PoC Number (assigned by ETSI): #22

PoC Project Name: NTN, 5G SRv6 integration for TSN by Artificial Intelligence

PoC Project Host: CNR ISTI

PoC Team Members

	Organization name	ISG ENI participant (yes/no)	Contact (Email)	PoC Point of Contact (see note 1)	Role (see note 2)	PoC Components
1	CNR ISTI	Yes	Pietro Cassarà pietro.cassara@isti.cnr.it		Academic	- User Stories / Use Cases definition - PoC development - PoC documentation - PoC demos
2	CNIT	Yes	Fabrizio Granelli fabrizio.granelli@unitn.it		Academic	-Help with concept proof
3	Huawei	Yes	Aldo Artigiani aldo.artigiani@huawei.com		Vendor	-Help with architecture design, protocol design, intelligent solution, and platform verification
4	University of Luxenburg	Yes	Jorge Querol Borras, Mohammed Al-Ansi		Academic	Help with architecture design, protocol design,
5	CTTC	Yes	Pol Henarejos, Xavier Artiga		Academic	Help with architecture design, protocol design,

NOTE 1: Identify the PoC Point of Contact with an X.
NOTE 2: The Role will be network operator/service provider, infrastructure provider, application provider or other as given in the Definitions of ETSI Classes of membership.

REFERENCE SCENARIO

Figure 1 shows the reference topology and the protocol stack investigated in this document for the analysis of the performance of the SRv6 protocol for TSN scenarios. Our study aims to analyze how the SRv6 protocol can be configured to manage routing in an NTN serving as IP-RAN. Specifically, we analyze the use of the SBFD mechanism to address sudden changes in the NTN topology and evaluate its impact on QoS in terms of latency added to the typical latency of the NTN link.

The upper part of the figure illustrates a use case where a user equipment (UE) connects to the network through an NT-RAN. Here, the satellite-based access adheres to the 5G NR standards defined by ETSI [1]–[3], with satellites communicating via optical Inter-Satellite Links (ISLs). The IPv6 protocol underpins the IP routing between network nodes. A non-terrestrial gateway also connects satellites and data networks through an optical or radio-based feeder link.

The lower part of Figure 1 presents the IPv6-based Plane, showcasing the protocol stack architecture employed for Layer 3 routing in NT nodes using the SRv6 protocol. Traffic from the Uu interface of the ground cell (C#1) to the 5G Core's User Plane Function (UPF) via the NT-Gateway is initially routed along a predefined path, labeled Path #1 {SAT #1, SAT #2, SAT #3}, utilizing a GTP-U tunnel. If link failure or topology changes disrupt this path, the SBFD mechanism facilitates traffic rerouting to an alternative path, Path #2 {SAT #1, SAT #3}. This mechanism is crucial for detecting link failures within segmented paths. By leveraging SBFD, the analysis focuses on its efficiency in handling topology changes and ensuring effective rerouting within the NT-RAN.

To evaluate the performance of the topology introduced above, we consider that the connection between the UE and the access satellite is maintained using the handover protocol outlined in ETSI's 5G technical specifications [1], [4]. Furthermore, the analysis focuses on the performance of the SBFD mechanism in managing traffic paths within the satellite constellation. This includes inter-satellite and intra-satellite links and the feeder link connecting the final satellite to the NT gateway.

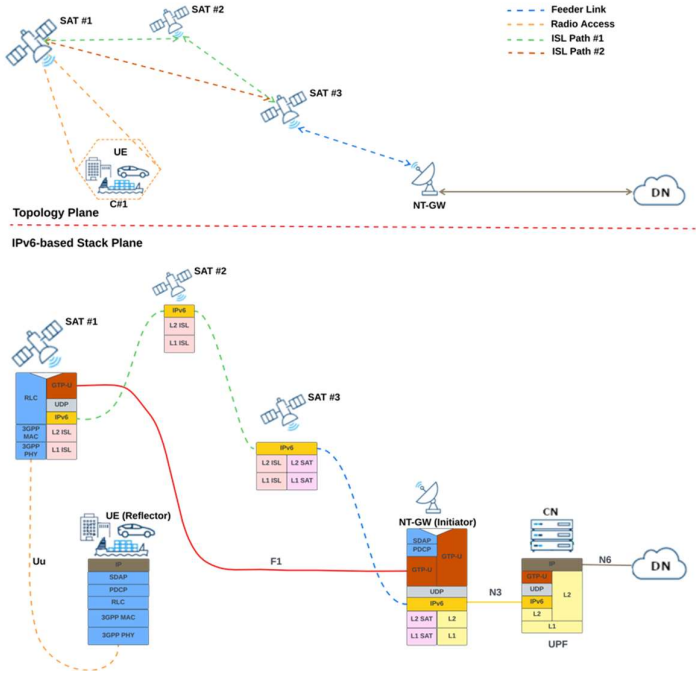

Figure 1: 5G-NTN Integrated Reference Scenario

Handover on 5G-NTN integrated infrastructures

As the benefits of Conditional Handover (CHO) gained widespread recognition within the 3GPP community, with its value analyzed and validated in various research and technological white papers, the approach has been extended to other applications of cellular communication. This chapter explores the potential adoption of CHO in vertical domains, specifically for use cases beyond Mobile Broadband (MBB).

Leveraging cellular technology via satellite networks presents an opportunity to deliver global coverage, even to regions previously lacking mobile communication infrastructure. However, this advancement introduces new challenges, particularly in addressing NTN-specific mobility issues. This is especially relevant for Low-Earth Orbit (LEO) satellites, which typically operate at altitudes of 600 to 1200 kilometers. The characteristics of these satellites, such as significant propagation delays and their continuous movement, complicate timely handover (HO) execution. Unlike Terrestrial Networks (TNs), where base stations are stationary, both User Equipment (UEs) and gNBs in NTNs are in motion.

In conventional CHO, handover execution criteria are based on the radio conditions (e.g., signal strength or quality) of the source and target cells. In NTNs, additional criteria are required to account for satellite movement and temporary coverage. The traditional reliance on signal strength and quality may be insufficient, as the differences between the cell center and edge measurements are often too minor to predict coverage loss effectively. Therefore, two new handover execution criteria have been introduced for NTNs [6]:

1. **Timing-Based Handover Execution**
   The predictable movement of satellites allows networks to determine the availability of specific cells over a given area. Using this information, the UE can be configured with CHO conditions specifying a time window [t1, t2] for accessing a target cell. However, relying solely on timing without considering radio conditions could negatively impact handover reliability. Therefore, timing-based CHO is typically implemented alongside radio-based measurement events. The UE first evaluates the radio conditions after t1 and executes the CHO if these criteria are met before t2.
2. **Location-Based Handover Execution**
   Similarly, CHO can be triggered based on the UE's proximity to a reference point associated with each NTN cell. For single-beam satellite coverage, the reference point may correspond to the beam center. Handover is initiated when the UE's distance to the current cell's reference point exceeds a threshold and its distance to the target cell falls below another threshold. As with timing-based execution, location-based CHO works best when combined with radio measurement criteria.

In standard CHO implementations, handovers are prepared for immediate transitions (i.e., single-cell changes). However, in NTNs, it may be possible to anticipate a sequence of future handovers due to the predictable orbital paths of satellites and the relatively static positions of UEs. This enables a chain of CHOs, where multiple handover configurations are provided to the UE in a single procedure. Figure 2 illustrates this concept, showing how the UE can be pre-configured with CHO settings for multiple upcoming cells. This approach minimizes signaling overhead, reduces unnecessary resource pre-allocation, and enhances mobility performance by lowering the likelihood of ping-pong handovers and mobility failures.

Integrated Access Backhaul (IAB) is a feature of New Radio (NR) technology that allows wireless communication to be relayed through the Radio Access Network (RAN). This architecture includes the **IAB-donor**, which serves as both a backhaul traffic node and a gNB, and **IAB-nodes**, which function as Distributed Units (gNB-DUs) while also supporting UE functionalities (referred to as IAB-MT in this context). The detailed architecture and procedures for IAB are documented in [7].

In 3GPP Release 17, enhancements to the IAB framework include CHO for improved reliability. CHO is applied to IAB-MTs, following a process similar to the UE handover procedure defined in Release 16. However, CHO is also being considered for IAB-node migration, which involves transitioning an IAB-node to a different parent node under the same IAB-donor Centralized Unit (CU). Unlike IAB-MTs, which are mobile, IAB-nodes are stationary, so their migration conditions differ from those of traditional handovers.

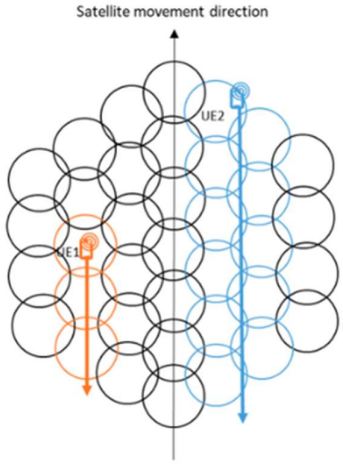

Figure 2: Example of CHO Chain

MULTIPATH MANAGEMENT

The Multi-Access Packet Data Unit (MA-PDU) Connectivity Service, established by the Technical Specification Group Services and System Aspects in Release 18 of the 5G System (5GS) standard under the Third Generation Partnership Project (3GPP), enables the simultaneous use of 3GPP and non-3GPP access networks for exchanging Packet Data Units (PDUs) between User Equipment (UE) and the Data Network (DN). This service is integral to the Access Traffic Steering, Switching, and Splitting (ATSSS) functionality.

A 3GPP access network refers to networks defined by 3GPP, such as 5G or 6G cellular networks, while non-3GPP access networks include technologies like Wi-Fi or Non-Terrestrial Networks (NTN). The MA-PDU service can also operate when a 3GPP access network connects to the 4G Evolved Packet Core (EPC) and a non-3GPP access network connects to the 5G Core (5GC). However, it does not support configurations where a 3GPP access network links to the 5GC and a non-3GPP network links to the EPC.

In the MA-PDU service framework, PDUs are transmitted via two separate N3/N9 tunnels that connect the PDU session anchor (PSA) in the User Plane Function (UPF) to the radio access or other access network (RAN/AN). This connectivity supports traffic types including IPv4, IPv6, dual-stack IPv4/IPv6, and Ethernet, as illustrated in Figure 3 of the standard.

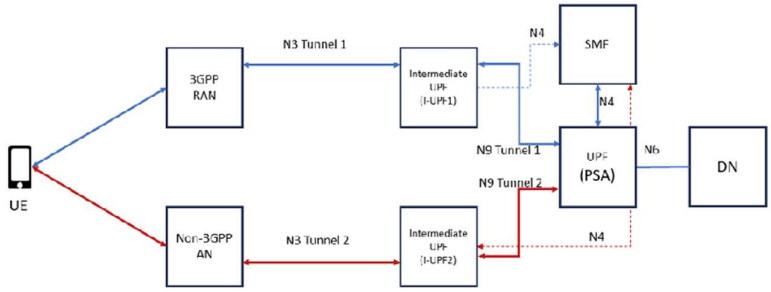

Figure 3: Multipath Architecture

TESTBED SETUP

The testbed was developed using the VIAVI MTS-5800 testing tool, three Huawei NetEngine 8000 M1C routers, two Huawei ATN905S routers, and MATLAB's Satellite Toolbox. Satellite dynamics were simulated using the MATLAB Satellite Toolbox in conjunction with the Starlink constellation's TLE dataset [5]. From this simulation, latency and link availability distributions were derived for the network topology. These distributions were then used in SRv6 traffic engineering to emulate the latencies and availabilities of the links connecting the routers.

Leveraging SRv6's capability to dynamically update the SIDs that define a path, backup link policies were configured to ensure seamless rerouting when a primary link failed or its quality fell below a specified threshold. The rerouting paths were informed by topology data generated from the Starlink constellation simulation and were selected using a custom routing application based on predefined policies. It is worth noting that routing policy design lies outside the scope of this phase of the study. Huawei's NetEngine 8000 M1C and ATN905S routers were selected for their compatibility with LEO satellite deployments, supporting IPv6 and SRv6 functionality while meeting bandwidth requirements.

Figure 4 illustrates the testbed setup analyzed in this study. The testbed configuration includes three Huawei NetEngine 8000 M1C routers, emulating SAT #1, SAT #2, and SAT #3, connected in a meshed topology to mimic the redundancy of the Starlink satellite network. Two Huawei ATN905S routers represent the User Equipment (UE) and the Non-Terrestrial Gateway (NT-GW). This meshed topology ensures fault tolerance by providing redundant paths, with SRv6 enabling the exploitation of these redundant paths. The SBFD mechanism was implemented to monitor link conditions and enable swift fault detection.

UE traffic flow originating from ground cell C#1 was emulated using the VIAVI MTS-5800 as highlighted by the red dotted square in Figure 4. The VIAVI MTS-5800 was also employed to measure real-time network latency and link availability, enabling the analysis of traffic restoration speed following path switching. During parameter tuning, particular attention was given to calibrating the SBFD detection timer. This timer controls the responsiveness of fault detection and was optimized to achieve maximum sensitivity while accounting for typical LEO satellite communication latencies, which generally range from a few milliseconds to around ten milliseconds.

To minimize false alarms caused by environmental delays, best practices were applied to set the SBFD detection timer at twice the average latency. This adjustment balances prompt fault detection with the need to avoid false positives, ensuring reliable and rapid responsiveness in detecting and addressing link failures.

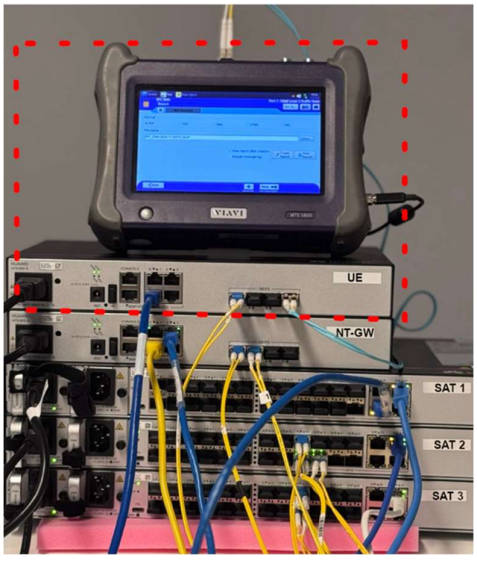

Figure 4:Testbed Configuration

PERFORMANCE ANALYSIS

The numerical analysis involves traffic measurements on Path #1 and Path #2 during switching procedures triggered by the two types of generated events:

- **E1**: A link in Path #1 gradually degrades due to satellite movement;
- **E2**: A link abruptly fails because of unexpected issues such as radio malfunctions or obstructions, resulting in a sudden disconnection of Path #1.

For **E1**, a soft disconnection is simulated through Traffic Engineering (TE) policies using pre-generated data from MATLAB. In **E2**, a hard disconnection mimics hardware failure. Traffic generated by the VIAVI MTS-5800 was transmitted at a bit rate of 500 Mbps, with a standard MPDU size of 1500 bytes, totaling 65 million frames. The SBFD timer for link fault detection was set to 20 ms. The Starlink constellation topology was modeled using TLE data, calculating all possible three-hop paths between endpoints over a 2-hour window. This duration covers one complete revolution of access satellite SAT #1 and accounts for all handovers over cell C#1. SAT #1 provides coverage for C#1 with elevation angles between 35° and 45°, and an aperture of 110°, in line with Starlink's Dishy specifications.

The latency distributions for inter- and intra-satellite links are presented in Figure 5b, while feeder link latencies are shown in Figure 5a. TE policies were derived from these distributions to adjust link characteristics dynamically. Figure 5b identifies two latency modes: an average of 7 ms for inter-satellite links within the same orbital plane and 12 ms for links across different planes.

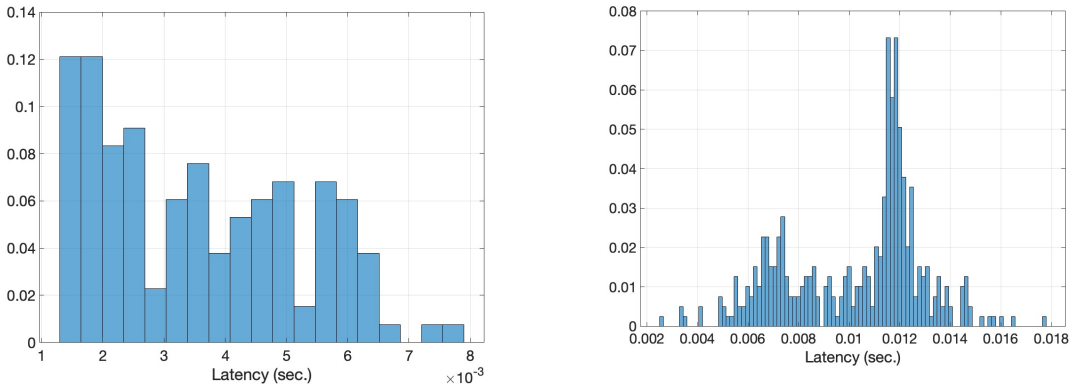

(a) (b)

Figure 5: Latency Probability Distribution

Event E1 Results

Table I summarizes the performance metrics during switching triggered by E1. Tests were conducted on 2024/11/05 at 11:00 UTC for 170 minutes, generating 50 switching events. The average throughput remained close to the maximum, with end-to-end latency proportional to the number of satellites traversed.

- **Path Parameters**: The throughput was measured as a percentage of the maximum link rate, while latency was recorded in milliseconds.

- **SBFD Analysis**: The average latency introduced by the SBFD mechanism was 222 µs, with a maximum jitter of 3.8 µs, which is three orders of magnitude smaller than satellite link latency.
- **OoS and Frame Loss**: Some Out-of-Sequence (OoS) and frame loss events occurred, but their probabilities were negligible. OoS events were attributed to SRv6 buffering, which aims to minimize frame loss during switching.

Table I: E1 Results

	MPDU (Bytes) 1518		**Tx frames** $65 \cdot 10^6$		
Path Parameters		**SBFD Analysis**			
Thrpt. (%)	**Lat.** (ms)	**Lat.** (μs)	**max Jit.** (μs)	**OoS** (%)	**frame Loss** (%)
98.9	37.2	222	3.8	$6.135 \cdot 10^{-6}$	$4.71 \cdot 10^{-4}$

Event E2 Results

Table II details the results for E2, conducted on 2024/11/05 at 15:00 UTC for the same duration as E1. Despite the abrupt nature of E2, SRv6 switching maintained near-maximum throughput and latency consistent with the number of satellite hops.

- **Path Parameters**: Throughput and latency remained comparable to E1, but SBFD latency increased by 18%, while jitter remained nearly unchanged and negligible relative to satellite link latency. The increase in SBFD latency reflects SRv6's buffering, designed to reduce frame loss.
- **OoS and Frame Loss**: Frame loss increased by one order of magnitude but remained minimal. No OoS events were observed, likely due to buffer depletion caused by the abrupt link failure, leading to a temporary loss of all buffered data.

In both cases, SRv6 demonstrated robust performance in maintaining throughput and minimizing disruptions during switching, even under abrupt failure conditions.

Table II: E2 Results

	MPDU (Bytes) 1518		**Tx frames** $65 \cdot 10^6$		
Path Parameters		**SBFD Analysis**			
Thrpt. (%)	**Lat.** (ms)	**Lat.** (μs)	**max Jit.** (μs)	**OoS** (%)	**frame Loss** (%)
98.6	35.8	262	3.6	0	$1.79 \cdot 10^{-3}$

REFERENCE

[1] G. RAN. (2023) Technical specification 38.300 (ver.17.5.0) nr; nr and ng-ran overall description; stage-2. [On-line]. Available: https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails\\.aspx?specificationId=3191

[2] X. Lin, S. Rommer, S. Euler, E. A. Yavuz, and R. S. Karlsson, “5g from space: An overview of 3gpp non-terrestrial networks,” IEEE Communications Standards Magazine, vol. 5, no. 4, pp. 147–153, 2021.

[3] M. Majamaa, “Toward multi-connectivity in beyond 5g non-terrestrial networks: Challenges and possible solutions,” IEEE Communications Magazine, vol. 62, no. 11, p. 1–7, 2024.

[4] G. SA. (2024) Technical specification 23.502 (ver.18.7.0) procedures for the 5g system (5gs). [On-line]. Available: https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails\\.aspx?specificationId=3145

[5] celestrak.org. Two-line element (tle) data constalletion starlink. [Online].Available: http://celestrak.org/NORAD/elements/supplemental/

[6] 3GPP Technical Specification 38.300 “NR and NG-RAN Overall description; Stage-2”, section 9.2.3.2, v16.8.0, December 2021

[7] 3GPP Technical Report 38.874 “Study on integrated access and backhaul”, v16.0.0, January 2019