Internet Protocol Routing in Cis-Lunar Space

James DiCorcia	Mike Frandsen	Steve Dawson	Victor M. Aguero
Cambrian Works, Inc.	Cambrian Works, Inc.	Cambrian Works, Inc.	Cambrian Works, Inc.
Los Gatos, CA, USA	Los Gatos, CA, USA	Los Gatos, CA, USA	Los Gatos, CA, USA
james@cambrianworks.com	michael@cambrianworks.com	steve@cambrianworks.com	victor@cambrianworks.com

Over the last five years, integrating space systems into the global Internet paradigm has advanced significantly. Space constellations and networks are now commonplace; however, current implementations remain limited to networks managed as single terrestrially orchestrated systems or orchestrated hybrid mixtures of interacting space networks. This centralized orchestration approach constrains autonomous operations, requires continuous ground station visibility, and prevents direct peer-to-peer communication between space nodes. The need for standards-compliant inter-networking enabling space systems to function as direct Internet extensions is increasingly clear. Once achieved, space systems gain continuous real-time access to distributed data sources, on-orbit analytics, and artificial intelligence capabilities.

Cambrian Works has developed hardware and software applying Internet functionality in space without requiring terrestrial code rewrites. This technology allows constellation and in-space resources to serve as direct Internet extensions, supporting edge computing and distributed in-space processing. Space nodes equipped with this software interact with terrestrial Internet nodes identically to any other computer on the Internet. Applications include real-time integration of observations into terrestrial or space-based analyses regardless of processing location, continuous access to network-hosted data worldwide independent of ground station passes, and leveraging existing IP-based networking applications and security frameworks without modification.

This paper presents lessons learned from on-orbit space routers flown in late 2025 aboard multiple LEO missions. Performance evaluation used digital twin emulation tools characterizing software, hardware, and network behavior for multi-orbit constellation inter-networking. Flight validation demonstrated routing protocols, autonomous network reconfiguration, and operational procedures for Internet-extended space architectures.

The validated technology brings space architectures into direct functional alignment with the terrestrial Internet, eliminating custom ground segment middleware and enabling autonomous distributed operations. Expected outcomes include operational deployment for real-time space-ground data integration and standards-compliant networking throughout cis-lunar space.

Keywords—space networking, GigRouter, cis-lunar, Internet Protocol, edge computing, LEO constellations, digital twin, MISSE-21

1. Introduction

As space systems become more distributed, the value of standards-based networking increases. Future constellations and beyond-LEO architectures will need more than fixed point-to-point links and custom ground middleware. They will need practical ways to move data across changing topologies, to operate under intermittent contact, and to take advantage of existing networking tools where mission constraints allow. That is the larger motivation behind IP-based networking in space.

This work focuses on a more immediate question: what does it take to fly and operate an Internet-compatible networking payload in space today? The GigRouter mission on MISSE-21 aboard the ISS provided an opportunity to answer that question in a real operating environment. The mission established measured on-orbit behavior for IP-routing-capable hardware under real mission constraints and identified practical issues that matter when deploying this class of system in space. That focus is distinct from the way many customers use GigRouter today, where it is deployed as a standard product for onboard compute and internal data handling rather than for the broader networking use case emphasized here.

Read in that light, the mission data provide a practical record of how this class of hardware behaves on orbit. They show how the flight article performed across network and storage benchmarks, how it operated within a constrained power envelope, how the system behaved during a storage-side anomaly, and how onboard environmental monitoring helped guide interpretation. Those results are directly relevant not only to future routed space networks, but also to customers using GigRouter today as an onboard compute and internal data management platform. They also matter because larger multi-node architectures will ultimately depend on how real hardware behaves under limits on power, thermal rejection, telemetry, recovery cadence, and operator access.

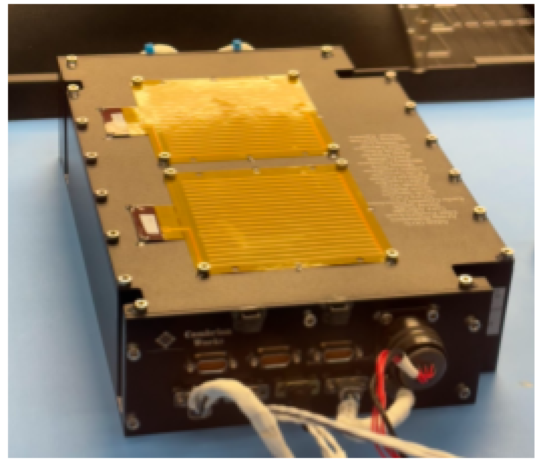

Figure 1. GigRouter flight article used for MISSE-21 / ISS operations.

2. Flight Article and Mission Context

The GigRouter flight article is an NVIDIA Orin-based platform that combines edge compute resources with standard IP networking interfaces, local solid-state storage, and supporting spacecraft interfaces. The key point here is not the full feature list, but that this class of hardware was flown and operated on orbit as a working compute-and-networking payload.

The unit is flying on MISSE-21 aboard the ISS, and the mission remains ongoing as of this writing. In that role, it is operating in vacuum, microgravity, radiation, and repeated thermal cycling associated with orbital sun and eclipse transitions. The mission also imposes practical operational limits. The flight article is mounted to a general-purpose enclosure that provides a non-ideal thermal interface. It operates under a 50 W continuous power ceiling, and is accessed through Ethernet after selecting high-bandwidth network connectivity over serial console access, because the MISSE integration supported one primary access path rather than both. Those constraints are central to the mission because they make the results operationally informative. They show what it takes to operate, monitor, and interpret a networking payload in space under realistic host-platform limits. [1]

Mission data from the GigRouter show

continuous operation exceeding three months at the time of this writing, successful over-the-air software installation and reconfiguration, and end-to-end exercise of the telemetry pipeline. Together, those results establish the mission as a sustained on-orbit operation rather than a brief functional checkout.

Figure 2. GigRouter connectorization prior to integration into MISSE 21 launch container (blue, foreground). [1]

3. Operating Conditions and Test Approach

The value of the results depends in part on the conditions under which the payload was operated. The mission worked within a 50 W continuous power ceiling, constrained commanding and uplink capacity, a weekly cadence for text and log telemetry downlink, and Ethernet-only operator access. Those conditions required logging, telemetry, benchmark design, and recovery planning to be built into the operating approach from the start. In future beyond-LEO missions, those same concerns are likely to matter even more.

The benchmark data used in this paper come primarily from the MISSE Benchmarks Report generated on 24 April 2026. That report includes weekly Ethernet loopback testing, input power tracking, thermal context during benchmark windows, SSD bandwidth measurements, and RADFET dosimetry. The report also maps the most relevant sensors to the hardware regions under test. The network controller chip's temperature sensor is associated with the interface panels (front panel), including the 10 GbE interface pair, while the PCIe_Chip sensor is the more relevant proxy for the rear panel interfaces and one of the solid-state drives. That sensor mapping matters when interpreting the storage-side anomaly discussed later in this paper.

The time history of these measurements produced a useful combination of performance data and operating context. The network results show what the interfaces did on orbit. The power data show whether those tests stayed inside the flight envelope. The storage data show both high-rate local data handling and device-availability behavior during automated testing. The thermal and dosimetry data show the environmental context that was also being recorded when interpreting what happened. Together, these measurements move the mission beyond a binary flight-success claim and provide a more useful record of on-orbit system behavior. [2]

Pair	Type	Runs	Fails	TX min/mean/max (Mb/s)	RX min/mean/max (Mb/s)
eth1i/eth3i	parallel	132	0	940 / 942 / 942	939 / 941 / 941
eth1i/eth3i	sequential	141	0	471 / 940 / 942	471 / 939 / 941
eth4x/eth5x	parallel	132	0	6292 / 8426 / 9266	6291 / 8425 / 9265
eth4x/eth5x	sequential	141	0	7659 / 9003 / 9398	7659 / 9003 / 9398
eth6i/eth7i	parallel	132	0	941 / 942 / 942	940 / 941 / 941
eth6i/eth7i	sequential	141	1	941 / 941 / 942	941 / 941 / 941
eth8i/eth9i	parallel	128	0	940 / 942 / 942	939 / 941 / 941
eth8i/eth9i	sequential	141	0	941 / 941 / 942	941 / 941 / 942

Figure 3. Weekly Ethernet loopback benchmark summary showing repeated on-orbit testing across four interface pairs. [2]

4. On-Orbit Results

4.1 Network interface performance

The clearest quantitative result from the mission is the Ethernet loopback benchmark set. 1,087 of 1,088 sub-tests met or exceeded their defined bandwidth thresholds; one sub-test on the eth6i/eth7i sequential pair fell below threshold. The 1 GbE-class interface pairs operated near line rate, with representative mean values around 941–942 Mb/s. The 10 GbE pair showed substantially higher throughput, with mean per-direction throughput of 8.43 Gb/s in parallel tests and 9.00 Gb/s in sequential tests, and maxima up to 9.398 Gb/s. These results show that the flight article's network interfaces were exercised repeatedly on orbit and produced strong, measurable performance under flight conditions.

The single failed sub-test indicates that one eth6i/eth7i sequential run fell below its defined bandwidth threshold. It should be interpreted as a benchmark threshold miss within an otherwise consistent loopback campaign, not as evidence of a persistent Ethernet-interface limitation. The network benchmark results are therefore best understood as interface-level loopback measurements collected during a broader automated test campaign. They show that an Internet-compatible networking payload can be benchmarked repeatedly on orbit, that the results are consistent enough to inform engineering decisions, and that the mission generated data relevant to both future architectures and current operational use cases.

▼eth6i/eth7i (PCIe Chip) – Per-Test Detail (273 tests)

Time (UTC)	Type	eth6i TX	eth6i RX	eth7i TX	eth7i RX	Pass?	PCIe_Chip temp start	PCIe_Chip temp during	PCIe_Chip temp end	PCIe_Chip temp after 2m	Max Δ°C
2026-04-22 23:20 UTC	S	941	941	941	941	PASS	57.7	59.0	60.4	61.4	3.6
2026-04-22 23:31 UTC	par	942	941	941	941	PASS	62.8	63.3	63.8	64.5	1.7
2026-04-23 00:20 UTC	S	941	941	941	941	FAIL	54.5	54.3	54.1	54.0	0.5
2026-04-23 00:31 UTC	par	942	941	942	941	PASS	53.8	53.8	53.8	53.7	0.0
2026-04-23 01:20 UTC	S	941	941	942	941	PASS	61.4	60.9	60.4	60.0	1.5
2026-04-23 01:31 UTC	par	941	941	942	941	PASS	59.2	59.0	58.9	58.4	0.8
2026-04-23 02:20 UTC	S	942	941	941	941	PASS	55.0	56.2	57.4	58.4	3.4
2026-04-23 02:31 UTC	par	942	941	942	941	PASS	60.2	60.8	61.4	62.2	2.0

Figure 4. Ethernet loopback table excerpt showing the single eth6i/eth7i sequential sub-test that fell below the defined bandwidth threshold. The marked result is a benchmark threshold miss within an otherwise consistent loopback campaign, not evidence of a persistent Ethernet-interface limitation. Corresponding temperature columns are included for operating context. [2]

4.2 Power behavior

The power record strengthens the network results because it ties them directly to the flight operating envelope. Over the benchmark interval from 17 April to 24 April 2026, the report shows mean input power of 39.1 W, a minimum of 32 W, and a maximum of 53 W. The summary also notes that no sustained excursions above 50 W were recorded during the reported period using a threshold of at least 20 seconds. The payload performed benchmark activity while remaining inside the intended sustained

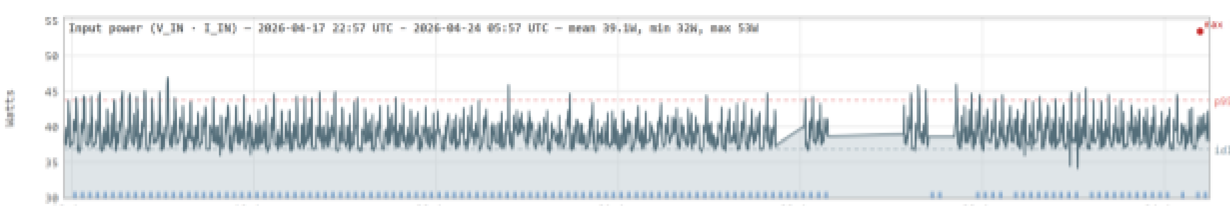

Figure 5. Measured input power during the benchmark interval. The mission remained within the intended sustained power envelope while network and storage benchmarks were performed on orbit. [2]

operating envelope.

This is one of the mission's most operationally relevant results. It demonstrates that high-rate networking and associated system activity can be exercised on orbit within defined spacecraft resource limits. This is true for future hosted payloads and integrated spacecraft networks.

4.3 Storage-side anomaly during automated testing

Separately from the Ethernet loopback threshold miss, the storage benchmark campaign encountered a storage-side anomaly in which one SSD dropped from the system during high-cadence automated testing. The payload remained operational after the drive dropped, and the test sequence continued to return useful information. The event also exposed a dependency in the telemetry collection architecture: some collection applications depended on files located on the dropped drive, so temperature and power data were no longer collected in parallel for portions of the campaign.

This behavior is operationally important. Flight systems are not evaluated only by nominal throughput; they are also evaluated by how they behave when a subsystem becomes unavailable. In the first storage-drop event, the payload continued operating and data collection continued sufficiently to preserve the test campaign. In the latest event, the payload remained operational, but much of the background telemetry collection stopped because the dropped drive contained binaries required by the collection process. A commanded reboot did not restore the drive, but the drive came back online after a requested complete power cycle. In both known cases, the drive loss was recoverable, but the second event showed that telemetry collection can still depend on where collection software and supporting files reside.

The root cause remains under engineering review. The observed behavior appears consistent with a recoverable driver, interface, or device-state condition rather than permanent storage loss. Thermal cycling remains a plausible contributing factor under engineering review, but the current data do not establish thermal causation. The operational lesson is that storage state, driver behavior, thermal environment, and recovery procedures need to be considered together when interpreting on-orbit anomalies.

Time (UTC)	Read (MB/s)	Write (MB/s)	Duration	Power (during test)
2026-04-24 18:34 UTC	skipped / mount unavailable			
2026-04-24 17:34 UTC	skipped / mount unavailable			
2026-04-24 16:34 UTC	skipped / mount unavailable			
2026-04-24 15:34 UTC	skipped / mount unavailable			
2026-04-24 14:34 UTC	skipped / mount unavailable			
2026-04-24 13:33 UTC	skipped / mount unavailable			
2026-04-24 12:34 UTC	skipped / mount unavailable			
2026-04-24 11:34 UTC	skipped / mount unavailable			
2026-04-24 10:34 UTC	skipped / mount unavailable			
2026-04-24 08:36 UTC	4183	2073	107s	–
2026-04-24 07:36 UTC	4190	2070	107s	–
2026-04-24 06:37 UTC	4184	2078	109s	–
2026-04-24 05:37 UTC	4174	2065	108s	42W avg / 45W max
2026-04-24 04:37 UTC	4206	2072	107s	44W avg / 47W max
2026-04-24 02:36 UTC	4184	2106	107s	43W avg / 46W max
2026-04-24 00:37 UTC	4182	2160	110s	43W avg / 47W max
2026-04-23 23:36 UTC	4181	2107	107s	43W avg / 45W max

Figure 6. SSD benchmark sequence and representative pre-anomaly test context. [2]
(a, left) Hourly SSD benchmarks show repeated successful read/write tests followed by automated skips once the SSD mount became unavailable. Successful runs immediately before the unavailable-mount condition produced approximately 4.17–4.21 GB/s read throughput and 2.07–2.16 GB/s write throughput.
(b, below) A representative successful run at 2026-04-24 05:37 UTC shows 4174 MB/s read and 2065 MB/s write, with 42 W average and 45 W maximum power during the test. The plotted temperature and power traces provide operating context, but do not establish root cause for the later mount-unavailable condition.

4.4 Storage subsystem performance

Storage behavior was relevant both as a benchmark result and as part of the anomaly investigation. For SSD0, the report shows mean read throughput of 3977 MB/s and mean write throughput of 2238 MB/s, with corresponding ranges across the reported runs. The report also includes SSD1 runs in the same weekly period at approximately 2917–2923 MB/s read and 3291–3298 MB/s write. These measurements show that the flight article remained capable of high-rate local data handling on orbit under the same constrained power and thermal conditions seen elsewhere in the mission.

The storage anomaly also showed the operational value of automated testing and continued data collection after a device-state change. After one SSD dropped from the system, the payload continued operating and continued returning core test outputs, although parallel telemetry collection was degraded where collection applications depended on the dropped drive. A prior similar event recovered after a full power-off/power-on cycle, and the latest case also recovered after a complete power cycle. This behavior does not by itself prove a fully redundant storage architecture, but it does show that the payload could continue operating through a storage-side anomaly without ending the test campaign.

4.5 Radiation dosimetry

The onboard RADFET measurements add another important layer of environmental context. The benchmark report shows that dosimetry was collected on orbit alongside performance and housekeeping telemetry, with approximately hourly sampling sent to ground downsampled from a 5-minute sampling rate. That matters because it shows the flight article was instrumented to capture environmental conditions on orbit alongside operational performance data in a situation where environmental conditions directly influence operational data interpretation.

Figure 7 shows an inverse relationship between the raw RADFET-derived dose values and board temperature when the data are plotted on a highly zoomed scale. The RADFET sensors include temperature compensation, but the remaining temperature sensitivity is still visible at small mrad-level variations. The expected accumulated dose for the week in this environment is small, on the order of 10 mrad, while the sensors are intended to measure much larger lifetime exposure, up to tens of krads. On a 0–10 krad scale, the plotted variation would be effectively flat, consistent with the low accumulated dose expected in this orbit.

For future missions, appropriate RADFET sensors should be selected that are either very sensitive to low levels of radiation or are appropriate to the expected radiation exposure over the life of the mission. These RADFET sensors remain appropriate when the goal is detecting larger accumulated dose or rare higher-dose events. For MISSE-21, where the accumulated dose over the reported week was low, the primary lesson is not that the sensors were inappropriate, but that highly zoomed raw RADFET-derived values should be interpreted in the context of sensor precision and board temperature. If the goal is simply anticipating component lifetimes, seeing that exposure didn’t grow significantly is adequate to have confidence that components haven’t been subjected to high radiation amounts. But, appropriate sizing becomes especially relevant for networking payloads beyond LEO, where longer recovery timelines, less frequent operator access, and more mission-specific radiation environments increase the value of correlating performance data with local environmental telemetry.

Figure 7. RADFET dosimetry and board ambient temperature, week-over-week comparison. The RADFET data provide onboard environmental context for the flight article. The comparison shows why short-window dose trends should be interpreted alongside board temperature, especially when the expected accumulated dose over the interval is small relative to temperature-sensitive variation. [2]

5. Extending the Lessons to Larger Network Architectures

The on-orbit mission addresses a specific set of questions: can this class of payload be operated, observed, and benchmarked in space under real flight constraints? Larger network architectures introduce different operational concerns, including how route options evolve as geometry changes, how traffic continuity is maintained across changing topologies, and how similar systems scale beyond a single hosted payload.

Cambrian Works uses PlexusMesh and Retina to explore those larger-architecture questions. PlexusMesh is the operational mesh-routing software, while Retina is the visualization and digital-twin environment used to observe and analyze network behavior across larger scenarios. In this paper, their role is to connect the MISSE flight results to the broader architectural problem: flight data establish how GigRouter behaves on orbit, while emulation helps evaluate how similar hardware and software patterns behave as topology, route availability, and network scale become more complex.

Those roles remain distinct. The ISS results provide measured on-orbit behavior for the hardware and operations discussed in this paper. PlexusMesh and Retina provide architectural context for extending those lessons to larger multi-node and future beyond-LEO systems.

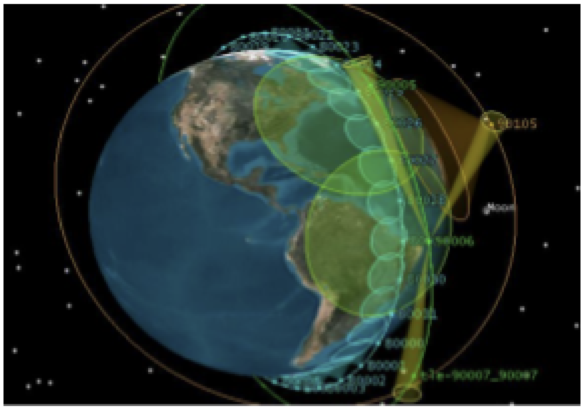

Figure 8. Retina visualization of a multi-orbit PlexusMesh network scenario in cis-lunar space. [3]

6. Operational Lessons

Several operational lessons stand out from the mission.

First, diagnostic and recovery access matters. The payload was successfully operated through Ethernet access, which was selected for its higher bandwidth and suitability for nominal networking operations. In the MISSE integration, only one primary access path could be supported, so Ethernet was selected instead of integrator-provided serial console access. That trade supported nominal operations, but reduced access to a live low-level terminal console independent of the network stack. Such access is especially useful for diagnosing boot, configuration, software-update, and network-stack failure modes where Ethernet connectivity may not be available.

Second, power-aware operations have to be built in from the beginning. The 50 W continuous power ceiling was not just a requirement on paper. It shaped what could be tested and when. Executing network and storage benchmarks within that envelope makes clear that compute, routing, telemetry, and maintenance activity must be planned together in any operational mission.

Third, automated test design and recovery planning matter as much as nominal benchmark performance. Both times a drive dropped out, the payload continued operating and recording core test outputs even though many parallel metric recordings stopped in the second case. Immediate intervention was not necessary for core functionality to continue. That outcome is operationally relevant: the system did not stop producing useful data simply because one local storage device became unavailable. It also highlights the importance of storage layout, backup copies of critical files, and avoiding unnecessary dependencies between telemetry collection and non-critical storage paths.

Fourth, anomaly interpretation depends on correlating device state, software state, environmental telemetry, and recovery history. A prior similar storage-drop event recovered after a full power-off/power-on cycle, while a commanded reboot alone did not restore the drive in the latest case. As with the first case, the drive was restored to full functionality after a power cycle. Driver issues, power issues, firmware behavior, or operation near the high end of thermal cycling all remain plausible contributing factors under engineering review, but the current data do not yet establish specific root cause(s). The broader lesson is that storage telemetry, driver state, thermal context, recovery procedures, and data-return cadence need to be designed into the operating concept together.

Fifth, environmental monitoring needs to be designed for interpretation, not just collection. The RADFET data show that on-orbit dosimetry was captured alongside performance telemetry, but also that simply looking at the raw data is misleading if it's not interpreted in the context of the accuracy of the sensors. For future missions, LEO and VLEO especially, RADFETs can be selected that are appropriate for expected dosimetry trends when radiation exposure data are used to support performance or degradation analysis. The RADFETs that were used on this mission may be appropriate for GEO or other higher radiation orbits.

Taken together, these lessons extend the mission's value beyond the benchmark numbers alone. They show not only what was measured, but what matters when deploying networking hardware as a practical flight system: constrained power, automated testing, recovery posture, environmental telemetry, and the ability to continue collecting useful data when a subsystem enters an anomalous state.

7. Conclusion

The central result of this work is that an Internet-compatible networking payload was flown, operated, and benchmarked on orbit under practical mission constraints. The mission produced repeated on-orbit network performance measurements, power data that tied those results to the operating envelope, storage benchmarks that showed broader system capability, and environmental measurements that provided context for interpreting system behavior.

The results establish that GigRouter can support high-rate networking and local data handling while operating within a constrained hosted-payload environment. They also show how much of the operational value comes from the surrounding system design: automated test execution, compact telemetry products, recovery planning, environmental monitoring, and the ability to continue collecting useful data when a subsystem enters an anomalous state.

The storage-side anomaly is part of that result. During high-cadence automated storage testing, one SSD became unavailable. The payload remained operational, but the event exposed a storage-layout dependency: some background telemetry collection depended on files located on the dropped drive, so portions of the parallel temperature and power data collection stopped. A prior similar event recovered after a full power-off/power-on cycle, and the latest case also recovered after a complete power cycle; root cause remains under engineering review. That experience is operationally relevant because future space networks will depend not only on nominal throughput, but on whether the underlying hardware can be operated, observed, and recovered within real limits on power, telemetry cadence, recovery cadence, environmental exposure, and operator access.

In that sense, the ISS mission provides operational evidence for the practical utility of IP-based space networking systems. The flight results establish how this class of hardware behaves on orbit today. The emulation work extends those lessons to larger multi-node and future beyond-LEO architectures, where networking payloads will need to combine high-rate data movement, local processing, environmental awareness, and robust recovery posture across mission-specific operating environments.

Citations

[1] Cambrian Works, “Cambrian Works Announces Upcoming GigRouter Launch to the International Space Station,” *SpaceNews*, 2025.
[2] Cambrian Works, “MISSE Benchmarks Report — 20260424,” internal on-orbit benchmark report, Apr. 24, 2026.
[3] Cambrian Works, “PlexusMesh and Retina Architecture Overview,” internal technical note, 2026.