TU Delft

Delft University of Technology

Document Version
Final published version

Licence
CC BY

Citation (APA)
Shreedhar, T., Zavodovski, A., Kumar, A., & Tuomela, S. (2026). Rebooting the Internet from Orbit. In *SIGCOMM 2026 - Proceedings of the 2026 ACM SIGCOMM 2026 Conference* (pp. 1880-1886). (SIGCOMM 2026 - Proceedings of the 2026 ACM SIGCOMM 2026 Conference). Association for Computing Machinery (ACM). https://doi.org/10.1145/3789240.3827587

Important note
To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright
In case the licence states “Dutch Copyright Act (Article 25fa)”, this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership.
Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse
Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy
Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

This work is downloaded from Delft University of Technology.

Rebooting the Internet from Orbit

Tanya Shreedhar*
Delft University of Technology
Delft, The Netherlands
t.shreedhar@tudelft.nl

Aleksandr Zavodovski*
University of Oulu
Oulu, Finland
aleksandr.zavodovski@oulu.fi

Abhishek Kumar
University of Jyväskylä
Jyväskylä, Finland
abhishek.k.kumar@jyu.fi

Sanna Tuomela
University of Oulu
Oulu, Finland
sanna.tuomela@oulu.fi

ABSTRACT

Low Earth Orbit (LEO) megaconstellations are the first network in three decades being built greenfield at a planetary scale, and the window for shaping their architecture is closing. Once operators lock their inter-satellite and ground-gateway data planes to IP conventions, the field will have lost its last realistic opportunity to deploy a non-legacy inter-domain architecture without coordinating a billion endpoints. The terrestrial Internet has resisted architectural change for two decades, as IPv6 illustrates: standardized in 1998, it crossed 50% on Google's measurement only in 2026. Clean-slate designs such as ICN, SCION, and XIA, and extensibility frameworks such as the Trotsky Processor (TP) and Plutarch, have not yet broken this rigidity. We argue that non-terrestrial networks (NTNs), in synergy with 6G, offer the last realistic substrate for the long-anticipated revitalization of the Internet's architecture, and that TP is the most credible primary vehicle for that reboot, complemented at the edges by translation approaches. We sketch two concrete blueprints, with TP either native to the orbital segment ("Trotsky in the Sky") or anchored on the ground and using the NTN as a globe-spanning logical pipe between L3.5 nodes ("Trotsky in the Trenches"). A capacity analysis of the Starlink constellation shows that the V3 generation entering service in 2026 makes capacity a non-issue: fewer than 2K V3-class satellites would absorb today's entire yearly inter-domain traffic, well within the ≈10K currently in orbit.

CCS CONCEPTS

• **Networks** → **Network architectures**; *Network design principles*; *Network protocol design*; *Naming and addressing*; *Routing protocols*; *Network mobility*.

KEYWORDS

Internet architecture, non-terrestrial networks, LEO satellites, clean-slate design, network evolvability, 6G

ACM Reference Format:
Tanya Shreedhar, Aleksandr Zavodovski, Abhishek Kumar, and Sanna Tuomela. 2026. Rebooting the Internet from Orbit. In *4th Workshop on LEO Networking and Communication (LEO-NET '26), August 17–21, 2026, Denver, CO, USA.* ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3789240.3827587

*Equal co-first authors; these authors contributed equally.



This work is licensed under a Creative Commons Attribution 4.0 International License.
SIGCOMM '26, August 17–21, 2026, Denver, CO, USA
© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2467-1/26/08
https://doi.org/10.1145/3789240.3827587

1 INTRODUCTION

Starlink operates over 10K active satellites [28] and has crossed 9M subscribers [36]. Its monopoly is mitigated by Eutelsat OneWeb [18], which operates a LEO constellation of over 600 satellites alongside Eutelsat's heritage GEO fleet, and by Amazon Leo, which has begun deploying toward a 3K+ satellite constellation [2]. Taken together, these operators are assembling the first planetary-scale network built greenfield in three decades, and the engineering decisions that will define it are being made right now. Today, that lock-in is still pliable: satellite generations refresh on multi-year cycles, inter-satellite link protocols remain proprietary and operator-defined, and ground software stacks are being rewritten as constellations scale. Within a single hardware refresh, however, IP semantics tend to harden into ASICs, ISL framing, and billing systems, and from that point onward displacing them faces the same coordination cost that has kept the terrestrial Internet on IPv4 for two decades. Once those commitments are baked into the next satellite generation, the field will have spent its last realistic chance to deploy a non-legacy inter-domain architecture without coordinating a billion endpoints. We argue that the satellite networking community should treat the next three to five years as an architectural deadline, not a research direction.

The case for the terrestrial Internet's non-evolvability was made by Clark et al. [13] and developed by others over two decades. The mechanisms are now understood: backward-compatibility constraints, the ISP coordination problem, and the absence of any actor with both the reach and the incentive to deploy something new. IPv6 illustrates the timescale: specified in 1998, it crossed 50% on Google's measurement only in March 2026, eighteen years after public measurement began [7, 19]. This inertia is not unique to IPv6: even backward-compatible upgrades diffuse slowly, as longitudinal measurements of Multipath TCP show its adoption confined to a small fraction of hosts years after standardization [4, 33]. In response, researchers have proposed clean-slate designs that redesign the Internet from the ground up. Information-Centric Networking (ICN) [21, 40] shifts the unit of communication from hosts to named content; SCION [11] provides path control and isolation for inter-domain routing;[1] and the eXpressive Internet Architecture (XIA) [3] supports multiple communication models in a single

[1]To our knowledge, SCION is the only practically deployed next-generation architecture, used in Switzerland primarily by financial institutions [34].

header. None is backward-compatible with IP, and each requires substantial coordinated adoption to reach its full potential, a daunting prospect at Internet scale. Where new protocols have gained ground, indirection was often the enabler: QUIC circumvented transport-layer ossification by encapsulating in UDP and encrypting its headers [22], and now carries diverse web, cloud-storage, and video workloads at scale [32]. To sidestep these barriers, two indirection-based families have emerged. The Trotsky Processor (TP) [27] introduces a thin coexistence layer at L3.5 that lets multiple Internet architectures run side by side, with TP-aware nodes wrapping and unwrapping packets so that legacy IP segments behave as transparent overlays. Translation-oriented approaches such as Plutarch [14], in contrast, embrace heterogeneity through explicit boundary functions that convert between architectural "contexts". The two families have different failure modes: TP preserves end-to-end semantics but requires participating endpoints, whereas translation works at the boundary yet risks losing context, especially security context, in transit. Neither family has yet broken the Internet's rigidity, primarily for want of coordinated adoption across a fragmented ecosystem.

These constellations sit within the broader class of Non-Terrestrial Networks (NTNs), which also encompasses high-altitude platforms (HAPs) and unmanned aerial vehicles (UAVs), and together extend connectivity to regions where terrestrial infrastructure is sparse or absent. Crucially, this reach lets NTNs establish direct, long-range links between almost any two points on Earth, bypassing legacy routers and forming logical pipes that align naturally with TP's overlay model for inter-domain traffic. Growing onboard compute further enables satellites to host sophisticated software, performing routing, caching, and general-purpose edge computing in orbit [8, 9, 38]. Taken together, these capabilities suggest that satellites can serve as native deployment platforms for clean-slate architectures, extensibility points, and translators, allowing the direct deployment of protocols such as ICN or SCION in orbit. 6G strengthens this outlook by integrating terrestrial and non-terrestrial networks and supporting direct communication between handheld devices and satellites, with edge compute available at both ends of the path [16, 39].

Our position is that, among the candidate vehicles surveyed above, TP is the most credible primary driver of an NTN-based reboot of Internet architecture [6], with translation approaches playing a complementary role at the edges where heterogeneity is irreducible. Section 2 contrasts three candidate deployment models for NTN, namely direct deployment of a single clean-slate design, TP-style universal extensibility, and Plutarch-style translation, and argues that TP enables "permanent revolution" where translation alone risks losing security context in transit. Section 3 turns this position into a deployment blueprint, distinguishing an in-orbit variant ("Trotsky in the Sky") from a ground-anchored variant ("Trotsky in the Trenches") that uses the constellation as a long-haul logical pipe between L3.5 nodes. Section 4 grounds the vision quantitatively, showing that the V3 generation entering service in 2026 reduces the capacity question to a non-issue, and closes with a values-oriented reflection on why keeping Internet architecture open to change is worth the effort.

2 ENGENDERING EVOLUTION

Repeated failures in architectural redesign have left the Internet research community with a kind of collective fatigue: some see further attempts as futile, others see initiating change as worthwhile in itself. Our stance is the second one, and the goal is not only technical: it is to give future Internet users a choice of communication paradigm, on technical and value grounds alike.[2] Several deployment paths can exploit the NTN substrate to drive architectural change. We examine the three most plausible options: direct deployment of a single clean-slate design, TP-style universal extensibility, and translation-based interoperability between heterogeneous architectures. Each strikes a different balance between fidelity to the new architecture, backward compatibility with the incumbent IP world, and the realistic prospect of incremental adoption.

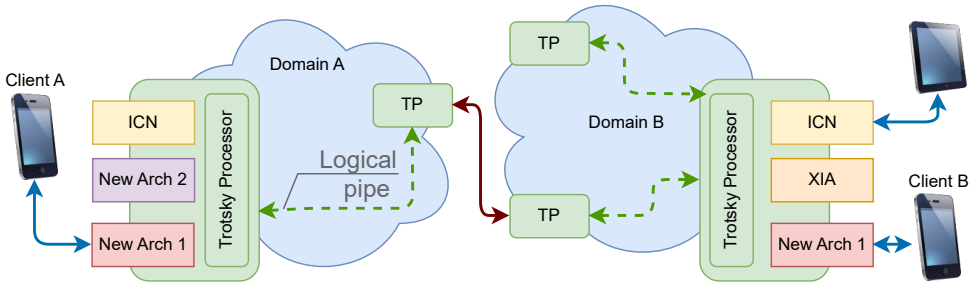
Direct deployment. The simplest option is to pick a single clean-slate architecture and deploy it natively across the orbital segment. An ICN-style deployment fits naturally when the constellation acts as a globe-spanning CDN, with content names rather than host addresses driving caching and forwarding decisions on board [8]. SCION, with its path control, source authentication, and resilience against prefix hijacks, suits operators selling guaranteed connectivity to security-sensitive customers. In both cases, the satellite is a first-class node of the new architecture rather than a transit element: ICN caches and forwarders sit on board adjacent to the radio interface, and SCION beaconing, path construction, and packet authentication run inside the constellation's own control plane. The administrative homogeneity of a single operator's constellation becomes an asset rather than a liability: the entire in-orbit system can be upgraded synchronously, at planetary scale, and without inter-AS negotiation. The downside is being locked into one architecture managed by a single, most likely commercial, entity. A multi-operator world makes this worse: a Starlink committed to SCION, an Eutelsat OneWeb running native ICN, and an Amazon Leo holding to IPv6 would re-create the inter-domain coordination problem at orbital scale, without the half-century of operational practice that has kept BGP working between terrestrial ASes. Proposals for the democratization and decentralization of space-based NTNs [29, 30] may eventually alleviate this concern, but the opposite is also plausible: NTNs may become entrenched, just as the terrestrial Internet has, and the limits of any once-promising clean-slate design will then surface earlier than expected.

The Trotsky Processor (TP). Presented schematically in Figure 1b, TP overcomes the Internet's architectural ossification by introducing a new layer, L3.5, on which multiple clean-slate architectures can coexist [27]. L3.5 decouples inter-domain from intra-domain communication and hosts pluggable network processors, each implementing a different next-generation architecture (e.g., ICN or XIA) on top of an unmodified IP forwarding plane. A network processor is a self-contained module that owns its own packet format, addressing, and forwarding logic; a TP node located at an architectural boundary demultiplexes incoming traffic to the right processor and re-encapsulates it for the next IP-only hop. TP's two defining properties, backward compatibility with the installed IP

[2] A recent initiative in this spirit is CATRIN (Controllable, Accountable, Transparent: the Responsible Internet) [10, 20], which emphasizes transparency of network management and user control over data handling.



a) Translation, as suggested in [31]: a SCION network operates in native mode, while IP packets are translated at the boundary to and from SCION.



b) The Trotsky Processor [27]: a pluggable L3.5 layer hosts network processors that implement next-generation architectures (here "New Arch 1" and "New Arch 2") alongside legacy IP, enabling "permanent revolution".

Figure 1: Two approaches to accelerating the evolution of the Internet.

base and extensibility through the simultaneous hosting of multiple clean-slate designs, together enable incremental deployment without breaking changes. L3.5 nodes that share a clean-slate architecture are stitched together by logical pipes and overlays that traverse legacy IP segments, preserving the native packet format end-to-end. NTNs, augmented by upcoming 6G deployments, offer an unusually clean placement for L3.5: their global reach can connect almost any two domains in a single hop at the inter-domain boundary, precisely where TP needs its extensibility points to sit. Growing onboard satellite compute, together with mobile-oriented edge compute in the 6G ground segment, provides the processing budget required to run multiple network processors at once.

Translation approaches. Illustrated in Figure 1a, translation-oriented designs enable communication across heterogeneous architectures by rewriting protocol headers at the boundaries between architectural regions. A pioneering example is Plutarch [14], which introduces *contexts*, distinct regions with their own addressing, routing, and naming, and *interstitial functions* that translate packets, identifiers, and lookups from one regime to the next. XIA [3] takes a complementary route, equipping every packet with a list of typed identifiers (host, service, content, …) so that any node can fall back to whichever identifier type it understands, effectively translating at the packet header rather than at a dedicated boundary. Such designs are valuable wherever diverse networks (the legacy terrestrial Internet, novel clean-slate islands, and emerging NTNs) must interoperate without converging on a single unified protocol. They do, however, pay a structural price: information that has no counterpart in the target architecture is, by construction, lost in translation.

2.1 Trotsky vs. Translators

For the future Internet, the desired condition is that clients are free to select their communication paradigm. Direct deployment (Section 2) offers limited support for this principle, so we concentrate on TP and translation, juxtaposed in Figures 1a and 1b. Ribeiro et al. [31] advocate translation as the driver of revitalization and take direct issue with TP's implementation strategy [27]. The critique centers on "IP-centrality": the logical pipes and overlays that TP uses to forward traffic across IP domains are seen as undermining crucial properties of clean-slate designs, illustrated through SCION, whose secure routing conflicts with overlay links. Their proposed remedy is direct translation: IP traffic is converted into SCION on entry to a SCION-native domain and back to IP on exit, letting the SCION domain operate in native mode and retain its resistance to prefix hijacks and DDoS attacks.

This translation approach has two problems. First, the IP client originating the traffic is unaware of SCION's existence and cannot exercise the path-selection control that is a defining feature of SCION. Second, from the SCION domain's own security perspective, it would arguably be safer not to forward translated IP traffic of unknown origin at all. Consider also the more realistic reverse scenario: a SCION host reaches another SCION host through an IP-only transit segment, forcing the boundary routers to translate SCION to IP on egress and back on ingress. The security context, including path proofs, source authentication, and packet authenticators, is dropped at the first translator and cannot be reconstructed at the second. This is strictly worse than the logical-pipe approach criticized above, which would have carried the SCION packet across the IP segment intact. And from a practical perspective, SCION networks are operated by security-sensitive entities [34] who are unlikely to allow pass-through IP-translated traffic, especially given the abundance of IP-native routes that already connect any pair of IP hosts worldwide.

Springboards and scaffolding. TP's overlays are not an endorsement of "IP-centrality" but a "necessary evil" in an IP-dominated world, where any new architecture must traverse the legacy substrate to reach the next node that speaks its language. Logical pipes and overlays are deployment scaffolding, designed to be retired hop by

hop as the clean-slate fabric accumulates enough native deployment to stand on its own. Crucially, they preserve the native packet format end-to-end: a SCION packet entering an overlay leaves it as the same SCION packet, with its path authenticators and source proofs intact, while only the IP envelope around it changes hop by hop. This also blunts a sharper objection: that path selection is forfeited wherever clean-slate traffic crosses IP, leaving islands of SCION stranded in seas of IP. The NTN substrate is precisely what drains those seas, since two L3.5 nodes can be joined by a single operator-defined satellite hop rather than a long chain of IP-only ASes, confining any loss of path control to one link instead of a multi-AS path. Translation cannot offer the same guarantee: whenever an advanced architecture is downgraded to a simpler one, or when two architectures lack a shared abstraction, the unmatched information is irrecoverably dropped at the boundary.

Contrasting objectives. Translators are necessary in a highly diverse and fragmented environment: where subtleties can be lost, connectivity is found. In this sense, the scenario of Ribeiro et al. [31] (Figure 1a) is ahead of its time and becomes plausible once the "thin waist" of the Internet, the IP layer, is eventually unbuckled, leaving small islands of IP within a clean-slate Internet that connect mainly through translators. The objective of translators is to handle heterogeneity that already exists; the TP framework, conversely, aims to *break* the homogeneity of the current Internet. Universal extensibility makes TP the primary driver of an NTN-inspired revitalization.

Synergy. The most defensible outcome is a synergistic one, in which TP supplies the substrate for "permanent revolution" and translation handles the irreducible heterogeneity that accumulates on top of it. TP is well-positioned to absorb translation as a special case: its L3.5 nodes can host translation-oriented network processors, as the original prototype already demonstrated by hosting XIA [27]. In a mature deployment, the same NTN-anchored L3.5 fabric can simultaneously carry native ICN, native SCION, and legacy IP, with translators inserted only where architectures genuinely cannot share a common abstraction, for instance at the boundary of a residual IP island.

3 TECHNOLOGY BLUEPRINT

In what follows, by TP we mean not the prototype described in the original article,[3] but the abstraction it embodies: a network element that retains backward compatibility with IP while exposing universal extensibility through L3.5. The concrete realization of TP in NTNs depends on its placement (in-orbit, ground-station, or cloud-based) and on the performance characteristics demanded by the supported architectures. We next examine the two most plausible placements.

3.1 Trotsky in the Sky

Depicted as variant (a) in Figure 2, this is the most optimistic scenario for enabling "permanent revolution", and at first glance it may appear "too good to be true". Onboard compute is growing steadily, prompting proposals for placing mobile core functions [26], CDNs [8], and general-purpose edge compute in orbit [24, 25, 38]. Deploying TP in orbit to handle ICN caching, SCION path control, or XIA routing would bring next-generation architectures directly to satellite-terminal owners and to mobile clients with satellite connectivity.[4] Such clients could rely fully on clean-slate designs for the satellite-to-satellite segment, bypassing the legacy Internet entirely. In-orbit deployment would also provide low latencies, in-network caching and compute, and access to advanced services in underserved areas, narrowing the digital divide.

Launching Trotsky into orbit brings substantial challenges. The most pressing is constrained satellite compute: clean-slate architectures can be resource-hungry, as illustrated by SCION's per-packet cryptographic operations. A split architecture can mitigate this, deploying only latency-critical and lightweight TP components in orbit while offloading heavier processing to ground facilities; similarly, a dual-layer caching scheme with smaller on-board caches and larger ones at ground stations can ease the storage constraints of in-space ICN. Frequent satellite handovers also pose a challenge for stateful architectures, which depend on context and client state being available at the serving node, requiring an architecture-specific context-and-state handover mechanism between successive satellites [38]. The concrete state to migrate is architecture-specific: for ICN, the Pending Interest Table and a hot subset of the content store; for SCION, active path-segment registrations and authentication contexts; for any L3.5 framework, the per-flow processor state that ties a connection to its serving processor instance. Energy and thermal budgets compound the problem: per-packet cryptographic verification and large on-board caches both compete for the same constrained satellite bus, and an operator unwilling to dedicate that power to clean-slate processing will quietly default to plain IP forwarding.

However, the primary concern with space deployment is not technical. First, it entrusts the evolution of the Internet to a single dominant operator, raising clear monopoly risks: Starlink is at present the only LEO constellation offering Internet service directly to consumers at significant scale, and proposals for democratization and decentralization [15, 29, 30] are at best long-term mitigations. A related open question is which next-generation architectures the operator(s) would choose to host on board. A more flexible service model would be on-demand deployment, in which clients express interest in a specific architecture and it is provisioned dynamically on the satellites entering their coverage area; in practice, however, such an offering is likely to collapse to a small, operator-curated set of paradigms, with non-trivial control overhead. Second, in-orbit computation remains expensive, and the business case for hosting unfamiliar clean-slate architectures is highly uncertain; in the current geopolitical climate, business considerations may nevertheless be overshadowed by the security and resilience advantages that new architectures bring. A clearer commercial incentive may come from differentiation: an operator can package native SCION as security-as-a-service or in-orbit ICN as content-delivery acceleration, turning architectural openness into a premium product rather than a pure cost.

[3] The prototype implementation described in [27] is not publicly available.

[4] Starlink already offers a *Direct to Cell* service for unmodified 4G LTE and 5G handsets [35]; bandwidth and feature set remain limited today, and the situation is expected to improve as 6G reaches mobile phones.

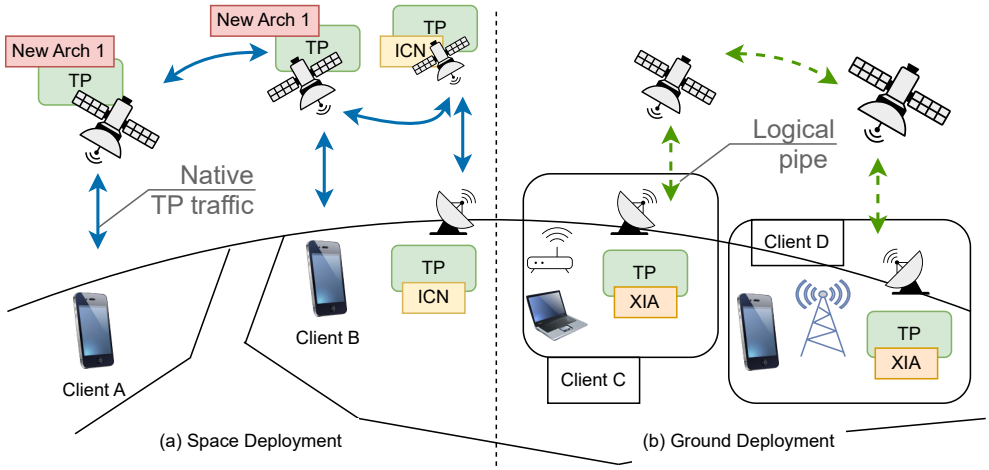


Figure 2: Two outlooks for TP deployment: (a) natively in orbit, or (b) on the ground, using the NTN as a global logical pipe between L3.5 nodes.

3.2 Trotsky in the Trenches

In practice, the orbital deployment described above may never materialize, owing to a mix of technical, economic, and political factors. A more pragmatic alternative is to keep TP on the ground and treat space-based NTNs as a globe-spanning logical pipe between L3.5 nodes, as depicted in variant (b) of Figure 2. A satellite-connected client can then reach any ground-based TP node in the world along a route that traverses only the satellite network, sidestepping the multi-hop terrestrial Internet altogether. The price of this pragmatism is the loss of some in-network benefits of clean-slate architectures, most visibly ICN's in-network caching, which now happens only at the endpoints rather than along the path.

Even so, this scenario retains most of the advantages that motivated the deployment in the first place. The global reach of LEO constellations remains intact: regions that terrestrial networks struggle to serve are still covered, which matters for bridging digital divides and supporting global-scale applications. LEO satellite links can outperform long-haul terrestrial fiber over intercontinental distances, lowering the round-trip cost of TP traffic and improving responsiveness for latency-sensitive workloads. Satellite pipes also offer alternative paths for TP traffic, improving resilience against terrestrial outages and geopolitical disruptions; this is especially valuable for security-oriented architectures such as SCION and for critical use cases such as emergency communications. Finally, the resulting TP topology is simpler than the multi-layered terrestrial Internet: with fewer intermediate ISPs, each carrying its own policies and performance variability, network behavior becomes more predictable and easier to manage. Concretely, consider a SCION-secured transaction between two financial institutions on different continents. With ground-anchored TP and an NTN logical pipe, the SCION packet leaves the originating ground TP node, traverses the constellation as the opaque payload of the NTN's native (IP-encapsulated) data plane, and reaches the SCION peer at the remote ground TP node intact, preserving the path proofs and source authentication that motivated the choice of SCION in the first place [34]. The same pattern serves ICN: an Interest issued from a satellite-connected handset reaches the nearest ground TP node, is forwarded natively across the inter-domain TP fabric to a producer or off-path cache, and the resulting Data packet returns along the symmetric route, with the NTN providing the long-haul leg in both directions.

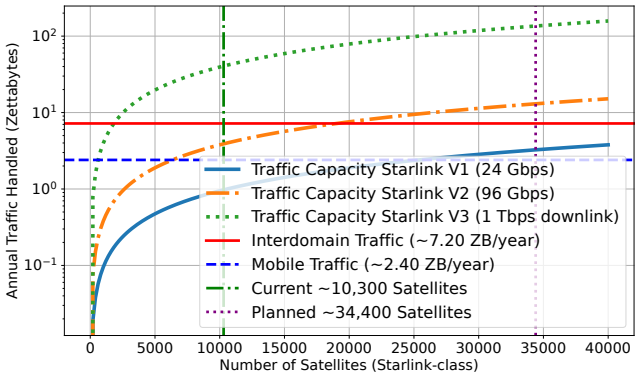


Figure 3: Annual aggregate throughput as a function of constellation size, for V1 (24 Gbps), V2 (96 Gbps), and V3 (≈1 Tbps downlink) Starlink generations [36, 37], compared against 2026 estimates of global inter-domain (≈7.2 ZB/year [12]) and mobile (≈2.4 ZB/year [17]) traffic. Vertical lines mark the ≈10K active and ≈34K planned Starlink constellation sizes [28].

4 CONCLUSIVE DISCUSSION

A natural question is whether space-based NTNs can carry meaningful volumes of next-generation traffic. To investigate, we analyze the throughput capacity of the Starlink constellation, since it publishes the most detailed performance data of any operator [36]. We present the results in Figure 3, comparing per-generation aggregate throughput against ≈2.4 ZB/year of global mobile traffic (extrapolated from Ericsson's Q4 2025 figure of ≈200 EB/month [17]) and ≈7.2 ZB/year of inter-domain traffic (extrapolated from Cisco's Annual Internet Report [12]). The original V1 generation is no longer sufficient on its own: at the full planned ≈34K satellites it carries only ≈3.2 ZB/year, enough to absorb yearly mobile demand but

well short of inter-domain volumes. The V2 generation, already in deployment, crosses the inter-domain line at roughly 19K satellites, comfortably within the planned envelope. The V3 generation, entering service in 2026 with ≈1 Tbps downlink per satellite [37], makes the capacity question essentially trivial: fewer than 2K V3-class satellites suffice to absorb the entire yearly inter-domain traffic, and the currently active constellation of ≈10K satellites [28] would, once upgraded, exceed inter-domain demand by more than five times. We read this as an order-of-magnitude feasibility bound rather than a claim that any one constellation should carry all traffic: aggregate downlink is an upper bound on, not a substitute for, deliverable inter-domain transit, and the ground-anchored variant in Section 3 in particular depends on inter-satellite and backhaul capacity rather than user-facing downlink. The conversion to ZB/year likewise assumes full utilization, with many satellites over oceans at any instant, yet even a conservative duty-cycle derating leaves the V3 generation comfortably above yearly inter-domain demand.

Satellite megaconstellations such as Starlink and OneWeb pose environmental challenges, especially for radio and visual astronomy [23], suggesting that in a pessimistic scenario their growth and applications could be curtailed. Mitigations have already been proposed, at least for nanosatellites [1]. At the same time, Earth was never an ideal site for astronomical observation, and space technology has steadily lifted instruments into orbit, with Hubble and the James Webb Space Telescope as the best-known examples; the proposed Lunar Crater Radio Telescope (LCRT) on the radio-quiet far side of the Moon takes this even further [5]. Paradoxically, by adding to terrestrial radio interference, megaconstellations may accelerate this migration off Earth, which is in any case becoming inevitable as instrument sensitivity outgrows what natural conditions on the ground can support.

Trotsky's approach also has its own limits: over a long-enough horizon, its L3.5 abstractions may run out of flexibility and fail to accommodate the newest designs without breaking changes. At present, however, it remains the most promising vehicle for the "permanent revolution", and our next step is to build a practical implementation of TP and release it as open source. Several adjacent open questions deserve community attention. The prototype of [27] is not public, so the design's claims about backward compatibility and multi-processor coexistence have not been independently re-evaluated on contemporary traffic; an open reference implementation would also let in-orbit benchmarks attach concrete numbers to the per-packet costs of SCION cryptography and ICN content-store lookups against current and projected satellite power budgets. Handover protocols that migrate architecture-specific state between adjacent satellites are similarly undefined, with the closest analogues being mobile core context transfer [26] and inter-satellite handover frameworks for edge workloads [38]. A further open question is the inter-architecture security model: when traffic crosses an L3.5 boundary into an encapsulating overlay or, at the edges, into a translator, the trust assumptions of both endpoints must be carried with it, and today there is no shared vocabulary for expressing, attenuating, or auditing those assumptions across heterogeneous processors. The standardization path is the least developed of all: 3GPP, the IETF, and the constellation operators themselves each hold partial stakes, but no forum currently combines them, and the window in which they could coordinate is the same three to five years during which the orbital data plane is being frozen.

This paper aims to open a dialogue on the unique opportunity that the rapidly developing NTNs offer for the long-awaited revitalization of Internet architecture. Since NTNs are still young and transformative, their developers may be more open to architectural innovation and to adopting extensibility solutions such as Trotsky, taking a leadership role in establishing new communication paradigms. It is also time to overcome the "learned helplessness" of clean-slate research, in which the hope of practical impact appears to have faded, and to view communication as more than a utilitarian medium in which only latency and throughput matter. Communication paradigms also reflect values; our mission is to empower Internet users to choose freely how they communicate, which is not possible without flexibility at the architectural layer.

ACKNOWLEDGEMENTS

This research is supported by a research grant from the University of Oulu and the Research Council of Finland under Profi6 (grant 336449), and has been supported in part by the Business Finland 6GSoft project under Grant 8541/31/2022.

REFERENCES

[1] Ali Abedi, Joshua Sanz, Mariya Zheleva, and Anant Sahai. 2024. From foe to friend: the surprising turn of mega constellations in radio astronomy. In *Proceedings of the 23rd ACM Workshop on Hot Topics in Networks*. Association for Computing Machinery, New York, NY, USA, 10–16.

[2] Amazon. 2026. Project Kuiper. https://www.aboutamazon.com/news/innovation-at-amazon/what-is-amazon-project-kuiper.

[3] Ashok Anand, Fahad Dogar, Dongsu Han, Boyan Li, Hyeontaek Lim, Michel Machado, Wenfei Wu, Aditya Akella, David G Andersen, John W Byers, et al. 2011. XIA: An architecture for an evolvable and trustworthy Internet. In *Proceedings of the 10th ACM Workshop on Hot Topics in Networks*. Association for Computing Machinery, New York, NY, USA, 1–6.

[4] Florian Aschenbrenner, Tanya Shreedhar, Oliver Gasser, Nitinder Mohan, and Jörg Ott. 2021. From Single Lane to Highways: Analyzing the Adoption of Multipath TCP in the Internet. In *2021 IFIP Networking Conference (IFIP Networking)*. IEEE, Espoo, Finland, 1–9.

[5] Harry Baker. 2025. NASA plans to build a giant radio telescope on the 'dark side' of the moon. Here's why. https://www.livescience.com/space/space-exploration/nasa-plans-to-build-a-giant-radio-telescope-on-the-dark-side-of-the-moon-heres-why.

[6] Hari Balakrishnan, Sujata Banerjee, Israel Cidon, David Culler, Deborah Estrin, Ethan Katz-Bassett, Arvind Krishnamurthy, Murphy McCauley, Nick McKeown, Aurojit Panda, et al. 2021. Revitalizing the public internet by making it extensible. , 18–24 pages.

[7] Internet Society Pulse Blog. 2026. 18 Years Later, IPv6 Reaches Majority. https://pulse.internetsociety.org/en/blog/2026/04/18-years-later-ipv6-reaches-majority/.

[8] Rohan Bose, Saeed Fadaei, Nitinder Mohan, Mohamed Kassem, Nishanth Sastry, and Jörg Ott. 2024. It's a bird? it's a plane? it's CDN!: investigating content delivery networks in the LEO satellite networks era. In *Proceedings of the 23rd ACM Workshop on Hot Topics in Networks*. Association for Computing Machinery, New York, NY, USA, 1–9.

[9] Rohan Bose, Jinwei Zhao, Tanya Shreedhar, Jianping Pan, and Nitinder Mohan. 2026. Investigating Web Content Delivery Performance over Starlink. In *Proceedings of the ACM Web Conference 2026* (United Arab Emirates) *(WWW '26)*. Association for Computing Machinery, New York, NY, USA, 5165–5176. https://doi.org/10.1145/3774904.3792227

[10] Catrin. 2026. A Responsible Internet to Increase Trust in the Digital World. https://www.catrin.nl/.

[11] Laurent Chuat, Markus Legner, David Basin, David Hausheer, Samuel Hitz, Peter Müller, and Adrian Perrig. 2022. *The Complete Guide to SCION: From Design Principles to Formal Verification*. Springer, Cham, Switzerland. https://doi.org/10.1007/978-3-031-05288-0

[12] Cisco. 2020. Cisco Annual Internet Report (2018–2023) White Paper. https://www.cisco.com/c/en/us/solutions/collateral/executive-perspectives/annual-internet-report/white-paper-c11-741490.html.

[13] David D Clark, Karen Sollins, John Wroclawski, and Ted Faber. 2003. Addressing reality: an architectural response to real-world demands on the evolving internet.

In *Proceedings of the ACM SIGCOMM workshop on Future directions in network architecture*. Association for Computing Machinery, New York, NY, USA, 247–257.

[14] Jon Crowcroft, Steven Hand, Richard Mortier, Timothy Roscoe, and Andrew Warfield. 2003. Plutarch: an argument for network pluralism. *ACM SIGCOMM Computer Communication Review* 33, 4 (2003), 258–266.

[15] Leo Eichhorn, Tanya Shreedhar, Aleksandr Zavodovski, and Nitinder Mohan. 2021. Distributed Ledgers for Distributed Edge: Are We There Yet?. In *Proceedings of the Interdisciplinary Workshop on (de)Centralization in the Internet (IWCI '21)*. Association for Computing Machinery, New York, NY, USA, 26–33. https://doi.org/10.1145/3488663.3493687

[16] Ericsson. 2024. 6G: Follow the journey to future networks. https://www.ericsson.com/en/6g.

[17] Ericsson. 2025. Ericsson Mobility Report, November 2025. https://www.ericsson.com/en/reports-and-papers/mobility-report/reports/november-2025.

[18] Eutelsat Group. 2026. Eutelsat OneWeb. https://oneweb.net/.

[19] Google. 2026. IPv6 Statistics. https://www.google.com/intl/en/ipv6/.

[20] Cristian Hesselman, Paola Grosso, Ralph Holz, Fernando Kuipers, Janet Hui Xue, Mattijs Jonker, Joeri De Ruiter, Anna Sperotto, Roland van Rijswijk-Deij, Giovane CM Moura, et al. 2020. A responsible internet to increase trust in the digital world. *Journal of Network and Systems Management* 28, 4 (2020), 882–922.

[21] Van Jacobson, Diana K Smetters, James D Thornton, Michael F Plass, Nicholas H Briggs, and Rebecca L Braynard. 2009. Networking named content. In *Proceedings of the 5th international conference on Emerging networking experiments and technologies*. Association for Computing Machinery, New York, NY, USA, 1–12.

[22] Adam Langley, Alistair Riddoch, Alyssa Wilk, Antonio Vicente, Charles Krasic, Dan Zhang, Fan Yang, Fedor Kouranov, Ian Swett, Janardhan Iyengar, et al. 2017. The QUIC transport protocol: Design and internet-scale deployment. In *Proceedings of the conference of the ACM special interest group on data communication*. Association for Computing Machinery, New York, NY, USA, 183–196.

[23] Andy Lawrence, Meredith L Rawls, Moriba Jah, Aaron Boley, Federico Di Vruno, Simon Garrington, Michael Kramer, Samantha Lawler, James Lowenthal, Jonathan McDowell, et al. 2022. The case for space environmentalism. *Nature Astronomy* 6, 4 (2022), 428–435.

[24] Israel Leyva-Mayorga, Marc Martinez-Gost, Marco Moretti, Ana Pérez-Neira, Miguel Ángel Vázquez, Petar Popovski, and Beatriz Soret. 2023. Satellite edge computing for real-time and very-high resolution earth observation. *IEEE Transactions on Communications* 71, 10 (2023), 6180–6194.

[25] Chengcheng Li, Yasheng Zhang, Renchao Xie, Xuekun Hao, and Tao Huang. 2021. Integrating edge computing into low earth orbit satellite networks: Architecture and prototype. *IEEE Access* 9 (2021), 39126–39137.

[26] Yuanjie Li, Hewu Li, Wei Liu, Lixin Liu, Yimei Chen, Jianping Wu, Qian Wu, Jun Liu, and Zeqi Lai. 2022. A case for stateless mobile core network functions in space. In *Proceedings of the ACM SIGCOMM 2022 Conference*. Association for Computing Machinery, New York, NY, USA, 298–313.

[27] James McCauley, Yotam Harchol, Aurojit Panda, Barath Raghavan, and Scott Shenker. 2019. Enabling a permanent revolution in internet architecture. In *Proceedings of the ACM special interest group on data communication*. Association for Computing Machinery, New York, NY, USA, 1–14.

[28] Jonathan C. McDowell. 2026. Jonathan's Space Pages: Starlink Statistics. https://planet4589.org/space/con/star/stats.html.

[29] Veronica Muriga, Swarun Kumar, Akshitha Sriraman, and Assane Gueye. 2024. A road map for the democratization of space-based communications. In *Proceedings of the 23rd ACM Workshop on Hot Topics in Networks*. Association for Computing Machinery, New York, NY, USA, 17–24.

[30] Seoyul Oh and Deepak Vasisht. 2024. A call for decentralized satellite networks. In *Proceedings of the 23rd ACM Workshop on Hot Topics in Networks*. Association for Computing Machinery, New York, NY, USA, 25–33.

[31] Guilherme Ribeiro, Luis Pedrosa, Salvatore Signorello, and Fernando MV Ramos. 2024. Internet architecture evolution: Found in translation. In *Proceedings of the 23rd ACM Workshop on Hot Topics in Networks*. Association for Computing Machinery, New York, NY, USA, 300–307.

[32] Tanya Shreedhar, Rohit Panda, Sergey Podanev, and Vaibhav Bajpai. 2022. Evaluating QUIC Performance Over Web, Cloud Storage, and Video Workloads. *IEEE Transactions on Network and Service Management* 19, 2 (2022), 1366–1381. https://doi.org/10.1109/TNSM.2021.3134562

[33] Tanya Shreedhar, Danesh Zeynali, Oliver Gasser, Nitinder Mohan, and Jörg Ott. 2022. A Longitudinal View at the Adoption of Multipath TCP. arXiv preprint arXiv:2205.12138. https://doi.org/10.48550/arXiv.2205.12138

[34] SSFN. 2021. SCION-powered SSFN: STRENGTHENING THE SWISS FINANCE SECTOR. https://www.scion.org/ssfn-scion/.

[35] Starlink. 2024. Direct to Cell. https://www.starlink.com/business/direct-to-cell.

[36] Starlink. 2025. 2025 Progress Report. https://www.starlink.com/public-files/starlinkProgressReport_2025.pdf.

[37] Tom's Hardware. 2025. SpaceX shows off massive new V3 Starlink satellites with 1 Tbps downlink and 60 Tbps per-launch capacity. https://www.tomshardware.com/service-providers/network-providers/spacex-shows-off-massive-new-v3-starlink-satellites-expanded-technology-will-deliver-gigabit-internet-to-customers-for-the-first-time-and-enable-60-terabits-per-second-downlink-capacity.

[38] Renchao Xie, Qinqin Tang, Qiuning Wang, Xu Liu, F Richard Yu, and Tao Huang. 2020. Satellite-terrestrial integrated edge computing networks: Architecture, challenges, and open issues. *Ieee Network* 34, 3 (2020), 224–231.

[39] Aleksandr Zavodovski, Nitinder Mohan, Abhishek Kumar, Sasu Tarkoma, Jussi Kangasharju, and Ari T Pouttu. 2025. Transforming the Internet with 6G: Towards Architectural Extensibility. In *2025 Joint European Conference on Networks and Communications & 6G Summit (EuCNC/6G Summit)*. IEEE, IEEE, Poznań, Poland, 1–6.

[40] Lixia Zhang, Alexander Afanasyev, Jeffrey Burke, Van Jacobson, KC Claffy, Patrick Crowley, Christos Papadopoulos, Lan Wang, and Beichuan Zhang. 2014. Named data networking. *ACM SIGCOMM Computer Communication Review* 44, 3 (2014), 66–73.