

Distributed Mobility Management
Internet-Draft
Intended status: Informational
Expires: 31 August 2026

Z. Lai
Y. Hou
Q. Wu
H. Li
Tsinghua University
27 February 2026

Mobility Management for Inter-SNO Sharing in LEO Satellite Networks
draft-lai-dmm-sno-collaboration-00

Abstract

Deploying a low-Earth orbit (LEO) satellite network typically requires substantial financial investment and long development cycles. To shorten deployment timelines and alleviate economic burdens, collaborative constellation deployment has gained attention. In such a model, multiple satellite network operators (SNOs) contribute their respective constellations to jointly deliver LEO connectivity services. Despite these potential benefits, the high mobility of LEO satellites introduces operational complexity. As satellites move rapidly across coverage regions, terrestrial users are frequently reassigned to different access satellites, which may be connected to distinct SNO core networks. These cross-operator handovers, referred to as inter-SNO handovers, can negatively affect service continuity and overall network robustness. To address these challenges, this document outlines a mobility management framework aimed at supporting cooperative LEO constellation operation. The framework separates the satellite access layer from individual SNO core infrastructures, thereby permitting flexible mapping between shared access satellites and multiple operator cores. Through this decoupled architecture, user sessions can remain associated with their original SNO core networks whenever operational conditions permit. The framework incorporates a dynamic SNO association strategy that seeks to limit the frequency of inter-SNO handovers. Furthermore, in scenarios where such handovers cannot be avoided, a proactive optimization procedure is employed to shorten interruption duration and enhance the efficiency of the handover process.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 31 August 2026.

Copyright Notice

Copyright (c) 2026 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	3
2. Characteristics of LEO Satellite Networks	4
2.1. LEO satellite networks	5
2.2. Constructing an LEO network is costly and time-consuming	5
3. Impacts of LEO Mobility on Infrastructure Sharing	5
3.1. Telecom infrastructure sharing	6
3.2. Collaboration in LEO networks	6
3.3. New Challenges of Inter-Constellation Sharing	6
3.4. Limitations of existing handover optimizations	8
4. The Proposed Mobility Management Framework	8
4.1. Shared LEO access network	9
4.2. Independent SNO cores	10
4.3. Adaptive space-ground associations	10
4.4. Key techniques behind the proposed framework	11
5. Reducing Disruption Time During Handovers	11
5.1. Baseline inter-SNO handover optimizer	11
5.2. Optimizing handovers by proactive authentication and tunneling	13
6. Conclusion	14
7. IANA Considerations	14
8. Security Considerations	14
9. References	15
9.1. Normative References	15
9.2. Informative References	15

Acknowledgements	16
Contributors	16
Authors' Addresses	16

1. Introduction

Over the past few years, a number of emerging commercial space enterprises have initiated the deployment of Low-Earth orbit (LEO) satellite constellations. Representative examples include SpaceX's Starlink, Eutelsat OneWeb, and Amazon's planned LEO constellation. These constellations rely on satellites operating at high orbital speeds relative to the Earth's surface, typically around 8 km/s. As a result, a single satellite remains within the visibility range of a specific terrestrial location only for a limited duration, often on the order of several minutes. To provide continuous and global Internet connectivity under such mobility constraints, operators must deploy constellations with a large number of satellites distributed across multiple orbital planes. Achieving this scale of deployment is operationally demanding. It requires completion of multiple regulatory and engineering processes [satellite_constellation_launch], including orbital slot coordination, spectrum authorization, manufacturing, and repeated launch campaigns. Even for established operators such as Starlink and OneWeb, several years, typically between three and six, are required before global service availability can be realized.

In order to shorten deployment timelines and control operational expenditures, constellation-level cooperation among satellite network operators (SNOs) has recently attracted increasing attention [decentralized_satellite_networks]. Under this model, multiple SNOs share selected infrastructure resources and jointly provide LEO connectivity services to terrestrial users. The underlying concept is similar to infrastructure sharing practices in terrestrial mobile networks, where operators co-utilize physical facilities or network components to improve resource efficiency and reduce capital investment. Although such collaboration can lower entry barriers and accelerate rollout, its application to LEO constellations introduces additional complexity. Unlike terrestrial networks, LEO constellations exhibit strong infrastructure-level dynamics due to rapid satellite movement. Consequently, users may be frequently reassigned across access satellites that are connected to different SNO core networks. These repeated inter-SNO handovers pose substantial challenges to maintaining service continuity and overall network stability in a shared LEO environment.

To address the service disruptions introduced by inter-SNO handovers and to support stable constellation cooperation, this document specifies an inter-networking architecture for collaborative SNO

operation. At the architectural level, the framework separates the satellite access layer from individual SNO core infrastructures. This separation permits shared access satellites to be dynamically mapped to different operator core networks, rather than being permanently bound to a single SNO. As a result, user sessions can continue to be served by their original SNO core networks whenever network conditions allow. The framework is realized through two complementary mechanisms. The first mechanism governs dynamic SNO association, enabling access satellites to be selectively connected to appropriate core networks in order to limit unnecessary inter-SNO handovers. The second mechanism focuses on optimizing the handover procedure itself. In scenarios where inter-SNO handovers cannot be avoided, a prediction-assisted proactive process is applied. By utilizing pre-established SNO association information, the network-level handover operations can be prepared in advance, thereby shortening service interruption duration and improving handover efficiency.

The aim of this document is to describe the challenges posed by inter-constellation sharing in emerging LEO satellite networks, particularly the frequent inter-SNO handovers caused by infrastructure-level dynamics, and to describe a mobility management framework to mitigate their negative impacts. Based on our analysis of collaborative SNO deployments and handover behaviors, we provide design insights and architectural considerations for enabling seamless inter-constellation sharing. We hope this document will serve as a useful reference for future standardization efforts on inter-SNO inter-networking mechanisms in shared LEO satellite networks.

2. Characteristics of LEO Satellite Networks

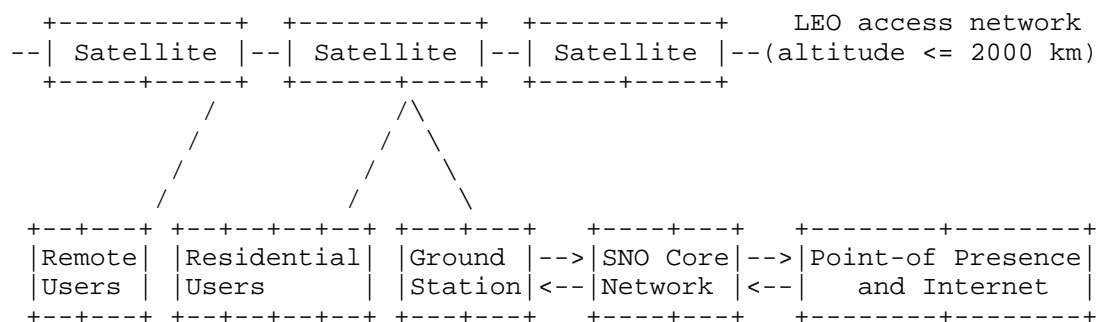


Figure 1: Networking architecture of operational LEO networks.

2.1. LEO satellite networks

Figure 1 illustrates a representative networking structure adopted by contemporary LEO satellite network operators (SNOs) [democratizing_leo_satellite_network_measurement]. A typical LEO constellation can be divided into two primary components. The first component is a satellite-based access segment, formed by a large constellation of LEO satellites that deliver last-mile connectivity to terrestrial users. The second component is a terrestrial core network, which performs operator-specific control and management functions, including user authentication, billing, IP address management, and traffic exchange with the public Internet. When a user connects to the Internet via a satellite terminal, the data path generally proceeds as follows. User traffic is transmitted through one or more LEO satellites to a ground station, forwarded to the corresponding SNO core network for processing, and then routed to the broader Internet. The reverse path follows a similar sequence. In geographically remote regions where direct satellite-to-ground visibility may be limited, inter-satellite links (ISLs) enable traffic to be relayed across multiple satellites before reaching an appropriate ground gateway.

2.2. Constructing an LEO network is costly and time-consuming

Since LEO satellites travel at high speeds relative to the Earth, sustained service availability in any given region requires continuous satellite replacement within the coverage footprint. To achieve this level of persistence, an SNO must deploy a sufficiently large constellation before commercial operation can begin. In practice, this implies the construction of a globally distributed satellite network. Reaching such deployment scale is neither straightforward nor rapid. Establishing a mega-constellation involves multiple regulatory, technical, and logistical steps [satellite_constellation_launch]. Operators are required to obtain approval for orbital configurations from national regulators, such as the Federal Communications Commission, and to secure spectrum usage rights across different jurisdictions. In parallel, satellite production, integration, and launch scheduling must be coordinated, each of which introduces additional lead time. Consequently, SNOs often experience an extended pre-operational phase, which may span several years before full commercial services become available.

3. Impacts of LEO Mobility on Infrastructure Sharing

3.1. Telecom infrastructure sharing

Infrastructure sharing among operators has long been adopted in terrestrial mobile systems as a means of improving deployment efficiency. Through infrastructure sharing arrangements, multiple operators can jointly utilize assets such as spectrum licenses and access facilities to expand coverage, lower capital expenditures, and accelerate network rollout. Such cooperation has been implemented in various markets. For example, T-Mobile Poland and PTK Centertel entered into a long-term agreement, lasting approximately fifteen years, to jointly operate portions of their mobile access network [comarch-orange-T-Mobile]. In another case, major South Korean operators including LG Uplus, KT, and SK Telecom reached an agreement to share 5G network infrastructure in sparsely populated coastal and agricultural regions [analysysmason-network-sharing].

3.2. Collaboration in LEO networks

Drawing on the experience of infrastructure sharing in terrestrial mobile systems, similar concepts are now being considered in the context of satellite Internet. Given the substantial financial investment and lengthy timelines required to deploy LEO constellations, both satellite network operators (SNOs) and the research community have begun to investigate inter-constellation cooperation [satellitetoday-iridium-oneweb], [decentralized_satellite_networks]. Under such arrangements, participating SNOs coordinate the use of spectrum and other network resources to jointly construct and operate LEO infrastructures. This cooperative model offers several practical benefits. By pooling resources, individual SNOs are no longer required to independently deploy a fully standalone global constellation, which can significantly reduce the time needed to reach service readiness. In addition, coordinated spectrum planning and joint infrastructure utilization may help streamline regulatory processes, thereby accelerating overall deployment schedules.

3.3. New Challenges of Inter-Constellation Sharing

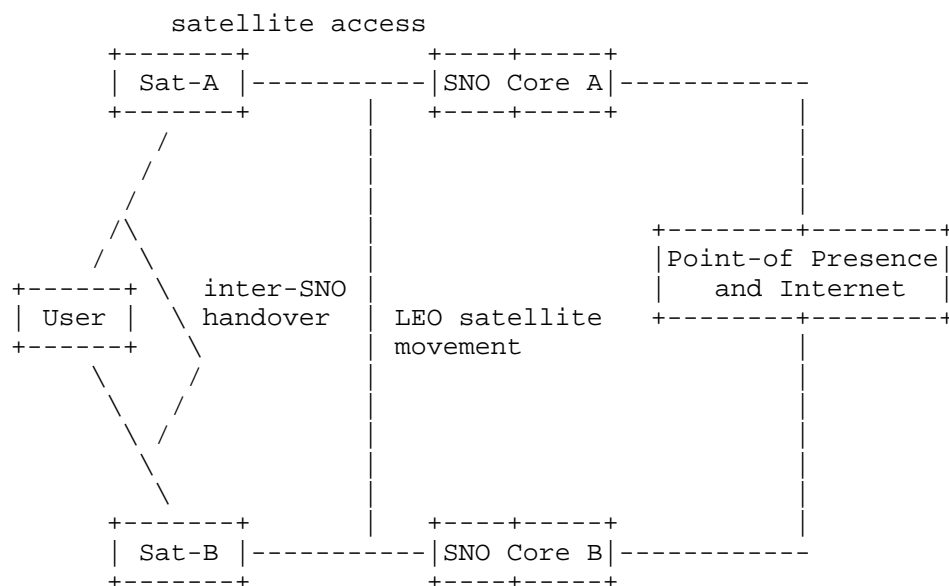


Figure 2: Current constellation sharing.

Figure 2 illustrates a representative constellation-sharing architecture as described in [decentralized_satellite_networks]. In this model, collaborating operators such as SNO A and SNO B retain primary control over their respective user bases. Users of SNO A preferentially access satellites owned by SNO A, and only associate with satellites of SNO B when local coverage from SNO A is temporarily unavailable. In such deployments, the satellite access layer remains tightly bound to the corresponding operator's core network. Due to the rapid orbital movement of LEO satellites, satellite visibility for each operator fluctuates over time across a given geographic region. Consequently, users may repeatedly handover between satellites affiliated with different SNOs. These handovers give rise to inter-SNO handovers at the network level. Compared with handovers occurring within a single operator domain, inter-SNO handovers involve additional signaling and control operations. Procedures such as user re-registration, authentication re-establishment, and reassignment of IP addresses must be executed before service can resume. These extra steps introduce short-term service interruptions, which can degrade session continuity and reduce the perceived stability of the shared LEO network.

3.4. Limitations of existing handover optimizations

From a network perspective, repeated inter-SNO handovers can be interpreted as users continuously alternating between their home LEO network and a partner operator's infrastructure. Such behavior resembles roaming across administrative domains, but occurs at a much higher frequency due to the orbital dynamics of LEO constellations. Mobility management has been extensively studied in the broader mobile networking community. Prior efforts include the design of mobility management protocols, proxy-assisted handover acceleration techniques, and adaptations of mobility mechanisms for LEO environments. For instance, protocols such as [RFC6275] and [RFC5944] are primarily concerned with minimizing service interruption during individual handover events. However, these mechanisms generally address the latency of a single handover and do not attempt to reduce how often such handovers occur. Proxy-based optimization approaches, including those described in [practical_traffic_management_lte_wifi], typically depend on stateful proxy entities deployed at relatively stable access points in terrestrial networks. In contrast, access satellites in LEO constellations are inherently mobile. As user attachment points continuously change, maintaining persistent proxy state becomes impractical, which limits the applicability of such designs in satellite scenarios. More recently, mobility schemes tailored for LEO networks have been proposed, such as [skycastle_leo_mobility]. These solutions focus on addressing mobility within a single operator's constellation. When extended to inter-constellation collaboration environments, however, they remain vulnerable to frequent cross-operator handovers and the associated disruptions. Taken together, the high frequency of inter-SNO handovers and the constraints of existing mobility optimization mechanisms highlight the need for approaches specifically designed to support seamless operation in shared LEO constellations.

4. The Proposed Mobility Management Framework

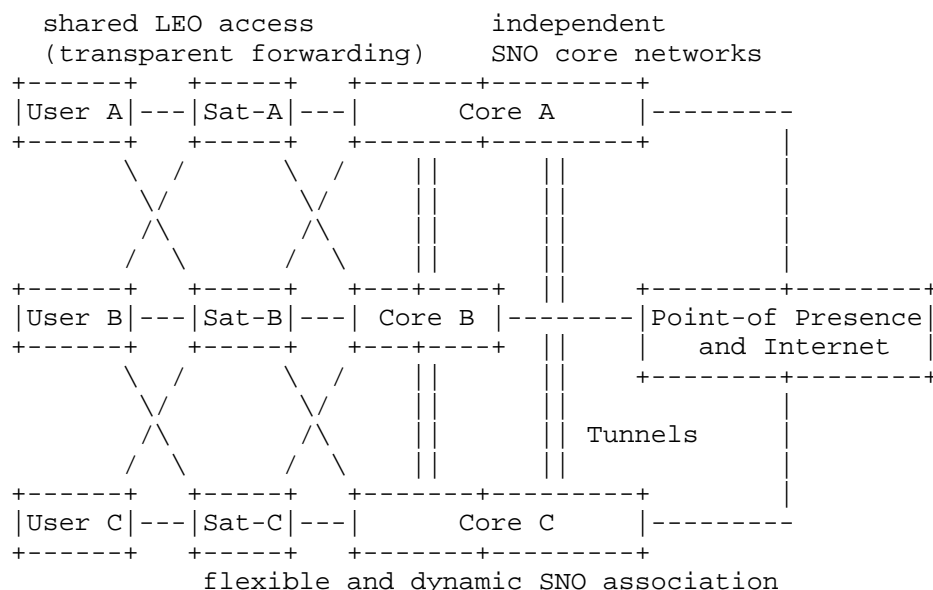


Figure 3: The proposed mobility management framework overview.

This document outlines an architectural approach for constellation cooperation that aims to limit the operational impact of recurrent inter-SNO handovers and to support stable inter-constellation service integration. At a conceptual level, the architecture is composed of two principal elements: a common LEO-based access layer shared among participating operators, and a set of autonomous SNO core networks. The overall structure of this arrangement is depicted in Figure 3.

4.1. Shared LEO access network

Within this architecture, participating SNOs make their satellite resources and associated spectrum bands mutually available, potentially leveraging techniques such as those described in [spectrum_sharing_leo]. Collectively, these satellites form a unified LEO access layer that is shared across operators. Satellites in this shared layer operate in a stateless forwarding mode, allowing a single satellite to simultaneously provide access services to users subscribed to different SNOs. Because a variety of established mechanisms already exist for assigning satellite access points to terrestrial users, the framework does not mandate a specific selection policy. Instead, it permits collaborating SNOs to jointly determine a common satellite allocation strategy, which may consider factors such as signal quality, instantaneous traffic load, or other operational metrics. The resulting allocation decisions are then enforced at the satellite terminal level.

4.2. Independent SNO cores

Within each operator domain, the SNO core network continues to handle functions that are specific to that operator, including user authentication, address assignment, and other control-plane responsibilities. Accordingly, the architecture preserves the administrative independence of participating SNO core networks. To support inter-operator coordination, secure connectivity is established between the core networks of different SNOs over the terrestrial Internet. These protected tunnels provide a communication channel for exchanging signaling information and aligning operational decisions across domains. Each SNO maintains a set of geographically distributed ground stations that serve as gateways between its core infrastructure and the satellite layer. Building on this foundation, the framework enables dynamic satellite-to-operator mapping. At any given time, a satellite within the shared access layer may attach to the ground station and core network of a selected SNO, allowing association relationships to evolve in response to operational conditions.

4.3. Adaptive space-ground associations

In contrast to conventional collaboration models, such as the one shown in Figure 2, where an operator's LEO access segment remains tightly bound to its own core infrastructure, the proposed architecture separates these two layers. By removing the fixed coupling between access satellites and specific SNO cores, the framework allows connection relationships to be adjusted dynamically, with the objective of limiting unnecessary inter-SNO handovers experienced by end users. As depicted in Figure 3, users belonging to different SNOs attach to a common LEO access layer that is shared among participating operators. Because satellites move rapidly along their orbital paths, user terminals inevitably handover between different access satellites over time. To preserve continuity at the operator level, the framework continuously adapts the mapping between access satellites and SNO core networks. Through this adaptive association, user sessions are maintained with their original SNO whenever feasible, thereby reducing cross-operator handovers. In situations where a handover between SNO cores cannot be avoided, the architecture relies on inter-core tunneling mechanisms to forward user traffic back to the original SNO core network, mitigating the impact of the handover on ongoing sessions.

4.4. Key techniques behind the proposed framework

To mitigate the service degradation caused by inter-SNO handovers, the architecture incorporates two complementary mechanisms. The first mechanism regulates SNO association by dynamically assigning stateless access satellites to suitable core networks. Through adaptive satellite-to-core mapping, the frequency of cross-operator handovers experienced by terrestrial users can be effectively constrained. The second mechanism focuses on accelerating the handover procedure itself. By taking satellite trajectory information into account, a fast handover scheme is applied to shorten the duration of network-layer interruption during inter-SNO handovers.

5. Reducing Disruption Time During Handovers

By taking advantage of the deterministic nature of satellite orbital motion, the architecture computes in advance the mapping between individual satellites and participating SNO core networks. This advance planning reduces the likelihood of unnecessary cross-operator handovers. Building on these pre-established association results, the framework further implements an expedited inter-SNO handover procedure. Because the target operator relationship is known ahead of time, network-layer preparation can be initiated prior to the actual handover, thereby shortening service interruption during inter-SNO handovers.

5.1. Baseline inter-SNO handover optimizer

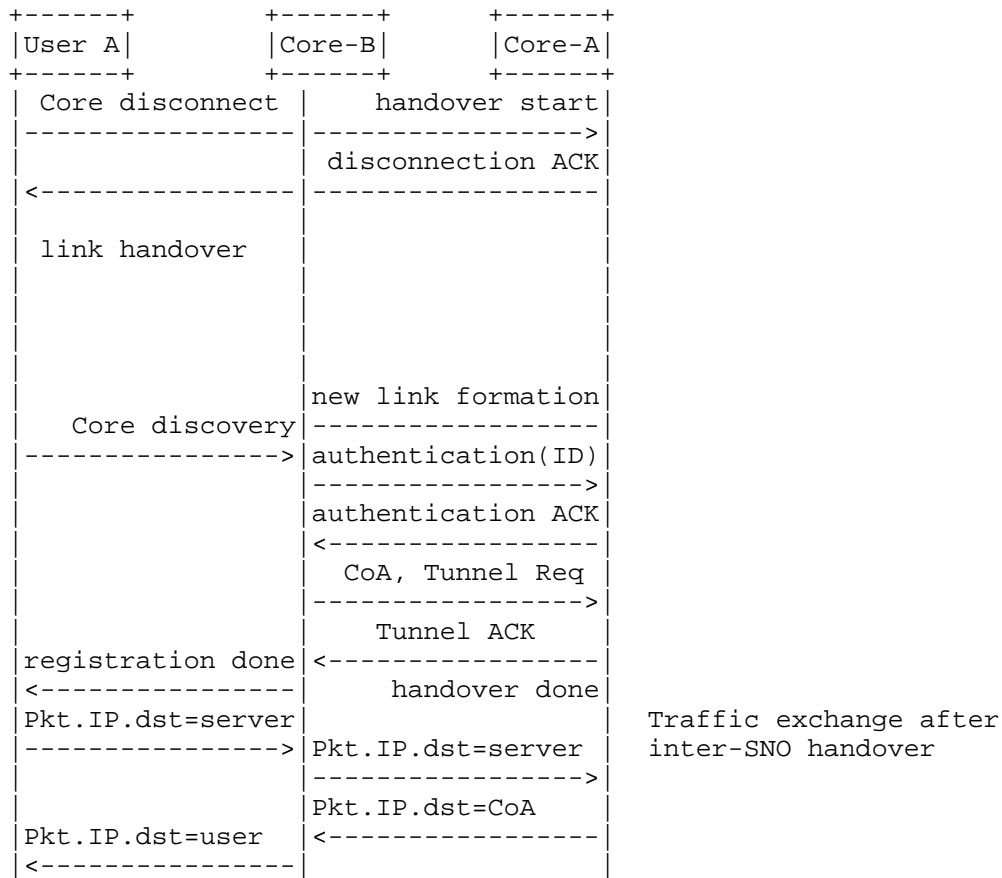


Figure 4: Baseline inter-SNO handover.

This section first describes a baseline handover optimization procedure that incorporates the Mobile IP mechanism defined in [RFC6275] into the LEO environment, as shown in Figure 4. The objective of this baseline design is to preserve the user's IP address across an inter-SNO handover. Consider a user initially served by SNO-A. The user accesses the Internet through the shared LEO access layer and is anchored at the core network of SNO-A. When satellite mobility causes the user to attach to a new access satellite associated with SNO-B, a cross-operator handover is triggered. Because SNO-A and SNO-B maintain independent IP address pools, a direct handover would normally result in address reassignment. To prevent this outcome, the baseline procedure proceeds as follows. Upon detaching from core A, the user initiates an unregistration process by sending a disconnect notification to core A. The user then establishes connectivity with the incoming

satellite, while link-layer mechanisms handle the physical disconnection and subsequent link setup. Once the new link becomes operational, the user broadcasts a core discovery message containing its unique identifier. After receiving this identifier, core B determines that the user is originally associated with SNO-A. Core B forwards the identifier to core A over the terrestrial network, allowing core A to authenticate the user. Following successful authentication, core B allocates a care-of address (CoA) for the user and sets up a secure tunnel between the two core networks. Core B then informs the user that registration has been completed. At this stage, the user regains Internet connectivity without changing its original IP address. During subsequent data exchange, packets transmitted from the user toward an external server are first routed through core B, then tunneled to core A, and finally delivered to the destination. In the reverse direction, because the destination IP corresponds to SNO-A, incoming packets are directed to core A. Core A encapsulates each packet with a CoA associated with SNO-B and forwards them through the inter-core tunnel. Core B decapsulates the packets and delivers them to the user.

5.2. Optimizing handovers by proactive authentication and tunneling

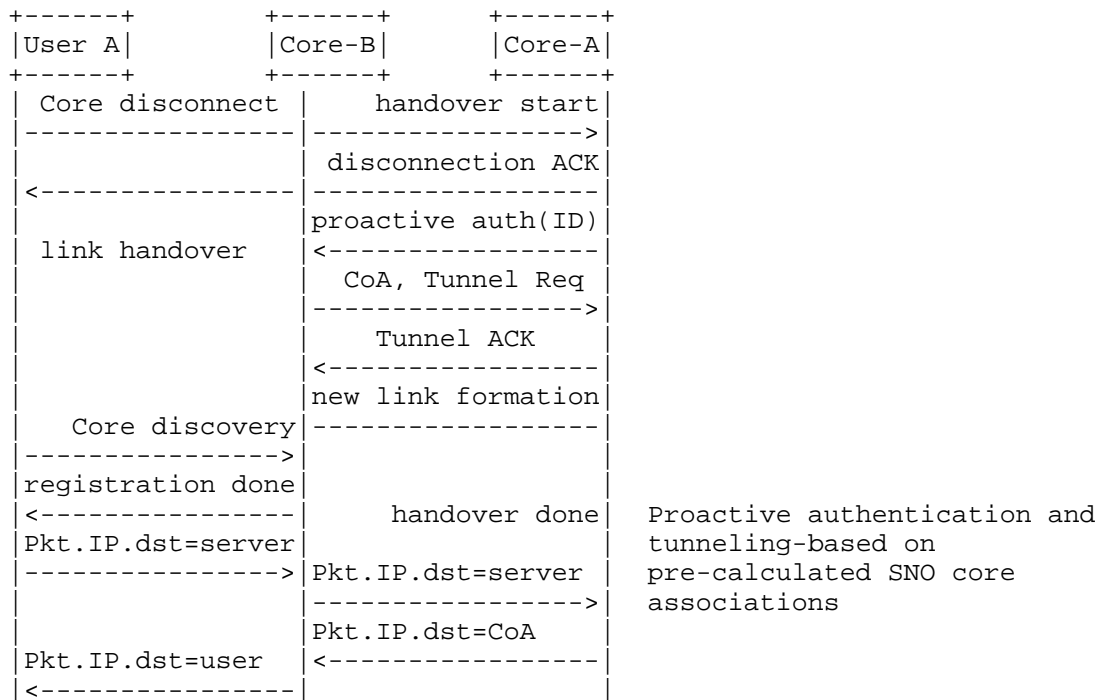


Figure 5: Optimized inter-SNO handover.

The baseline network-layer recovery procedure may introduce several seconds of service interruption. To further shorten this interruption interval, the architecture refines the baseline process by utilizing advance knowledge of satellite-to-SNO associations. Because these associations are computed ahead of time, the corresponding inter-SNO coordination steps can be partially prepared before the actual handover occurs. An overview of the enhanced handover procedure is provided in Figure 5. With the association results already available, core A can determine in advance the target core network to which the user will attach after the handover. Once the user detaches from core A, core A immediately transmits the user identifier to core B to initiate local authentication procedures. In parallel with the link-layer switching process, core B allocates a care-of address (CoA) and establishes the required inter-core tunnel. After the physical link to the new satellite becomes operational, the user issues a core discovery message. Core B validates the identifier and finalizes the registration with minimal additional delay. Because inter-core signaling and tunnel setup are executed concurrently with the underlying link-layer handover, the overall service disruption experienced during the inter-SNO handover can be significantly reduced.

6. Conclusion

This document presents an architectural model for cooperation among multiple SNOs with the objective of strengthening service continuity and operational stability in LEO environments. The design addresses inter-SNO mobility from two complementary perspectives: limiting the occurrence of cross-operator handovers and accelerating the handover process when such handovers are unavoidable. Through this combined strategy, the disruptive effects commonly associated with frequent inter-SNO handovers can be substantially alleviated. By outlining mechanisms that coordinate shared access infrastructure with independent core networks, the architecture establishes a technical basis for interoperable inter-SNO operation within collaborative LEO constellations. The concepts and procedures described herein are intended to contribute to ongoing and future discussions on standardizing inter-SNO inter-networking solutions.

7. IANA Considerations

This document includes no request to IANA.

8. Security Considerations

This document does not represent a change to any aspect of the TCP/IP protocol suite and therefore does not directly impact Internet security.

9. References

9.1. Normative References

- [RFC6275] Perkins, C., Ed., Johnson, D., and J. Arkko, "Mobility Support in IPv6", RFC 6275, DOI 10.17487/RFC6275, July 2011, <<https://www.rfc-editor.org/info/rfc6275>>.
- [RFC5944] Perkins, C., Ed., "IP Mobility Support for IPv4, Revised", RFC 5944, DOI 10.17487/RFC5944, November 2010, <<https://www.rfc-editor.org/info/rfc5944>>.

9.2. Informative References

- [satellite_constellation_launch]
Cornara, S., Beech, T., Bello-Mora, M., and A. M. D. Aragon, "Satellite constellation launch, deployment, replacement and end-of-life strategies", 1999.
- [decentralized_satellite_networks]
Oh, S. and D. Vasisht, "A call for decentralized satellite networks", 2024.
- [democratizing_leo_satellite_network_measurement]
Izhikevich, L., Tran, M., Izhikevich, K., Akiwate, G., and Z. Durumeric, "Democratizing LEO satellite network measurement", 2024.
- [comarch-orange-T-Mobile]
Comarch, "Comarch Fault Management Supports Orange and T-Mobile's Infrastructure Sharing Initiative in Poland", URL <https://www.comarch.com/telecommunications/customers/networks/>, 2026.
- [analysysmason-network-sharing]
Analysys Mason, "Operators View Cost Saving and Coverage as the Two Main Drivers for Network Sharing", URL <https://www.analysysmason.com/research/content/articles/cost-network-sharing-rdns0/>, 2026.
- [satellitetoday-iridium-oneweb]
Satellite Today, "Constellations Combined: Iridium and OneWeb Join Forces on New LEO Service", URL <https://www.satellitetoday.com/mobile-connectivity/2019/09/17/constellations-combined-iridium-and-oneweb-join-forces-onnew-leo-service/>, 2026.

`[practical_traffic_management_lte_wifi]`

Mahindra, R., Viswanathan, H., Sundaresan, K., Arslan, M. Y, and S. Rangarajan, "A Practical Traffic Management System for Integrated LTE-WiFi Networks", 2014.

`[skycastle_leo_mobility]`

Li, J., Li, H., Lai, Z., Wu, Q., Liu, W., Wang, X., Li, Y., Liu, J., and Q. Zhang, "Skycastle: Taming LEO Mobility to Facilitate Seamless and Low-Latency Satellite Internet Services", 2024.

`[spectrum_sharing_leo]`

Yun, J., An, T., Jo, H., Ku, B., Oh, D., and C. Joo, "Downlink Spectrum Sharing of Heterogeneous Communication Systems in LEO Satellite Networks", 2022.

Acknowledgements

Contributors

Thanks to all of the contributors.

Authors' Addresses

Zeqi Lai
Tsinghua University
30 ShuangQing Ave
Beijing
100089
China
Email: zeqilai@tsinghua.edu.cn

Yunan Hou
Tsinghua University
30 ShuangQing Ave
Beijing
100089
China
Email: houyn24@mails.tsinghua.edu.cn

Qian Wu
Tsinghua University
30 ShuangQing Ave
Beijing
100089
China

Email: wuqian@cernet.edu.cn

Hewu Li
Tsinghua University
30 ShuangQing Ave
Beijing
100089
China
Email: lihewu@cernet.edu.cn