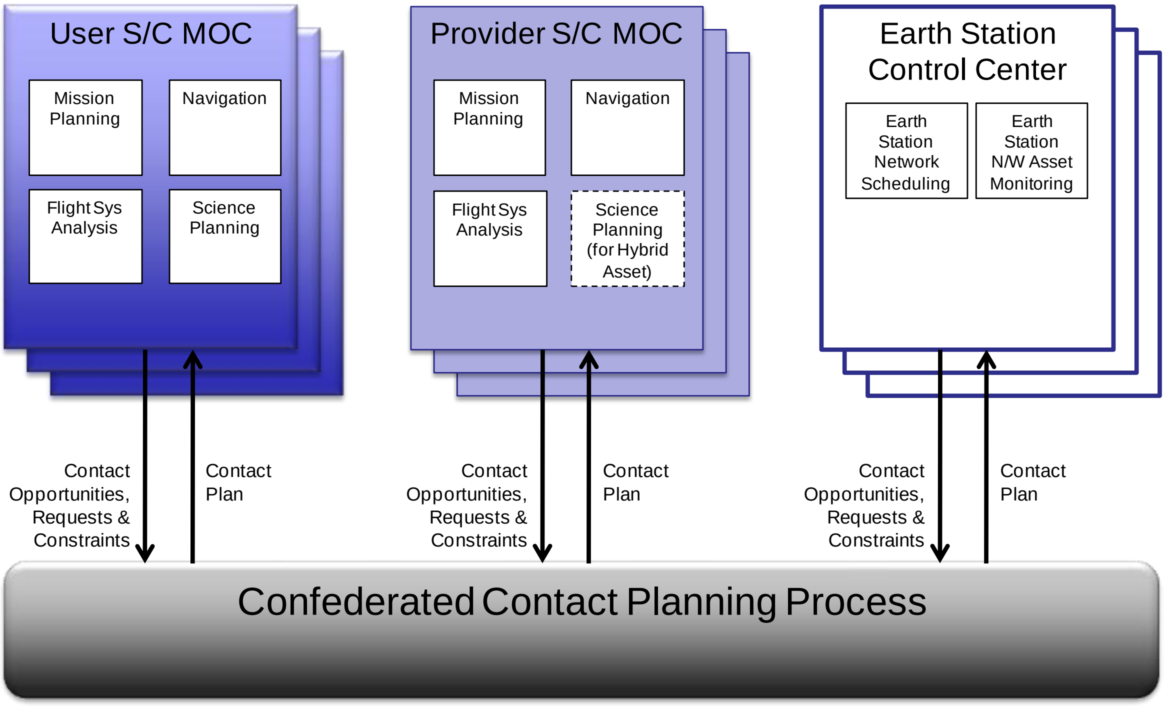

Figure 7-1: SSI Contact Planning Process

Similarly, Provider MOCs will supply the contact requests and constraints of both the Relay Spacecraft and Provider Spacecraft MOC. These inputs will also involve mission planning, navigation, and flight system analysis functions for the Relay Spacecraft, and for the case of hybrid science/telecommunications orbiters, must also include science planning inputs.

Earth Station Control Centers will also be key participants in this confederated contact planning process, capturing the contact requests and constraints of individual Earth Stations for space-to-Earth links. This scheduling must reflect the capabilities of the individual Earth Stations and the maintenance status of subsystem assets within each station. (For scenarios that involve a Planetary Station [i.e., Planet-Space link terminal], inputs from the Planetary Station would also be integrated into the confederated contact planning process.)

The contact planning process must take all these inputs into account to generate the final confederated contact plan that is then disseminated to all nodes. While the detailed process and tools for generating a contact plan are yet to be developed, the SISG's goal is to minimize the need for a highly centralized management activity. Hence, we envision an iterative contact planning process: first, agencies will work internally to establish intra-agency contact plans; and then interagency cross-support links will be added, based on the peering agreements discussed in Section 4.1, to form the confederated contact plan.

It is also likely that multiple, enterprise-specific, confederated contact planning processes will operate in parallel (where an “enterprise” in this context represents a set of inter-related missions, potentially spanning multiple agencies). For instance, one might envision one contact planning process for a set of Mars robotic missions, and a separate process for a set of human lunar exploration missions.

The final contact plan must then be disseminated to all SSI nodes. Individual User and Provider MOCs will be responsible for planning the transmission of contact plans to their Spacecraft nodes.

Once contact plans are in place, SSI network operations can proceed in a highly automated fashion. Figure 7-2 illustrates the run-time operation of the SSI for the scenarios of an end-to-end service between a User MOC and a User Spacecraft. An SSI service is initiated by an application in the User MOC. This application invokes network-layer services within the User MOC, utilizing IP or DTN functionality. Network level addressing will be used to target network data products to the User Spacecraft end destination.

Link-layer services will be established with the next hop, based on the contact plan. This may involve CCSDS Cross Support Transfer Services (CSTS, including Space Link Extension [SLE]), invoking CCSDS TC/TM (or AOS) protocols on the link from the Earth to space. Intermediate nodes, such as Relay Spacecraft, will receive these link-layer products and process them at the network layer, again utilizing IP or DTN functionality. Relay Spacecraft will utilize CCSDS link-layer services and protocols (e.g., Proximity-1 Space Link Protocol) to establish a link to the User Spacecraft. Network services will be terminated at the User Spacecraft, ultimately passing data to an application on the User Spacecraft to complete the end-to-end SSI application-layer services.

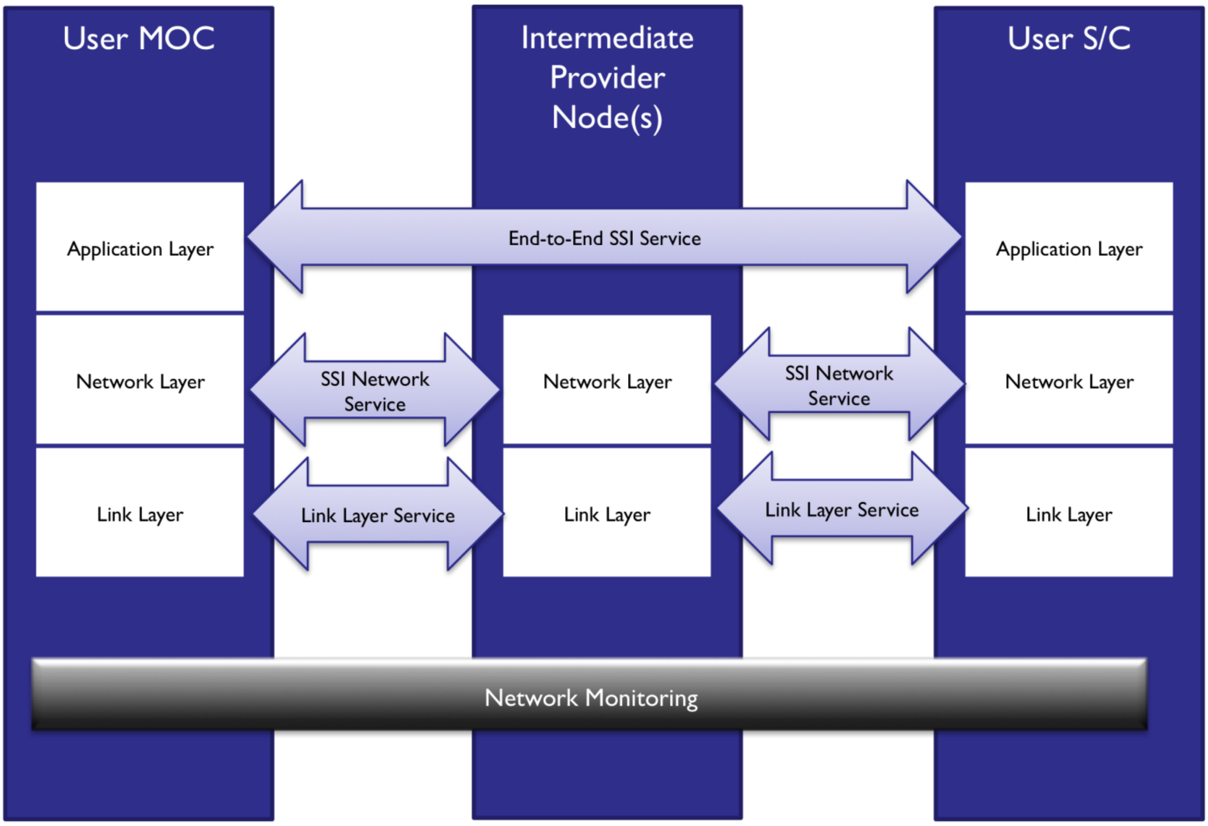

Figure 7-2: SSI Run-time Operations

Throughout this process, network monitoring functions (as described in Section 4.2.4) are operating to make available key performance metrics from each SSI node. These data will be available to SSI users to assess end-to-end service accountability, and to SSI providers to track performance and identify network anomalies.

Appendix A. Acronyms

AMS	Asynchronous Message Services
AOS	Advanced Orbital Systems
ASI	Agenzia Spaziale Italiana (Italian space agency)
BCH	Bose-Chaudhuri-Hocquenghem
BP	Bundle Protocol
BSP	Bundle Security Protocol
CCSDS	Consultative Committee for Space Data Systems
CFDP	CCSDS File Delivery Protocol
CNES	Centre National d'Etudes Spatiales (French space agency)
CSTS	Cross Support Transfer Services
DLR	Deutschen Zentrums für Luft- und Raumfahrt (German Aerospace Center)
DTN	Disruption Tolerant Network
ESA	European Space Agency
ESOC	European Space Operations Centre
ESTEC	European Space Research and Technology Centre
GSFC	Goddard Space Flight Center
GST	Global Science & Technology
HQ	Headquarters
ICD	Interface Control Document
IOAG	Interagency Operations Advisory Group
IOP	Interoperability Plenary
IP	Internet Protocol
ISO	International Standards Organization
ISP	Internet Service Provider
JAXA	Japan Aerospace Exploration Agency
JPL	Jet Propulsion Laboratory
JSC	Johnson Space Center
LEO	Low Earth Orbit
MAP	Multiplexer Access Point
MaROS	Mars Relay Operations Service
MER	Mars Exploration Rover
MEX	Mars Express
MOC	Mission Operations Center
MOU	Memorandum of Understanding
MRO	Mars Reconnaissance Orbiter
MSFC	Marshall Space Flight Center
MSL	Mars Science Laboratory
NASA	National Aeronautics and Space Administration
OSI	Open Systems Interconnection
PDU	Protocol Data Unit
PHX	Phoenix
RF	Radio Frequency
SANA	Space Addressing and Naming Authority
SCID	Spacecraft Identification
SISG	Space Internetworking Strategy Group
SLE	Space Link Extension

SSI	Solar System Internetwork
SSI-ISP	Solar System Internetwork-Internet Service Provider
TC	Telecommand
TM	Telemetry
UDD	User Defined Data
UHF	Ultrahigh frequency
WAN	Wide Area Network

Appendix B. Referenced Documents

---

[i] "Recommendations on a Strategy for Space Internetworking, SISG Phase 1 Report." Interagency Operations Advisory Group's (IOAG) Space Internetworking Strategy Group (SISG). July 28, 2008.

[ii] "Meeting of IOP-2 at the InterContinental Hotel in Geneva, Switzerland – Communiqué." December 10, 2008.

[iii] "Solar System Internetwork (SSI) Issue Investigation and Resolution." Interagency Operations Advisory Group's (IOAG) Space Internetworking Strategy Group (SISG). Date TBD

[iv] "Mars Relay Operations Service (MaROS): Rationale and Approach." Gladden, R. AIAA SpaceOps 2010 Conference, Huntsville, AL. April 2010.