



ELEVSOFT

SILICON

Cross-Vendor Networking Layer for Space

ELEVSOFT SILICON WHITEPAPER

DOCUMENT	ES-002
SUBJECT	Elevsoft Space NIC Overview and Frequently Asked Questions
ISSUED	9/20/2026
CLASSIFICATION	Public
AUTHORS	Raghav Sharma, Hirdae Sandhu, Leo Shim

Abstract

The growing number of low Earth orbit satellite constellations is increasing demand for reliable, high-capacity communication. However, these constellations largely operate as separate mesh networks, limiting opportunities to share capacity. Interconnecting compatible networks could give participating operators access to additional bandwidth and alternative routes, improving both capacity utilization and resilience. A standardized system which is designed to supplement the networks of constellations and enable spacecraft to share traffic among different networks where compatible communication links and commercial agreements are in place. Wider interconnection would offer additional delivery routes, give access to spare transport capacity, and increase resilience while at the same time reducing the need for repeated integration work. The long-term goal is to make interoperable satellite networking more accessible by means of a widely adopted onboard component and a commercially viable interconnection service. This greater accessibility could enable satellite operators to expand connectivity at lower integration costs, use existing capacity more efficiently, and generate additional revenue by carrying traffic for other networks.

Index

1. The Problem Description.....	4
1.1 Introduction.....	4
1.2 The Way LEO Satellite Networks Work.....	4
1.3 The Two Key Problems: Routing and Translation.....	5
1.3a The Routing Problem.....	5
1.3b The Translation Problem.....	5
2. Market Trends.....	6
2.1 Commercial Market Trends.....	6
2.2 Defense Market Trends.....	6
3. The Solution.....	6
3.1 The Space-NIC Overview.....	6
3.2 Simulation-Based Evaluation of Satellite Interconnection and Routing.....	7
Orbital data and simulation scope.....	7
Connectivity across constellation boundaries.....	8
Delivery between ground terminals.....	8
Ground computed forwarding schedules.....	9
Recovery from unexpected link failures.....	10
Variation across repeated trials.....	10
Failures affecting groups of links.....	11
Timeliness of ground updates.....	12
Selecting the forwarding interval.....	12
Storage and update requirements at larger scales.....	13
Performance when link capacity is constrained.....	14
Effect of deployment placement.....	14
4. Frequently Asked Questions.....	15
5. Acronyms and Abbreviations.....	16
6. References.....	17

1. The Problem Description

1.1 Introduction

Low Earth orbit (LEO) is filled with satellite networks that cannot communicate with each other. Public tracking catalogs list more than 16,000 satellites belonging to just a few operators, and each fleet operates in its own isolated environment, built on entirely separate systems. In contrast, terrestrial telecommunication providers eliminated this problem many generations ago by adopting common standards and using hardware that implemented them. This enabled customers to roam between different operators, improved connectivity, and made each carrier's network more valuable. Orbit still lacks interoperability; a packet sent on one operator's constellation cannot be transmitted to another operator's, even though the fleets do pass within line of sight of each other. The interconnected version of that sky is more valuable for everyone, but it is severely fragmented as it stands today.

1.2 The Way LEO Satellite Networks Work

For most of the industry's history, communications satellites have been geostationary space vehicles designed to operate at very high altitudes, offering extensive coverage but also high round-trip latency and low service capacity (kbps/km²). [3] However, this approach is no longer suitable for today's data-intensive economy and what lies ahead, because of growing demand for computing power and the need for real-time interactivity in both consumer and military applications. As a result, demand is growing, driving companies to invest heavily and launch satellites into low Earth orbit at altitudes roughly 60 times lower than geostationary satellites. Because the satellites fly much closer to the ground, latency decreases, and the capacity density increases. [3]

The trade-off is no continuous coverage since each LEO satellite can see only a small part of the Earth, and its flying in orbit. In order to maintain continuous coverage, it is necessary to have a group of low Earth orbit satellites with links between them to pass traffic among one another. Each spacecraft establishes simultaneous connections to several neighbors, thus forming a mesh network. Data traffic must pass through this mesh from satellite to satellite until it reaches its final destination, which may be another satellite or a gateway to the ground.




Figure 1: The bent pipe and the LEO mesh. A geostationary relay (on the left) retransmits between the two ground stations and makes no routing decisions. A LEO constellation (on the right) transfers a packet via optical inter-satellite links from a user to a satellite that can connect to one. Not to scale.

1.3 The Two Key Problems: Routing and Translation

Unlike terrestrial networks, where standardized protocols and commercially available hardware let independently operated networks interconnect, the LEO satellite ecosystem remains largely fragmented across independent constellations. [3, 4] Major space players, such as Starlink, have vertically integrated networking and routing capabilities within their own constellations. However, smaller and emerging operators do not have comparable ecosystems or the resources to develop one. As a result, many smaller constellations face essentially two fundamental challenges in participating in a broader, interoperable LEO network: dynamically routing traffic across a continuously changing inter-satellite topology and translating between the heterogeneous data formats used by different spacecraft, terminals, and network architectures. [3, 7, 9–11, 26] These challenges become particularly significant at the boundary between independently operated constellations, where both routing decisions and protocol translation must occur at the speed of inter-satellite links without imposing prohibitive computational, power, or hardware requirements. Addressing these two challenges requires a networking layer capable not only of adapting to dynamic LEO topology, but also of translating between otherwise incompatible systems, enabling satellites and constellations that were not originally designed to operate together to function as an interoperable network.

1.3a The Routing Problem

In terrestrial networks, link-state routing protocols such as Open Shortest Path First (OSPF) share reports of link changes so that each router can update its picture of the network and recalculate the routes. The time it takes for these updates to spread throughout the network and for the routes to stabilize is known as convergence. During this period, outdated routes may cause packet loss or forwarding loops. [5] In a LEO mesh, orbital motion can repeatedly change which links are available, making convergence a recurring operating cost. [3] One published emulation of a 48-satellite constellation reported 235 routing-update and acknowledgment packets for a single topology change, consuming inter-satellite capacity without carrying any user data. [6] In contrast to terrestrial networks, the relative motion and position of the satellites within a constellation are predictable and By using the orbital trajectories with planned link schedules, operators are able to compute routing tables on the ground and load them before they are required. If these tables are activated at scheduled transition times, reactive convergence can be avoided for expected changes, reducing control traffic and interruptions to packet delivery. [26] The same emulation reported no convergence-message overhead for predictable changes when orbit-aware routing was used. [6] This indicates possible savings in routing overhead. The extent of throughput gains will depend on link capacity, traffic load, update frequency, and the cost of distributing the tables. Unexpected failures still need to be detected and recovered. A routing-aware networking layer could provide this forwarding capability to constellations without onboard mesh routing, enabling them to take part in a wider interconnected network. Interoperability between constellations would make additional routes available across vendors, providing alternatives around congestion and failures and access to more ground gateways. To use these routes, traffic must pass between spacecraft and terminals that use different protocols and data formats, which brings the second challenge: protocol translation.

1.3b The Translation Problem

Networking equipment in space adheres to several standards, but no single standard governs the complete path from an optical terminal through the spacecraft's avionics to the ground. The U.S. Space Development Agency's Optical Communications Terminal (OCT) standard defines interoperability between compliant laser terminals and exists in several versions, but it does not make all spacecraft interfaces interchangeable. [7, 8] Along this path, equipment may use vendor-specific formats, CCSDS protocols for space communications, Ethernet, or on-board interfaces such as SpaceWire and SpaceFibre. [7, 9–11] Because these systems operate at different layers and serve different purposes, transferring traffic between them may require changes to framing, encapsulation, and flow control while preserving the original data. [9–11] It is therefore possible for even a single operator using equipment from two different suppliers to encounter an integration problem, a problem that will continue to exist across constellations whenever incompatible interfaces come together. A route is useful only if the equipment on it can exchange the traffic it carries. To achieve interoperability at scale, a common set of standards must be adopted and implemented using vendor-agnostic networking hardware, like Ethernet and IP standards currently enable equipment from different suppliers to connect in terrestrial networks. With this approach, operators could keep their internal architectures while using agreed-upon interfaces to exchange traffic at the network boundaries. [4, 12] It is also necessary for translation to maintain the required data rate within the spacecraft's processing and electrical power limitations, or else it will act as a bottleneck and reduce the benefits of the optical link. The German Aerospace Center's development of optical terminals designed specifically for small satellites shows how important it is to be able to handle high data rates within the limited resources of a spacecraft. [13] Future standards revisions and new equipment will also need to include implementations that can support additional formats after launch, within the constraints of their physical interfaces and processing capacity. The challenge therefore lies in combining common interoperability

standards with reusable, adaptable on-board hardware to reduce the custom integration needed for each terminal-spacecraft combination and make interconnection easier as more operators adopt compatible implementations.

2. Market Trends

2.1 Commercial Market Trends

Various operators are establishing LEO networks using their own separate architectural designs and supply chains. Amazon Leo (formerly Project Kuiper) has stated that its optical mesh is being designed entirely in-house, [14] Telesat has chosen MDA Space to manufacture its Lightspeed constellation, [15] and MDA in turn has selected Tesat to provide the optical terminals. [16] Kepler has also independently demonstrated its ability to establish optical links and to set up an IP mesh network between its Pathfinder satellites. [17] The extent of these investments shows how strongly there is a demand for onboard networking and indicates a potential opportunity for a standardised network layer. This would involve setting up optical links, co-ordinating network operations, and negotiating traffic-sharing agreements as well as supporting and standardising network framing. Greater connectivity and a reusable networking component could benefit satellite operators by expanding their routing options and reducing the need for custom integration.

In the near future, space-based data centers are likely to become yet another possible source of commercial demand. Google's Project Suncatcher is studying the possibility of a network of AI-computing satellites.[18] Axiom has flown early orbital computing nodes and, with Spacebilt, plans an optically linked ISS node for 2027. Axiom has described a connectivity plan that involves raising speeds from 2.5 Gbps to future 100 Gbps links, showing the level of network performance such services would require. [19] When these distributed orbital data centers are deployed at scale, demand will rise for high-bandwidth links between satellites to exchange workloads and data, as well as for ground connections to upload inputs and send back results.

2.2 Defense Market Trends

The growing need for resilient military satellite infrastructure is driving demand for controlled interconnection with commercial networks, which can serve as additional infrastructure by providing additional capacity and alternative communication routes. This is explicitly stated by the U.S. Space Force's own Commercial Space Strategy , [20] and its Space Data Network program is currently funding demonstrations of connectivity between different providers.[21] The GAO lists alternative communication methods as one advantage of the military adopting an integrated satellite communications architecture. [22] A pertinent example is the early internet: the defense forces had helped to develop shared networking technology while at the same time maintaining their own separate operational network through MILNET. [4] Similarly, military satellite networks could retain their own security boundaries while selectively using the wider transport network.

Defense and government agencies in the United States and Europe commit ample funding to secure satellite communications. SDA's approximately \$1.8 billion award for the PWSA Tranche 1 Transport Layer [23] and the EU's €6 billion funding commitment announced for IRIS² [24] illustrate this purchasing capacity, although these program budgets cover much more than onboard networking equipment. SDA's purchase of satellites and optical terminals from various manufacturers as part of the Proliferated Warfighter Space Architecture clearly indicates the networking needs a turnkey solution [23, 25]. To make these spacecraft function as a single network, interoperability is needed at two levels: the Optical Communications Terminal standard covers communication between compatible laser terminals, [7, 8] while NEBULA specifies network functions such as packet forwarding and the scheduled activation of forwarding tables. [26] These standards point to a clear demand for interoperable links and coordinated routing, and when the supported equipment is not compatible, interface adaptation is required.

3. The Solution

3.1 The Space-NIC Overview

Space-NIC is an onboard networking card designed to link a satellite's optical communication terminals with its internal systems. It will accept the data that is coming in, determine where each data packet should be sent, and convert the data formats it receives into those needed by the receiving interface. Although routing decides where a packet should go, translation prepares the packet for the next interface by interpreting the framing of the incoming data and modifying the packet's encapsulation without altering its payload. This should enable data to be exchanged between terminals and onboard systems from different suppliers with less custom integration, on the condition that their physical links are compatible. The functions would run in reconfigurable hardware, reducing the processing burden on the spacecraft's main

computer. [1] In order to support routing as the satellites move, the forwarding tables will be computed on the ground using predicted orbital positions and will be uploaded in advance, enabling the Space-NIC to switch between tables when scheduled connections become available or cease. [1] Space-NIC is designed to support integration and interconnection across different spacecraft and network architectures in three ways:

1. Reusable integration: Space-NIC combines packet forwarding and supported interface translation into a single component, simplifying integration across spacecraft while accommodating existing proprietary mesh networks.
2. For operators without a mesh network: Space-NIC and its supporting software can provide routing functions between satellites equipped with compatible inter-satellite links.
3. For operators with an existing mesh network: Internal traffic can continue using the operator's current routing system, while traffic crossing into another operator's network passes through Space-NIC at designated interconnection points. This preserves the existing network structure while enabling wider connectivity where beneficial.

3.2 Simulation-Based Evaluation of Satellite Interconnection and Routing

Orbital data and simulation scope

Elevsoft assessed its proposed networking approach through simulations based on orbital data for 11,452 satellites in the Starlink, OneWeb and Iridium NEXT constellations. [2] The simulations compared networks operating independently with networks sharing a common routing layer, conditional on compatible inter-satellite links. Figure 3.2.1 illustrates the modeled orbital environment and potential cross-operator connections. Orbital proximity alone does not establish that existing spacecraft can form these links; terminal compatibility and operational access remain separate requirements.




Figure 3.2.1. (Above) Modeled satellite positions and candidate cross-operator links. The highlighted connections are subject to the simulation's link assumptions, rather than verified interoperability between deployed terminals. Source: [2].

Connectivity across constellation boundaries

Within this modeled environment, enabling cross-constellation routing increased the proportion of connected cross-operator satellite pairs from zero in the isolated baseline to approximately 85% over the simulated orbit. [2] The gallery attributes the remaining gap partly to conservative link-assignment assumptions. This establishes a network-level connectivity result: it measures which pairs can reach one another, rather than how much bandwidth they receive or whether every packet arrives.




Figure 3.2.2. Cross-operator connectivity over time for isolated and interconnected simulated networks. The zero baseline follows from disabling cross-operator traffic in the model. Source: [2].

Delivery between ground terminals

To examine the effect on traffic delivery, city-to-city scenarios modeled packets travelling between ground terminals, with satellites relaying data as the spacecraft moved. Across the tested scenarios, the reported overall packet-delivery rate increased from approximately 71% to 99.9%, a gain of 28.9 percentage points. [2] Individual city pairs showed different improvements, as Figure 3.2.3 illustrates. These results show how additional routes could help deliver traffic when an isolated constellation cannot maintain a continuous path.




Figure 3.2.3. Packet delivery by city pair using modeled OneWeb ground-terminal access. The overall 71% and 99.9% figures summarize the tested scenarios; the gain is not identical for every city pair. Source: [2].

Ground computed forwarding schedules

Alongside connectivity, the simulations examined whether predictable orbital motion could simplify routing by allowing forwarding tables to be computed on the ground and uploaded in advance. In the 243-satellite test, the comparison protocol, Optimized Link State Routing (OLSR), generated approximately 1.9 MB of network-control traffic per second across the simulated network. [27] The scheduled approach required approximately 32.6 KB of forwarding-table storage and ground updates of approximately 3.5 kbps per satellite. [2] These figures describe different resource demands, illustrating a trade-off between continuous routing-information exchange and ground computation with table uploads.



Figure 3.2.4. Modeled inter-satellite routing-control traffic over ten minutes. The zero bar excludes ground uploads and does not mean zero total networking overhead. Storage, upload rate and network-wide traffic are distinct metrics. Source: [2].

Recovery from unexpected link failures

Unexpected failures can invalidate a planned route. The simulations therefore tested preloaded backup paths and local repair, allowing a spacecraft to redirect traffic without waiting for a new ground-computed schedule. With 5% of links unavailable, the strongest tested repair scheme increased delivery from approximately 62% to 80%, while conventional dynamic routing remained more resilient. [2] Figure 3.2.5 shows that the benefit of scheduling must be considered together with its response to departures from the predicted network state.




Figure 3.2.5. Packet delivery under increasing proportions of unavailable links. Local repair improves on unrepaired schedules, but OLSR maintains higher delivery in these stress tests. Tested failure rates are scenario inputs. Source: [2].

Variation across repeated trials

Repeated seeded trials explored whether the recovery results depended on a particular failure placement. Under the tested low-failure conditions, scheduled routing with local repair and OLSR both achieved median delivery of approximately 99.9%. [2] The distribution matters alongside the median: some repaired-schedule trials had lower delivery than the dynamic-routing comparison. Figure 3.2.6 therefore supports strong typical performance within these scenarios while retaining the evidence of less favorable outcomes.




Figure 3.2.6. Twenty trials per scheme at approximately 0.2% and 0.5% failed links. Points represent individual runs and horizontal bars mark medians; equal medians do not imply equal reliability in every run. Source: [2].

Failures affecting groups of links

The simulations also considered failures affecting an entire satellite or a regional group of links, rather than only scattered individual links. Figure 3.2.7 compares these failure patterns and shows that delivery depends on both the number and placement of unavailable links. [2] Local repair retained substantial connectivity in the tested clustered cases, but performance varied. These results extend the failure analysis without establishing that all regional outages or correlated failures would have similar effects.



Figure 3.2.7. Delivery for random link losses, failed satellites and a regional cluster in the 243-satellite model. Comparisons apply to the particular failure locations and scenarios tested. Source: [2].

Timeliness of ground updates

Because local repair does not recover every affected route, timely ground updates remain important. In the tested 5% link-failure scenario, longer delays before corrected forwarding tables arrived produced greater packet loss, even with local repair. [2] Figure 3.2.8 illustrates this relationship. Only correction delays through five minutes were measured; the longer-delay curves are linear extrapolations. Scheduled forwarding must therefore combine local recovery with a practical method for detecting changes, recomputing routes and delivering updates.




Figure 3.2.8. Packet loss versus correction delay. Solid measured results extend through five minutes; dotted segments beyond that range are extrapolations, not additional simulated observations. Source: [2].

Selecting the forwarding interval

The interval between scheduled forwarding states also affects performance. In the tested moving-network scenario, intervals of one to five seconds preserved higher delivery than a sixty-second interval. A five-second interval was reported to reduce the update stream by roughly fourfold for an approximately 0.25-percentage-point delivery trade-off relative to the tested finer setting. [2] Figure 3.2.9 makes this design choice explicit: schedule granularity must balance update requirements against the accuracy of the planned network state.




Figure 3.2.9. Delivery versus forwarding epoch length under normal and failure conditions. The preferred interval is specific to the modeled topology, traffic and update implementation. Source: [2].

Storage and update requirements at larger scales

A separate scaling test examined whether forwarding schedules remained practical as the network grew. It combined the real satellite fleet with synthetic expansion, reporting approximately 32.6 KB of schedule storage per satellite at 11,452 satellites and approximately 144 KB at 50,160 satellites, with an update stream of approximately 15 kbps per satellite at the larger scale. [2] These results suggest modest storage requirements for the modeled representation. They do not measure the total memory, processing or ground-computation requirements of a complete flight implementation.




Figure 3.2.10. Per-satellite schedule storage and update rate versus network size. The largest networks include synthetic growth; the two curves use different axes and units. Source: [2].

Performance when link capacity is constrained

A further test explored the interaction between routing-control traffic and user traffic on constrained links. With simulated links limited to 1 Mbps and approximately 680,000 packets offered, the scheduled approach delivered approximately 99.93% of packets, compared with 64.2% for the link-state comparison. [2] The gallery reports that control traffic competed with user data in this configuration. This demonstrates a possible benefit under the tested congestion conditions; it does not establish the same gain on higher-capacity optical links or under other protocol configurations.




Figure 3.2.11. Packet delivery in the constrained heavy-load experiment. The result depends on the 1 Mbps link limit, offered traffic and routing configuration. “Zero routing chatter” refers to the modeled inter-satellite control exchange. Source: [2].

Effect of deployment placement

Finally, deployment location influenced the connectivity benefit. Adding the proposed networking functionality at selected network boundaries on approximately 5% of satellites produced approximately 78% cross-operator connectivity in the model. [2] Random placement required greater coverage to approach the same result. Figure 3.2.12 illustrates why strategically located interconnection points may provide value before universal adoption. This is a topology result, with physical link compatibility and the implementation of supported interface translation still requiring separate validation.




Figure 3.2.12.

Cross-operator connectivity versus the share of satellites with the proposed functionality. Targeted and random placement are compared under the same model assumptions; this is not a hardware deployment or revenue forecast. Source: [2].

Taken together, these simulations offer preliminary evidence that shared connectivity can improve traffic delivery and that ground-computed schedules can reduce ongoing inter-satellite routing-control exchange. The additional tests clarify the importance of recovery mechanisms, timely updates, forwarding intervals and deployment placement. The findings remain conditional on the simulated links, traffic and failure assumptions. Physical link compatibility, hardware performance, security and interface translation must be validated separately before the results can support claims about an operational Space-NIC network.

4. Frequently Asked Questions

- a. **Is Space-NIC a router, a network management system, or an orchestrator?**
Space-NIC is most similar to an on-board router having an integrated interface translation facility; it forwards packets using tables prepared on the ground and adjusts the data formats it supports for the outgoing interface, the network planning, scheduling and monitoring still remaining the responsibility of the ground-based support systems and the card carrying out the forwarding decisions on board.[1]
- b. **Does Space-NIC replace an operator's existing internal network?**
Operators who already had a mesh could continue to use their present internal routing system and, at certain interconnection points, introduce Space-NIC in order to manage the traffic between networks; for operators who had no existing mesh, Space-NIC along with its supporting software could provide the onboard networking functions, on the understanding that their satellites were fitted with compatible communications links.
- c. **Does installing Space-NIC make a satellite compatible with every other network?**
For an interconnection to occur, you would need compatible physical links, supported protocols, available capacity, and agreement among the various operators. The translation layer developed by Space-NIC would adapt the framing and interface formats that it supports; it would not, however, eliminate incompatible optical hardware or automatically grant access to another company's network.[1, 7, 8]
- d. **Does customer data pass through Space-NIC?**
Yes, Space-NIC would handle and forward the customer's packets while the spacecraft was in flight; the ground system calculates routing tables and coordinates operations, a function that would not involve the customer payload data passing through Elevsoft's ground facilities.[1]
- e. **What happens if a scheduled table update is missed?**
Space-NIC continues using stored schedules, with preloaded backup routes and local repair providing fault tolerance when links fail. Simulations showed improved delivery with these mechanisms and unchanged delivery in one test using two-day-old orbital data. However, resilience depends on the stored scheduling horizon and actual network conditions; uninterrupted operation is not guaranteed. [2]

f. **What happens when a link fails unexpectedly?**

The design suggested would make use of preloaded backup routes in order to divert traffic locally whenever an alternative was still available. In the low-failure situations that were tested, scheduled routing combined with local repair resulted in a median packet-delivery rate of about 99.9%. Delivery rates dropped when more extensive failures occurred, and conventional dynamic routing proved to be the better option in the case of severe disruption.[2] Local repair will thus have to be accompanied by ground-based updates that take into account the changed network conditions.

5. Acronyms and Abbreviations

Acronym or abbreviation	Full Form/Meaning
CCSDS	Consultative Committee for Space Data Systems
DLR	Deutsches Zentrum für Luft- und Raumfahrt; German Aerospace Center
GAO	U.S. Government Accountability Office. Congress's audit agency; source of the PWSA figures cited.
GEO	Geostationary Earth orbit, as used in this paper
IP	Internet Protocol
IRIS ²	Infrastructure for Resilience, Interconnectivity and Security by Satellite. The EU's secure-connectivity constellation.
ISS	International Space Station
LEO	Low Earth Orbit
LFA	Loop-Free Alternate. A pre-computed backup next hop a satellite can switch to on its own when a link fails, chosen so traffic cannot circle back. Figure 3.2.5.
MILNET	Military Network
NEBULA	Network Established Beyond the Upper Limits of the Atmosphere. The SDA's network-layer standard, covering packet forwarding and scheduled routing tables.
OCT	Optical Communications Terminal. A laser terminal; also the SDA standard that makes such terminals interoperable.
OLSR	Optimized Link State Routing. An internet-style protocol in which nodes exchange messages about link changes; used as the comparison baseline in §3.2.
OSPF	Open Shortest Path First. The link-state routing protocol most terrestrial networks use.
PWSA	Proliferated Warfighter Space Architecture
SDA	Space Development Agency
SDN	Space Data Network, as used in this paper
SGP4	Simplified General Perturbations 4. The standard model for predicting a satellite's position from its TLE; used to generate the orbits in Figure 3.2.1.
TI-LFA	Topology-Independent Loop-Free Alternate. An extension of LFA that pre-computes a multi-hop detour when no single-hop alternative exists.
TLE	Two-line element set. The public orbital-parameter format published for every tracked satellite (here, from CelesTrak).

6. References

- [1] Elevsoft Silicon. *Technical Note ES-001, Rev. A: Space-NIC Architecture Overview*. September 6 2026, controlling technical document. <https://www.elevsoft.dev/architecture>.
- [2] Elevsoft Silicon. *Space-NIC Simulation Evidence Gallery*, version of 6 September 2026 (orbital-data snapshots of 4 and 6 September 2026, CelesTrak TLEs). <https://www.elevsoft.dev/evidence>.
- [3] C. Westphal, L. Han, and R. Li. "LEO Satellite Networking Relaunched: Survey and Current Research Challenges." arXiv:2310.07646, 11 October 2023. <https://arxiv.org/abs/2310.07646>
- [4] B. M. Leiner, V. G. Cerf, D. D. Clark, R. E. Kahn, L. Kleinrock, D. C. Lynch, J. Postel, L. G. Roberts, and S. Wolff. "A Brief History of the Internet." Internet Society, 1997. <https://www.internetsociety.org/internet/history-internet/brief-history-internet/>
- [5] J. Moy. *OSPF Version 2*. RFC 2328, IETF, April 1998. <https://www.rfc-editor.org/rfc/rfc2328>
- [6] T. Pan, T. Huang, X. Li, Y. Chen, W. Xue, and Y. Liu. "OPSPF: Orbit Prediction Shortest Path First Routing for Resilient LEO Satellite Networks." *Proceedings of the IEEE International Conference on Communications (ICC)*, Shanghai, May 2019, pp. 1–6. doi:10.1109/ICC.2019.8761611
- [7] Space Development Agency. *Optical Communications Terminal (OCT) Standard, Version 3.2.0*. Document 9100-001-11, 24 March 2025. https://www.sda.mil/wp-content/uploads/2025/04/SDA-OCT-Standard-v3.2.0_FINAL.pdf
- [8] Space Development Agency. *Optical Communications Terminal (OCT) Standard, Version 4.0.0*. 1 July 2024. https://www.sda.mil/wp-content/uploads/2024/07/SDA_OCT_Standard_4.0.0_final-20240701.pdf
- [9] Consultative Committee for Space Data Systems. *Space Packet Protocol*. CCSDS 133.0-B-2, Issue 2, June 2020, with Editorial Change 2, September 2024. <https://ccsds.org/Pubs/133x0b2e2.pdf>
- [10] European Cooperation for Space Standardization. *SpaceWire: Links, Nodes, Routers and Networks*. ECSS-E-ST-50-12C Rev. 1, 15 May 2019. <https://ecss.nl/standard/ecss-e-st-50-12c-rev-1-spacewire-links-nodes-routers-and-networks-15-may-2019/>
- [11] European Cooperation for Space Standardization. *SpaceFibre: Very High-Speed Serial Link*. ECSS-E-ST-50-11C, 15 May 2019. <https://ecss.nl/standard/ecss-e-st-50-11c-spacefibre-very-high-speed-serial-link/>
- [12] J. Postel, ed. *Internet Protocol*. RFC 791, IETF, September 1981. <https://www.rfc-editor.org/rfc/rfc791>
- [13] D. Giggenbach, F. Moll, C. Fuchs, and M. Brechtelsbauer. "Direct Optical High Speed Downlinks and Ground Station Networks for Small LEO Missions." *Proceedings of the 16th Ka- and Broadband Communications Conference*, Milan, 20–22 October 2010. <https://elib.dlr.de/66370/>
- [14] Amazon. "Amazon's Project Kuiper Completes Successful Tests of Optical Mesh Network in Low Earth Orbit." 14 December 2023. <https://www.aboutamazon.com/news/innovation-at-amazon/amazon-project-kuiper-oisl-space-laser-december-2023-update>
- [15] MDA Space. "MDA Selected as Prime Satellite Contractor for Telesat Lightspeed Advanced Constellation." 11 August 2023. <https://mda.space/article/mda-selected-as-prime-satellite-contractor-for-telesat>
- [16] Tesat-Spacecom. "TESAT Selected by MDA Space as Part of MDA AURORA Supply Chain to Deliver Optical Inter-Satellite Links." 28 May 2024. <https://www.tesat.de/news/press/941-tesat-selected-by-mda-space-to-deliver-oisl>
- [17] Kepler Communications. "Optical Inter-Satellite Link." 11 June 2024. <https://kepler.space/optical-inter-satellite-link/>

- [18] T. Beals. "Exploring a Space-Based, Scalable AI Infrastructure System Design." Google Research, 4 November 2025. <https://research.google/blog/exploring-a-space-based-scalable-ai-infrastructure-system-design/>
- [19] Axiom Space. "Axiom Space, Spacebilt Announce Orbital Data Center Node Aboard International Space Station." 16 September 2025. <https://www.axiomspace.com/release/axiom-space-spacebilt-announce-orbital-data-center-node>
- [20] United States Space Force. *U.S. Space Force Commercial Space Strategy*. April 2024. https://www.spaceforce.mil/Portals/2/Documents/Space%20Policy/USSF_Commercial_Space_Strategy.pdf
- [21] Space Systems Command. "Space Force Invests in Resilient Multi-Vendor Architecture to Build Next-Gen Space Data Network." 13 August 2026. <https://www.ssc.spaceforce.mil/Newsroom/Article-Display/Article/4572417/space-force-invests-in-resilient-multi-vendor-architecture-to-build-next-gen-sp>
- [22] U.S. Government Accountability Office. *DOD Satellite Communications: Reporting on Progress Needed to Provide Insight on New Approach*. GAO-25-107034, 4 March 2025. <https://www.gao.gov/products/gao-25-107034>
- [23] Space Development Agency. "Space Development Agency Makes Awards for Tranche 1 Transport Layer." 28 February 2022. <https://www.sda.mil/space-development-agency-makes-awards-for-tranche-1-transport-layer/>
- [24] European Space Agency. "ESA to Support the Development of EU's Secure Communication Satellites System." 16 December 2024. https://www.esa.int/About_Us/Corporate_news/ESA_to_support_the_development_of_EU_s_secure_communication_satellites_system
- [25] U.S. Government Accountability Office. *Laser Communications: Space Development Agency Should Create Links Between Development Phases*. GAO-25-106838, 26 February 2025. <https://www.gao.gov/products/gao-25-106838>
- [26] Space Development Agency. *Network Established Beyond the Upper Limits of the Atmosphere (NEBULA) Standard*. SDA_STD_NEBULA v3.05, 25 September 2024. <https://www.sda.mil/wp-content/uploads/2024/11/SDANEBULASStandardv3.05Final.pdf>
- [27] T. Clausen and P. Jacquet, eds. *Optimized Link State Routing Protocol (OLSR)*. RFC 3626, IETF, October 2003. <https://www.rfc-editor.org/rfc/rfc3626>