Space Development Agency Network Established Beyond the Upper Limits of the Atmosphere (NEBULA) Standard

Developed by the

Space Development Agency
1670 Air Force Pentagon
Washington, D.C. 20330

Email: OSD.SDA.Outreach@mail.mil

Date:	September 25, 2024
Document ID:	SDA_STD_NEBULA V3.05

DISTRIBUTION STATEMENT A. Approved for public release: distribution unlimited.

Approval Sheet

Approved by:

Touch.Joseph.D.ORC3020012386.ID

Digitally signed by Touch.Joseph.D.ORC3020012386.ID
Date: 2024.09.25 14:01:41 -07'00'

Dr. Joseph D. Touch — Date
SDA NEBULA Network Lead

KANTSIPER.BRIAN.LEE.1398801330

Digitally signed by KANTSIPER.BRIAN.LEE.1398801330
Date: 2024.10.10 16:01:29 -04'00'

Dr. Brian Kantsiper — Date
SDA Chief Engineer

Prepared by:

Dr. Joseph D. Touch
SDA NEBULA Network Lead
E-Mail: OSD.SDA.Outreach@mail.mil

Revision History

Date	Version	Description of Change
Nov 5, 2020	1.0	Draft NEBULA Specification Version 1.0 Added Scope and Assumptions sections (not yet WG reviewed) Added IP Addressing appendix (not yet WG reviewed) Removed the 'planes of operation' from the requirements tables
Feb 5, 2021	2.0	Comments incorporated and removed unnecessary scope or requirements
Oct 28, 2021	3.0	Revision for Tranche 1; not backward compatible with Tranche 0 - Changed network core from IP to MPLS - Added description of network architecture and SV internal reference architecture - Added missing Tranche 0 items: AQM, red/black ECN copy, virtual network support, GEP/ground interfacing - Clarified potential for cross-links/downlink entry at LERs
Jun 9, 2022	3.01	Added signature page Clarified terminology, notably "edge", "host", and "router" Updated ICMP to include ping support (echo request, echo response) Updated DSCP/TOS mapping to support MPLS codepoint-based ECN Removed TBRs and TBDs
Mar 17, 2023	3.02	Added on-path MPLS VPN support (NEBULA 20, NEBULA 21) Limited lack of HAIPE IPv6 support to red end system IPv6 in footnote 4 Removed RSVP-TE (NEBULA_18) Added MPLS to ICMP (NEBULA_32)
Dec 6, 2023	3.03	Revise distribution for public release. Added summary of interoperability between versions in Section 1.4. Added discussion of internetworking in Section 2.4. Added consideration for MPLS overhead (NEBULA_1). Clarified requirements with examples (NEBULA_3, NEBULA_20). Added consideration for all off-SV links (including K_a) and provided exception for explicitly rate-managed links (NEBULA_10). Clarified that MPLS multicast replicas use separate labels (NEBULA_17). Clarified that MPLS IP LERs support IP router features (NEBULA_19). Clarified that permitted IPv6 extension headers include fragmentation, as already required per NEBULA_30 (NEBULA_27). Clarified that "hosts" include IP-addressable routers (NEBULA_30 and NEBULA_33). Clarified references (NEBULA_34).

		Sanitized DSCP/TOS meaning for public release (NEBULA_35). Updated pending references for SDA-TM-0022 and SDA-TM-0023.
Mar 11, 2024	3.03	Corrected typo referring to NEBULA_17 as deprecated rather than NEBULA_18.
Jul 23, 2024	3.04	Added discussion of implications on adjacent networking (Sec. 2.5). Clarified that LER performs ingress and egress IP processing (NEBULA_19). Updated TCP reference from RFC793 to RFC9293 (no substantive change in protocol support intended).
Sep 23, 2024	3.04a	Reference to SDA-TM-0061 renumbered to SDA-TM-0077.
Sep 25, 2024	3.05	Updated NEBULA_4 to include specified extensions. Clarified ECN impact on MPLS (NEBULA_14, NEBULA_24).

Table of Contents

List of Tables

List of Figures

1 Introduction

The National Defense Strategy (NDS) acknowledges that space is vital to the U.S. way of life, our national security, and modern warfare. In an era of renewed great power competition, maintaining our advantage in space is critical to winning these long-term strategic competitions. Potential adversaries seek to undermine this goal by employing strategies that exploit real or perceived vulnerabilities in our current and planned National Security Space systems. In addition, these potential adversaries are developing and demonstrating multi-domain threats to national security. The Department of Defense (DoD) established the Space Development Agency (SDA) on 12 March 2019 as a response to this problem.

SDA is responsible for defining and monitoring the Department's future threat-driven space architecture and accelerating the development and fielding of new military space capabilities necessary to ensure our technological and military advantage in space for national defense. To achieve this mission, SDA will unify and integrate next-generation space capabilities to deliver the Proliferated Warfighter Space Architecture (PWSA), a resilient military sensing and data transport capability via a proliferated space architecture primarily in Low Earth Orbit (LEO). SDA will not necessarily develop and field all capabilities of the PWSA but rather orchestrate those efforts across DoD and fill in gaps in capabilities while providing the integrated architecture.

SDA's mission begins and ends with the warfighter. SDA recognizes that sufficient or "good enough" capabilities in the hands of a warfighter sooner may be better than delivering the perfect solution too late. SDA will deliver capabilities to our joint warfighting forces in two-year tranches, starting with Tranche 0 (T0).

Critical to the success of the proliferated space architecture is the ability to transmit large amounts of data with low latency throughout the constellation and associated ground systems. Each satellite in the PWSA is equipped with Optical Communications Terminals (OCTs) and will be part of the same network. This document defines the standards and requirements for the PWSA network, known as the Network Established Beyond the Upper Limits of the Atmosphere (NEBULA).

The NEBULA Standard establishes the networking requirements for the PWSA network, to include space-to-space and space-to-ground communications. The definition of requirements is intended to be expandable over time in a modular fashion. This version of the NEBULA Standard is not backward compatible with the NEBULA Standard v2.0. Interoperability between v3.x versions is detailed in Section 1.4.

1.1 Purpose and Scope

This document specifies the protocols and behaviors for the space network service provided by the current PWSA.

This document does not specify standards or protocols for the cryptographic capabilities that separate the red and black sides of space vehicle (SV). Such specifications may be included in future versions of the document if deemed necessary by SDA.

External users connect directly to NEBULA through Internet protocols interconnected at ground entry points (GEPs) or indirectly through gateways that act as NEBULA endpoints, e.g., through JREAP-C to Link-16. All future external connections, e.g., to other ground networks or to tactical users, are expected to use one of these two approaches.

1.2 Assumptions

1. All SDA developed SVs, starting with Tranche 1, will need to interoperate and participate in the NEBULA. Any other entity wishing to utilize the SDA Transport Layer to participate in the NEBULA must be compatible with this specification.
2. SVs may have both classified (red-side) and unclassified (black-side) domains.
3. Red-to-red communication between SVs is supported via controlled channels over a black network.
4. The security mechanisms that separate unclassified from classified data will have the appropriate security approvals to operate in the NEBULA.

1.3 Nomenclature and Definitions

1.3.1 Normative Text

The following conventions apply for the normative specifications in this Specification:

a. the words ‘SHALL’ indicates a binding and verifiable specification
b. the word ‘SHOULD’ indicates an optional, but desirable, specification
c. the word ‘MAY’ indicates an optional specification
d. the words ‘IS’, ‘ARE’, and ‘WILL’ indicate statements of fact

NOTE – These conventions do not imply constraints on diction in text except where capitalized.

1.3.2 Data and Symbol Rates

Data and symbol rates are expressed as bits-per-second (bps) and baud. Bps is defined as 1 bit/second. Similarly, baud is defined as 1 symbol/sec. SI-prefixes for these rates are expressed in base-10 and not in base-2. For example, 100 Mbps represents 100 x 10^6 bps or 10^8 bps.

1.3.3 Red-Side Network

The red-side network refers to all systems that operate on unencrypted classified data.

1.3.4 Black-Side Network

The black-side network refers to all systems that operate on unclassified data.

1.4 Interoperability

NEBULA v3.x versions are intended to be backward compatible with earlier 3.x versions. Later versions are designed to add capability that can be deployed and leveraged incrementally at individual nodes, where the new capability travels transparently through legacy nodes. The following provides some details:

- NEBULA v3.0 and v3.01 required support for NEBULA_18 (RSVP-TE), but this has been contractually relieved and is no longer required in NEBULA v3.02 and later versions.
- NEBULA v3.02 adds required support for IPv6 transit even if not supported by the HAIPE, so IPv6 traffic that ingresses into or egresses out of NEBULA may not be supported at nodes using earlier versions.
- NEBULA v3.02 adds required support for Label Switch Routers (LSR) Multiprototcol Label Switching (MPLS) VPNs, so VPNs that start or end at LSRs may not be supported at nodes using earlier versions.
- NEBULA v3.05 adds required support for specific HAIPE extensions.

Note that it is hoped that new capabilities can be added to systems initially deployed using earlier versions, e.g., via firmware or software update.

2 Overview

2.1 NEBULA Network Architecture

The NEBULA network architecture provides communication services throughout on-orbit Transport layer mesh, Tracking layer satellites, partner payloads, radio frequency (RF) air/ground relays, and ground entry points. It relies on established and widely deployed network protocol standards to ensure interoperation while mitigating integration risk. As shown in Figure 1, the NEBULA network architecture supports Internet protocol (IP) communication at the edges over a MPLS backbone.

The NEBULA uses IP for ubiquitous messaging between endpoints, whether space-space or space-ground, and includes a variety of additional protocol layers that compensate for packet loss, duplication, reordering and (where desired) congestion avoidance (e.g., TCP/CUBIC, UDP/QUIC). This includes red and black IP endpoints and ground black network and services.

The end-to-end IP service operates over an Ethernet physical interconnect on-board, to enable plug-compatible interoperability and support the potential for device interchange.

The IP service operates over MPLS when transiting between SVs or to ground entry points, to support backbone network path management and traffic engineering, coordinated using existing network management protocols (e.g., Netconf/YANG). The MPLS paths in the backbone are computed (on ground or in space) and distributed in advance, where each node automatically invokes the appropriate tables based on local time and ephemeris to adjust the backbone to account for orbital dynamics; the set of such tables is referred to as a ‘forwarding almanac’. The MPLS edge for ground links could occur in space (near the space GEP RF), on the ground receiver (near the GEP), or further into the ground IP network (e.g., if it supports MPLS as well), the latter, in particular, to support ground Battle Management Command, Control, and Communications (BMC^3) nodes.

The architecture relies on Active Queue Management (AQM) and Explicit Congestion Notification (ECN) support to reduce latency, as well as the Diffserv (differentiated services) architecture of edge policing and shaping with configurable core weighted queuing to support traffic prioritization.

Figure 2 shows the SV internal NEBULA reference architecture[1], including the common components (within the grey polygon) and components that vary by SV. This includes sensors and other edge devices (e.g., JREAP-C gateways) interconnected with BMC^3 processing using an on-board Ethernet switch. These devices can transfer information off-SV after HAIPE encryption using the MPLS backbone. IP packets enter the backbone at an MPLS Label Edge Router (LER), which adds initial MPLS tags and performs traffic marking, policing, and shaping according to a service profile. Traffic is switched between the backbone nodes over MPLS LSRs, which may both remove and add MPLS tags to control the path of a packet through the network. NEBULA links are directly connected to these LSRs for space-space, space-ground, and space-air interconnect, including RF and optical communication terminals (OCTs). Non-NEBULA links that enter the network via OCTs or K_a-band communications enter at LERs[2]. Additionally, links to devices that terminate at translation gateways acting as black hosts on the NEBULA network also enter the NEBULA at LERs; these devices include black router control interfaces and HAIPEs, the latter of which tunnel traffic for BMC^3 processing, sensors, and JREAP-C gateways that relay information to Link-16. Nearly all traffic between on-board and off-board systems traverses a HAIPE or other network encryptor (e.g., tunnel-mode IPsec).

[1] This internal architecture describes only NEBULA communication within the SV; other internal communication, e.g., for direct bus or payload control, may also be supported.

[2] OCTs and K_a links connecting directly to external (non-NEBULA) devices interface at the LER; those links connect to the LSR when connecting directly to other NEBULA devices.

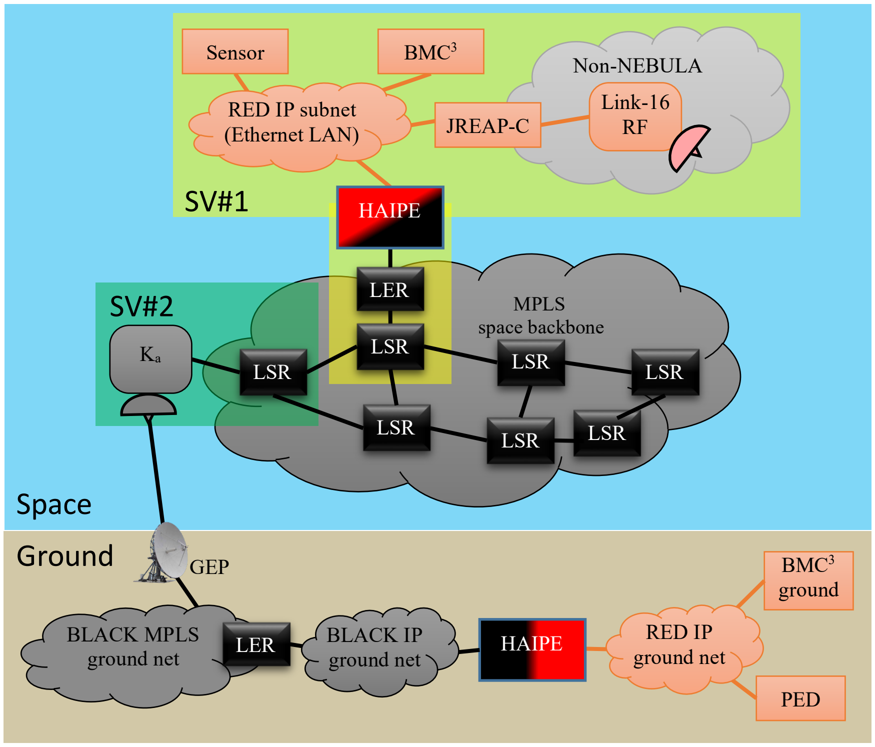

Figure 1: NEBULA network architecture (excepting non-NEBULA Link-16)

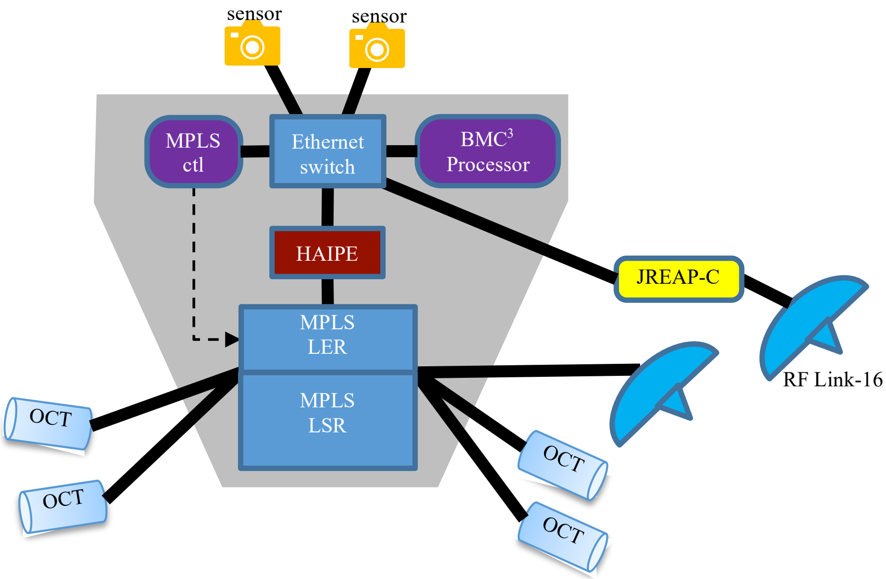

Figure 2: SV internal NEBULA reference

2.2 Security Domains

NEBULA provides data forwarding services to the PWSA system and partners such as ground and airborne systems. If other systems have classified (red-side) data to be exchanged via NEBULA, that data will be encrypted before being carried over a controlled channel over the black NEBULA network and will be decrypted at the appropriate security boundary. The black-side portion of NEBULA will be responsible for end-to-end (or crypto-boundary-to-crypto-boundary) delivery of data.

The red and black addressing domains and their relationship is shown conceptually in Figure 3 and is described in the sections that follow. All red-side assets are reachable (from other red-side assets) by controlled channels that are carried over the black network.

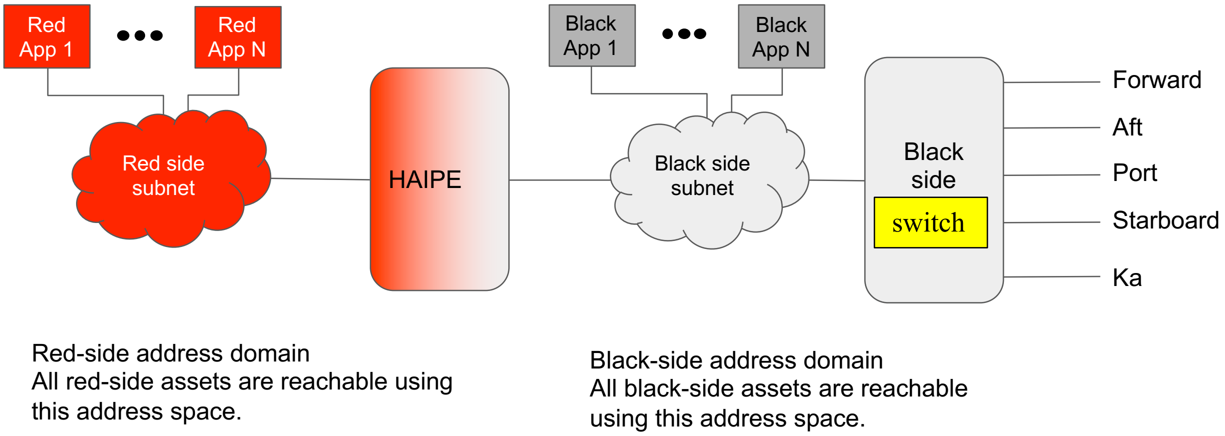
Red App 1
Red App N
Red side subnet
HAIPE
Black App 1
Black App N
Black side subnet
Black side
switch
Forward
Aft
Port
Starboard
Ka
Red-side address domain
All red-side assets are reachable using this address space.
Black-side address domain
All black-side assets are reachable using this address space.

Figure 3: Notional Red and Black Address Domains

2.3 Routing Domains

The red-side network maintains its own address domain and uses unique addresses for each red-side addressable entity (e.g., a host running a BMC[3] application instance). The exact address space to be used for the red-side will be determined later and is not part of this specification.

The black-side network similarly maintains its own address domain with unique addresses for all black-side addressable entities (e.g., the crypto on-board SV that separates the red and black sides). The exact address space to be used for the black-side will be determined later and is not part of this specification.

2.4 Internetworking

NEBULA network devices are intended to support both intranetworking and internetworking. For both SVs and GEPs, off-device links are reconfigurable to connect either to the MPLS LER (for external IP ingress/egress) or the MPLS LSR (for internal MPLS relay), as shown in Figure 2 (implied by how links connect at the boundary between LER and LSR). Reconfigurability is scheduled in the NEBULA Almanac [SDA-TM-0022].

Figure 4 shows the difference between internal (intranet) and external (internet) networking in NEBULA. The left side shows how SV#1 and SV#2 connect their off-device links to the LSR. On the right side, SV#2 and non-NEBULA SVs and GEPs connect to NEBULA via SV#3's LER, either over an OCT or K$_a$ links.

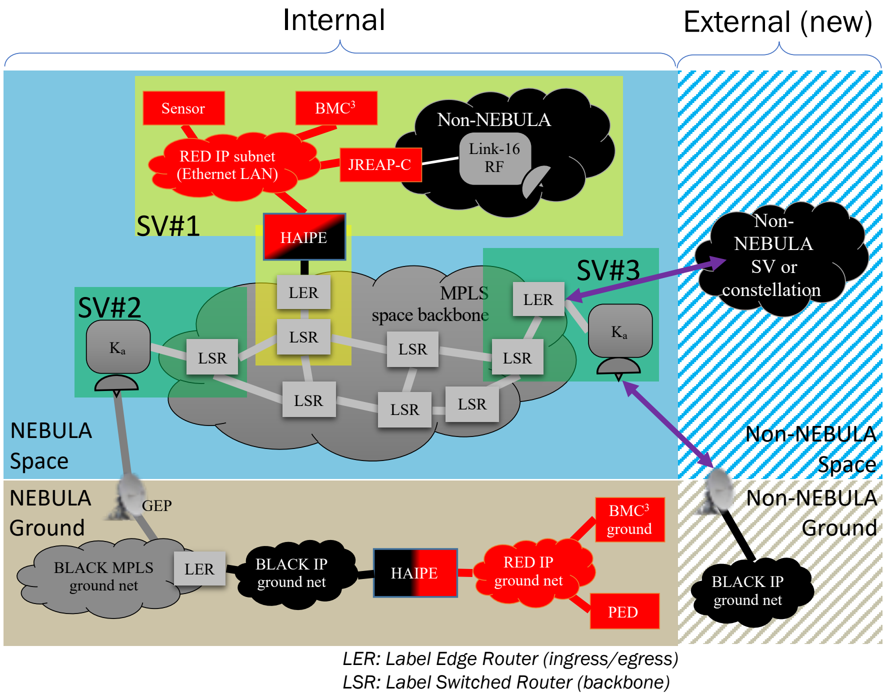

Figure 4: Internal and external networking with NEBULA

All internal and external links, including K_a and OCT, implement transparent Ethernet, i.e., they pass Ethernet traffic. Where external link capacity varies without explicit management, internal links between the black router and the onboard communication terminal also support PAUSE frames for flow control.

2.5 External impacts

The NEBULA standard focuses on the requirements of the space vehicles and its supporting SDA-managed ground network comprising the PWSA. To achieve low-latency end-to-end throughput, some of these requirements are also expected from peer and transit networks that interconnect with PWSA. In specific, these include but are not limited to the following requirements:

- Support for IP packet transit, consistent with the PWSA edge (NEBULA_1, NEBULA_27).
- Support for Diffserv queuing consistent with PWSA prioritization (NEBULA_13 and NEBULA_35).
- Support for low-latency buffering through the use of both ECN (NEBULA_14) and AQM (NEBULA_15), whether in MPLS or IP routers, to ensure that efficient multiplexing and burst accommodation does not result in accumulated latency.
- Where communication with red endpoints is needed, support for HAIPE traversal and endpoint identity (TLS, DTLS) consistent with PWSA (NEBULA_4, NEBULA_5, NEBULA_6, NEBULA_7, NEBULA_8).

3 NEBULA Network Requirements

3.1 General Requirements

Table 1: General Requirements

#	Requirement	Note
NEBULA_1	All network devices SHALL support unfragmented transfer of black IP packets up to a total length of 1500 bytes. The MTU of red IP packets must be reduced accordingly to accommodate HAIPE and MPLS packet overhead.	
NEBULA_2	All network devices SHALL be reconfigurable during regular operation to support changes in their addresses, forwarding tables, and queue parameters, without impact to packet processing.	
NEBULA_3	All network devices SHALL support grouped reconfigurations that are enacted as a set, i.e., as an atomic transaction.	

3.2 Information Assurance

Table 2: Interface Definitions – Information Assurance

#	Requirement	Note
NEBULA_4	<Space Vehicle Crypto> SHALL be compliant with all threshold requirements of HAIPE v4.2.5 and of the extensions listed below: Automated Security Association - IKEv2 Extension Traffic Protection-IKEv2 Suite B Cryptography Classified Traffic Protection-IKEv2 Suite B Cryptography (high side) Classified Traffic Protection-Suite B Cryptography (high side) Early Crossbow (high side) Remote Provisioning	
NEBULA_5	<Network devices> SHALL use TLS v1.3 endpoint authentication for TCP exchanges per [RFC8446].	
NEBULA_6	<Network devices> SHALL use DTLS v1.2 endpoint authentication for UDP exchanges per [RFC6347].	
NEBULA_7	<Space Vehicle Crypto> SHALL copy IP Differentiated Service Code Point (DSCP) markings from the red to black sides.	
NEBULA_8	<Space Vehicle Crypto> SHALL copy IP ECN markings both from the red to black sides and from the black to the red sides.	

3.3 Intra-node Ethernet Switches

Table 3: Interface Definitions – Intra-node Switches

#	Requirement	Note
NEBULA_9	<Space Vehicle Hosts> SHALL use Ethernet as per [IEEE802.1q] as the layer-2 (data link layer) protocol when sending between and among NEBULA links and endpoints within the space vehicle. The Ethernet uses no tags or MAC-in-MAC encapsulation.	Physical requirements appear in the Statement of Work
NEBULA_10	<Ethernet network devices> SHALL support Ethernet frame-based flow control (PAUSE frames), where such frames do not traverse off-SV links (e.g., OCT, K_a). Links that are explicitly rate-managed are excepted.	
NEBULA_11	<Space Vehicles> SHALL support native IP multicast over Ethernet.	

3.4 Inter-node MPLS Switches

Table 4: Interface Definitions – Inter-node Switches

#	Requirement	Note
NEBULA_12	<Space Vehicles and GEPs> SHALL use MPLS for inter-node communication as per [RFC3031].	
NEBULA_13	<Space Vehicles and GEPs> SHALL support MPLS differentiated services (Diffserv) per [RFC3270].	
NEBULA_14	<Space Vehicles and GEPs> SHALL support MPLS ECN marking at switches per [RFC5129]. Note that this includes copying ECN marking between labels during operations (PUSH, POP, SWAP), where TOS permits.	
NEBULA_15	<Space Vehicles and GEPs> SHALL support MPLS AQM at switches per [RFC7567].	
NEBULA_16	<Space Vehicles and GEPs> SHALL be capable of coordinating forwarding MPLS table configurations and link activations according to the provided schedule, referred to as a forwarding almanac.	
NEBULA_17	<Space Vehicles and GEPs> SHALL support native MPLS multicast, which uses separate labels per replica.	
NEBULA_18	(Requirement omitted; requirement label retained for numbering consistency across document versions).	[3]

[3] (requirement omitted; footnote retained for numbering consistency across document versions).

#	Requirement	Note
NEBULA_19	< Space Vehicles and GEPs> that support IP edge traffic SHALL include a MPLS Label Edge Router (LER), which partly supports IP router requirements [RFC1812] for both ingress and egress (post LSR penultimate hop pop) processing.	
NEBULA_20	<Space Vehicles and GEPs> LERs and LSRs SHALL support MPLS virtualization as per [RFC3032], in which stacked labels represent virtual paths.	
NEBULA_21	<Space Vehicles and GEPs> LERs and LSRs SHALL support at least 400 concurrent virtual networks.	
NEBULA_22	< Space Vehicles and GEPs> LERs SHALL support IP DSCP-based marking, policing, and shaping, including mapping IP DSCPs to MPLS on ingress.	
NEBULA_23	< Space Vehicles and GEPs> LERs SHALL support IP DSCP ingress functions for at least 20 concurrent flows.	
NEBULA_24	< Space Vehicles and GEPs> LERs SHALL support copying ECN signals to MPLS on ingress and from MPLS on egress. Note that this includes copying ECN marking involving multiple labels (PUSH2, POP2), where applicable.	
NEBULA_25	<Space Vehicles and GEPs> that support multiple backbone links SHALL include a MPLS Label Switch Router (LSR) to support MPLS traffic between those links.	
NEBULA_26	<Space Vehicles and GEPs> LSRs SHALL support MPLS QoS-based (TC) traffic prioritization.	

3.5 Hosts (red or black IP endpoints)

Table 5: Interface Definitions – Hosts

#	Requirement	Note
NEBULA_27	IPv4 [RFC791] and IPv6 [RFC8200] SHALL be used as the network-layer protocol data unit for the red and black side IP edge messaging to and from hosts, as specified in the following, with no need to support IPv4 options or IPv6 extension headers other than fragmentation and IPsec: [RFC791] Internet Protocol, [RFC1122] Requirements for Internet Hosts, [RFC2474] Definition of the Differentiated Serviced Field (DS Field) in the IPv4 and IPv6 Headers, [RFC3168] The Addition of Explicit Congestion Notification (ECN) to IP, [RFC3260] New Terminology and Clarifications for Diffserv And, for IPv6: [RFC8200] Internet Protocol, Version 6 (IPv6) Specification [RFC8504] IPv6 Node Requirements	4
NEBULA_28	<IP addressable hosts and routing devices> SHALL ensure the IP address space is unique and routable throughout the constellation and terminals accessing it. In particular, network address translation (NAT) functionality is NOT provided by the network itself.	5
NEBULA_29	<Hosts> SHALL support reconfiguration of the IP addresses both before and during operation.	6
NEBULA_30	<IP addressable hosts and routing devices> SHALL implement IP source fragmentation and IP destination reassembly per the above specifications.	
NEBULA_31	<Hosts> SHALL inhibit on-path IPv4 fragmentation on all traffic (set DF=1).	

[4] IPv6 red end system support is contingent on HAIPE IPv6 support.

[5] NAT may be provided by the terminals that access the constellation in order to support their individual missions.

[6] The intention here is that we are NOT specifying the red- or black-side address spaces in this document; only the use of IP. The program will determine the address space of the black-side transport network later in consultation with the ground segment and the space segment will need to accommodate that determination.

#	Requirement	Note
NEBULA_32	<Hosts and MPLS devices> SHALL support the Internet Control Messaging Protocol (ICMP) as specified in the RFCs below, limited to the following messages: Time-exceeded, Echo request, Echo reply, and others as indicated in "NEBULA Standard: ICMP Messages" [SDA-TM-0023]: For IPv4: [RFC792] Internet Control Message Protocol {for IPv4} [RFC4950] ICMP Extensions for Multiprotocol Label Switching And, for IPv6: [RFC4443] Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification	4
NEBULA_33	<IP addressable hosts and routing devices> SHALL implement the IPv4 Address Resolution Protocol (ARP) over Ethernet as specified in [RFC826].	7
NEBULA_34	<Hosts> SHALL use the IP differentiated services code point (DSCP) markings as described in [RFC2474] and [RFC3246] and its updates in the IP header to identify packet priority.	
NEBULA_35	<Hosts and MPLS devices> SHALL use strict priority queueing to treat / service packets according to their QoS markings. The mapping between IP DSCP, MPLS TOS, ECN support, and priorities SHALL be as indicated in [SDA-TM-0077]. Some TOS values do not support ECN and are intended for low-volume traffic. Other TOS values occur as pairs differing by a single bit; that bit is intended to change as needed to indicate ECN and are used for high-volume, congestion-reactive traffic.	8
NEBULA_36	<Hosts> SHALL support at least the Transmission Control Protocol (TCP) as specified in [RFC9293] and its updates and User Datagram Protocol (UDP) as specified in [RFC768] and its updates.	
NEBULA_37	<Hosts> SHALL provide a POSIX socket interface to applications.	
NEBULA_38	<Hosts> SHALL use IETF layer-4 (transport layer) or application mechanisms to remove duplicate packets.	9

[7] ARP traffic, which lacks an IPv4 header, is to be treated as implicitly marked with the highest DSCP priority.

[8] Other choices that could have been made here include e.g., 'standard' DSCP treatments including assured forwarding, expedited service, etc., or various forms of weighted round-robin or weighted fair queueing. This specification opts for requiring strict priority queueing..

[9] TCP, for example, provides a reliable, in-order, without duplication or omissions service to applications. UDP is a best-effort service where datagrams may arrive out of order, duplicated, or not arrive at all (due to loss).

3.6 Network Management

Table 6: Interface Definitions – Network Management

#	Requirement	Note
NEBULA_39	<Network devices> SHALL support telemetering of at least the following per-interface statistics: Number of bytes / packets sent. Number of bytes / packets received. Number of bytes / packets dropped. The status of each of the links (up/down) The current route schedules.	
NEBULA_40	<Network devices> SHALL support network management using Netconf protocols as per [RFC6241].	
NEBULA_41	<Network devices> SHALL support network management using YANG models as per [RFC6020], with specific models as indicated in "NEBULA Standard: YANG Model Definitions" [SDA-TM-0022].	10
NEBULA_42	<Space Vehicle network devices> SHALL support remote configuration from both on-board (e.g., BMC[3], local net controller) and off-board (e.g., remote space BMC[3], ground), used mutually exclusively.	
NEBULA_43	<Space Vehicle network devices> SHALL have an Operations, Administration and Maintenance (OAM) capability for continuous monitoring of the network characteristics to ensure meeting the committed service level agreement (SLA) for the network.	

[10] All SDA-TM-0022 YANG parameters marked as optional must be supported except as noted therein.

Acronyms and Terms

Acronym	Definition
AF	Assured Forwarding, a type of QoS supporting urgent information
Almanac	Set of forwarding tables, each with an activation time and location.
AQM	Active Queue Management
ARP	Address Resolution Protocol
baud	Rate of transmission, defined as 1 symbol/second
BCP	Internet Best Current Practice document
BE	Best Effort, a type of QoS supporting traffic that is not otherwise prioritized.
BMC^3	Battle Management Command, Control, and Communications
bps	Bits per second
CE	Congestion Experienced, a type of ECN marking
CS	Class Selector, a type of QoS supporting non-urgent information
CUBIC	A cubic-order TCP congestion control algorithm, not an acronym.
DF	Don't Fragment, a field of the IPv4 packet header
Diffserv	Differentiated Services
DoD	Department of Defense
DOI	Digital Object Identifier (becomes a persistent URL by prefixing https://)
DS	Differentiated Services
DSCP	Differentiated Service Code Point
DTLS	Datagram Transport Layer Security (TLS)
ECN	Explicit Congestion Notification
EF	Expedited Forwarding, a type of QoS supporting critical information, typically reserved for network control and operations
ESP	(IP) Encapsulating Security Payload
EXP	Originally the Experimental field in MPLS now used for QoS as "TOS"
GEP	Ground Entry Point
HAIPE	High-Assurance IP Encryptor
Host	An IP endpoint that creates or consumes IP packets; this includes tunnel endpoints that do so (e.g., the HAIPE as viewed from its black side), as well as BMC3, sensors, JREAP-C, and network management interfaces supporting Netconf
ICMP	Internet Control Messaging Protocol
IETF	Internet Engineering Task Force
IP	Internet Protocol

Acronym	Definition
IPsec	IP security protocol (note capitalization, per RFC 4301
JREAP-C	Joint Range Extension Applications Protocol, version C
K_a	A particular RF communication band of 26.7-40 GHz
Layer	A layer of the OSI Protocol Reference Architecture
LEO	Low Earth Orbit
LER	(MPLS) Label Edge Router
Link 16	The name of terrestrial RF communications system
LSP	Label switched path
LSR	(MPLS) Label Switch Router
MAC	(Ethernet) Media Access Control (address)
Mbps	Megabits per second
MIB	Management Information Base (network management database)
MPLS	Multiprotocol Label Switching
MTU	Maximum transmission unit, the largest network packet size
NAT	Network Address Translation
NDS	National Defense Strategy
NEBULA	Network Established Beyond the Upper Limits of the Atmosphere
Netconf	The name of a network configuration protocol
OAM	Operations and Management
OCT	Optical Communication Terminal
OSI	Open Systems Interconnect, a network reference architecture
PED	Processing, Exploitation, and Dissemination
POSIX	Portable Operating System Interface
PWSA	Proliferated Warfighter Space Architecture
QoS	Quality of Service (note capitalization)
QUIC	A protocol name now; originally “Quick UDP Internet Connections”
RF	Radio Frequency
RFC	Request for Comments, the specifications of the Internet
Router	An IP device that relays IP packets, modifying some fields (TTL, IP checksum); this includes the HAIPE as viewed from its red side
RSVP-TE	Resource reservation protocol – traffic engineering (extensions)
SDA	Space Development Agency
SLA	Service Level Agreement

Acronym	Definition
SV	Space Vehicle
T0	Tranche 0
T1	Tranche 1
TBD	To be decided (an unknown value that will be determined later)
TBR	To be resolved (a candidate value that will be confirmed or revised later)
TC	Traffic control
TCP	Transmission Control Protocol
TE	Traffic Engineering
TLS	Transport Layer Security
TOS	Type of Service, the field indicating MPLS QoS, formerly “EXP”
UDP	User Datagram Protocol
YANG	Yet Another Next Generation (data model)

References

The following publications contain provisions which, through reference in this text, constitute provisions of this document. At the time of publication, the editions indicated were valid. All publications are subject to revision, and users of this document are encouraged to investigate the possibility of applying the most recent editions of the publications indicated below.

1. [IEEE802.1q] IEEE Standard for Local and Metropolitan Area Networks—Media Access Control (MAC) Bridges and Virtual Bridged Local Area Networks 802.1Q-2014.
2. [RFC768] Postel, J., "User Datagram Protocol," RFC768, August 1980, <https://www.rfc-editor.org/info/rfc768>.
3. [RFC791] Postel, J., "Internet Protocol," STD 5, RFC 791, DOI 10.17487/RFC0791, September 1981, <https://www.rfc-editor.org/info/rfc791>.
4. [RFC792] Postel, J., "Internet Control Message Protocol," STD 5, RFC 792, DOI 10.17487/RFC0792, September 1981, <https://www.rfc-editor.org/info/rfc792>.
5. [RFC826] Plummer, D., "Ethernet Address Resolution Protocol: Or converting network protocol addresses to 48.bit Ethernet address for transmission on Ethernet hardware." STD 37, RFC 826, November 1982, <https://www.rfc-editor.org/info/rfc826>.
6. [RFC1122] Braden, R., Ed., "Requirements for Internet Hosts - Communication Layers," STD 3, RFC 1122, DOI 10.17487/RFC1122, October 1989, <https://www.rfc-editor.org/info/rfc1122>.
7. [RFC1812] Baker, F., Ed., "Requirements for IP Version 4 Routers," RFC 1812, DOI 10.17487/RFC1812, June 1995, <https://www.rfc-editor.org/info/rfc1812>.
8. [RFC2474] Nichols, K., Blake, S., Baker, F., and D. Black, "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers," RFC 2474, DOI 10.17487/RFC2474, December 1998, <https://www.rfc-editor.org/info/rfc2474>.
9. [RFC2597] Heinanen, J., Baker, F., Weiss, W., and J. Wroclawski, "Assured Forwarding PHB Group," RFC 2597, DOI 10.17487/RFC2597, June 1999, <https://www.rfc-editor.org/info/rfc2597>.
10. [RFC3031] Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol Label Switching Architecture," RFC 3031, DOI 10.17487/RFC3031, January 2001, <https://www.rfc-editor.org/info/rfc3031>.
11. [RFC3032] Rosen, E., Tappan, D., Fedorkow, G., Rekhter, Y., Farinacci, D., Li, T., and A. Conta, "MPLS Label Stack Encoding," RFC 3032, DOI 10.17487/RFC3032, January 2001, <https://www.rfc-editor.org/info/rfc3032>.

12. [RFC3168] Ramakrishnan, K., Floyd, S., and D. Black, “The Addition of Explicit Congestion Notification (ECN) to IP,” RFC 3168, DOI 10.17487/RFC3168, September 2001, <https://www.rfc-editor.org/info/rfc3168>.
13. [RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, “RSVP-TE: Extensions to RSVP for LSP Tunnels,” RFC 3209, DOI 10.17487/RFC3209, December 2001, <https://www.rfc-editor.org/info/rfc3209>.
14. [RFC3246] [RFC3246] Davie, B., Charny, A., Bennet, J., Benson, K., Le Boudec, J., Courtney, W., Davari, S., Firoiu, V., and D. Stiliadis, “An Expedited Forwarding PHB (Per-Hop Behavior),” RFC 3246, DOI 10.17487/RFC3246, March 2002, <https://www.rfc-editor.org/info/rfc3246>.
15. [RFC3260] Grossman, D., “New Terminology and Clarifications for Diffserv,” RFC 3260, DOI 10.17487/RFC3260, April 2002, <https://www.rfc-editor.org/info/rfc3260>.
16. [RFC3270] Le Faucheur, F., Wu, L., Davie, B., Davari, S., Vaananen, P., Krishnan, R., Cheval, P., and J. Heinanen, “Multi-Protocol Label Switching (MPLS) Support of Differentiated Services,” RFC 3270, DOI 10.17487/RFC3270, May 2002, <https://www.rfc-editor.org/info/rfc3270>.
17. [RFC3812] Srinivasan, C., Viswanathan, A., and T. Nadeau, “Multiprotocol Label Switching (MPLS) Traffic Engineering (TE) Management Information Base (MIB),” RFC 3812, DOI 10.17487/RFC3812, June 2004, <https://www.rfc-editor.org/info/rfc3812>.
18. [RFC4301] Kent, S. and K. Seo, “Security Architecture for the Internet Protocol,” RFC 4301, DOI 10.17487/RFC4301, December 2005, <https://www.rfc-editor.org/info/rfc4301>.
19. [RFC4303] Kent, S., “IP Encapsulating Security Payload (ESP),” RFC 4303, DOI 10.17487/RFC4303, December 2005, <https://www.rfc-editor.org/info/rfc4303>.
20. [RFC4443] Conta, A., Deering, S., and M. Gupta, Ed., “Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification,” STD 89, RFC 4443, DOI 10.17487/RFC4443, March 2006, <https://www.rfc-editor.org/info/rfc4443>.
21. [RFC4950] Bonica, R., Gan, D., Tappan, D., and C. Pignataro, “ICMP Extensions for Multiprotocol Label Switching,” RFC 4950, DOI 10.17487/RFC4950, August 2007, <https://www.rfc-editor.org/info/rfc4950>.
22. [RFC5129] Davie, B., Briscoe, B., and J. Tay, “Explicit Congestion Marking in MPLS,” RFC 5129, DOI 10.17487/RFC5129, January 2008, <https://www.rfc-editor.org/info/rfc5129>.
23. [RFC6020] Bjorklund, M., Ed., “YANG - A Data Modeling Language for the Network Configuration Protocol (NETCONF),” RFC 6020, DOI 10.17487/RFC6020, October 2010, <https://www.rfc-editor.org/info/rfc6020>.

24. [RFC6241] Enns, R., Ed., Bjorklund, M., Ed., Schoenwaelder, J., Ed., and A. Bierman, Ed., “Network Configuration Protocol (NETCONF),” RFC 6241, DOI 10.17487/RFC6241, June 2011, <https://www.rfc-editor.org/info/rfc6241>.

25. [RFC6347] Rescorla, E. and N. Modadugu, “Datagram Transport Layer Security Version 1.2,” RFC 6347, DOI 10.17487/RFC6347, January 2012, <https://www.rfc-editor.org/info/rfc6347>.

26. [RFC7567] Baker, F., Ed., and G. Fairhurst, Ed., “IETF Recommendations Regarding Active Queue Management,” BCP 197, RFC 7567, DOI 10.17487/RFC7567, July 2015, <https://www.rfc-editor.org/info/rfc7567>.

27. [RFC8200] Deering, S. and R. Hinden, “Internet Protocol, Version 6 (IPv6) Specification", STD 86, RFC 8200, DOI 10.17487/RFC8200, July 2017, <https://www.rfc-editor.org/info/rfc8200>.

28. [RFC8312] Rhee, I., Xu, L., Ha, S., Zimmermann, A., Eggert, L., and R. Scheffenegger, "CUBIC for Fast Long-Distance Networks,” RFC 8312, DOI 10.17487/RFC8312, February 2018, <https://www.rfc-editor.org/info/rfc8312>.

29. [RFC8446] Rescorla, E., “The Transport Layer Security (TLS) Protocol Version 1.3,” RFC 8446, DOI 10.17487/RFC8446, August 2018, <https://www.rfc-editor.org/info/rfc8446>.

30. [RFC8504] Chown, T., Loughney, J., and T. Winters, “IPv6 Node Requirements,” BCP 220, RFC 8504, DOI 10.17487/RFC8504, January 2019, <https://www.rfc-editor.org/info/rfc8504>.

31. [RFC9293] Eddy, W., Ed., “Transmission Control Protocol (TCP),” STD 7, RFC 9293, August 2022, <https://www.rfc-editor.org/info/rfc9293>.

32. [SDA-TM-0022] “NEBULA Standard: YANG Model Definitions,” SDA-TM-0022 V1.01, September 25, 2024.

33. [SDA-TM-0023] “NEBULA Standard: ICMP Messages,” SDA-TM-0023 V1.01, September 25, 2024.

34. [SDA-TM-0077] “NEBULA Standard: IP DSCP and MPLS TOS Map and Meanings,” SDA-TM-0077 V1.01, September 23, 2024.