Keeping Up with the Trends: Identifying Shifts in LEO Satellite Usage Patterns

Stepan Kravtsov
skravtsov3@gatech.edu
Georgia Institute of Technology
Atlanta, USA

Sherif Abdelrazek
sabdelrazek3@gatech.edu
Georgia Institute of Technology
Atlanta, USA

Vaibhav Bhosale
vbhosale6@gatech.edu
Georgia Institute of Technology
Atlanta, USA

Ahmed Saeed
asaeed@cc.gatech.edu
Georgia Institute of Technology
Atlanta, USA

Abstract

Low Earth Orbit (LEO) satellite networks, such as Starlink, are increasingly used for Internet access, emergency response, and backup connectivity. However, tracking when and where Starlink usage changes remains difficult because direct sources of evidence, such as operator telemetry and subscriber records, are not publicly available. We present a cross-dataset tool for detecting high-confidence shifts in Starlink usage from public measurements. The tool combines three complementary vantage points: APNIC estimates of visible users, M-Lab speed-test activity, and Cloudflare Radar traffic trends. It first identifies changes in user trends independently in each time series, then reports candidate events when multiple vantage points observe same-direction changes within a narrow temporal window. We evaluate the approach through country-level case studies of Jamaica, Ukraine, and Germany. In Jamaica, the tool identifies a sharp increase in Starlink usage around the Hurricane Melissa response. In Ukraine, it detects usage shifts consistent with reported changes in Starlink deployment and availability. In Germany, it avoids treating a prominent APNIC-only rise-and-fall pattern as a supported event. These results show that cross-source public measurements can provide the community with a practical and high-confidence way to track event-driven shifts in how LEO networks are used and relied upon over time.

CCS Concepts

• **Networks** → **Network measurement**; **Network dynamics**; *Wireless access networks.*

Keywords

LEO satellite networks, Starlink, Cross-source correlation

ACM Reference Format:
Stepan Kravtsov, Sherif Abdelrazek, Vaibhav Bhosale, and Ahmed Saeed. 2026. Keeping Up with the Trends: Identifying Shifts in LEO Satellite Usage Patterns. 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/10.1145/3789240.3827585

LEO-NET '26, Denver, CO, USA

ACM ISBN 979-8-4007-2467-1/2026/08
https://doi.org/10.1145/3789240.3827585

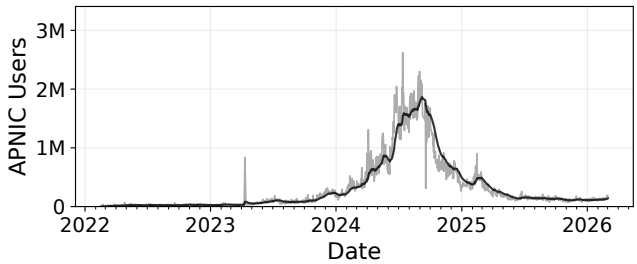

Figure 1: APNIC data for Starlink in Germany, showing a pronounced rise and fall in network usage that is not corroborated by other public signals.

1 Introduction

Low Earth Orbit (LEO) satellite networks are rapidly becoming an important part of the Internet access ecosystem. Systems such as Starlink are no longer niche deployments. They are increasingly used to extend connectivity to remote areas, provide access in areas with weak terrestrial infrastructure, and enhance resilience when conventional networks are disrupted. Recent measurement studies have shown both the growing operational relevance of Starlink and the need for better public visibility into how LEO networks behave in practice, including during failures, mobility scenarios, and emergency or fail-over settings [2, 14, 18, 20].

With a growing reliance on LEO networks in emergency scenarios and as a failover, understanding and following trends in LEO network adoption over time can be valuable for networking practitioners and operators, as well as policy makers [4, 6, 8]. However, tracking such usage shifts using public datasets remains a challenge. For example, consider the time series data shown in Figure 1. The data represents an estimate of the number of Starlink users in Germany, as estimated by the APNIC eyeballs dataset. The graph seems to indicate a sharp increase in users in the first half of 2024, followed by a sharp decline in the latter half of the same year. The trend exhibited by the data is very prominent and stable over a very long duration. However, as we will show later in this paper, no other data source that approximates or predicts network usage shows that trend. Thus, relying on a single source of information on network behavior, especially for satellite networks where changes in user adoption are expected, can be incomplete or even misleading.

To overcome this limitation, we propose tracking Starlink usage trends over long timescales using multiple public datasets. To improve the robustness of our tracker, we identify shifts only when there is consensus between multiple datasets on the timing and trajectory of the network usage trend. Specifically, we use three public datasets that capture Starlink usage from different perspectives: the APNIC "eyeballs" dataset, which estimates the number of Internet users associated with an Autonomous System (AS) based on large-scale online ad impressions; the daily number of unique M-Lab speed tests conducted by Starlink users in a given country; and the normalized number of HTTP requests observed by Cloudflare and reported through the Cloudflare Radar API.

Our goal is to detect two types of shifts: upward shifts, indicating a significant increase in Starlink usage, and downward shifts, indicating a significant decrease. To do so, we first identify candidate trend changes independently in each public time series using smoothed piecewise-linear segmentation, followed by post-processing that removes weak or short-lived changes. We then retain an event only when multiple datasets exhibit a same-direction change within a narrow temporal window. This consensus rule is intentionally conservative: it treats agreement across independent vantage points as stronger evidence than any single visually prominent trend.

Through three case studies, we show that this approach captures both positive and negative evidence. In Jamaica, the method surfaces a sharp increase in Starlink usage around a hurricane-related communications disruption. In Ukraine, it identifies large usage shifts consistent with reported changes in Starlink deployment and availability. In Germany, by contrast, a striking rise-and-fall pattern in APNIC is not supported by the other public signals, illustrating how cross-source agreement can prevent over-interpreting single-source artifacts. We next motivate why such LEO usage shifts matter and why public measurements require cautious, cross-source interpretation.

2 Background and Motivation

Why Study Changes in LEO Network Usage? Understanding how LEO network usage changes over time can reveal the operational role that LEO networks play during both normal conditions and disruptions. Under normal conditions, LEO networks such as Starlink provide connectivity not only in areas where terrestrial infrastructure is weak or unavailable (e.g., rural areas), but also in settings where users need flexible, quickly deployable, or mobility-supporting connectivity. More broadly, they can serve as an alternative or complementary access path when terrestrial options are limited by coverage, capacity, deployment cost, or temporary operational constraints. They can also act as backup connectivity during outages, support temporary fail-over when terrestrial paths degrade, and provide connectivity in mobility or emergency-response settings. Recent work has increasingly treated LEO networks not only as an access technology, but also as part of the Internet's resilience story, including their potential role in emergency or national failover scenarios [2, 20].

Because of this role, sudden increases or decreases in LEO network usage should be treated as meaningful signals of shifts in the networking infrastructure. A sharp increase may indicate that users are shifting onto satellite connectivity in response to a terrestrial failure or regional disruption, while a decrease may reflect recovery of terrestrial infrastructure, shifting demand, or limitations in the LEO system itself. Tracking such changes can therefore help researchers and operators reason about resilience, emergency communications, dependence on satellite infrastructure, and the extent to which LEO networks operate as a complement to or a substitute for terrestrial networks [10].

Public datasets are attractive but challenging. Public datasets offer a reproducible and accessible way to study network behavior without requiring privileged access to operator telemetry [7, 10–12]. Unlike proprietary internal data, public datasets can be queried by outside researchers, revisited over time, and compared across countries and providers. Analysis of such datasets have made it possible to glean significant information about changes in the Internet (e.g., outage detection [3]). However, public datasets often provide information from a single perspective or vantage point which might be affected by artifacts unique to its own collection methodology. As a result, an event that is observed in one single dataset may reflect a real-world event that can only be observed from that vantage point. However, it can also reflect a measurement artifact or biases unique to that dataset. Distinguishing between the two scenarios is very challenging and, in some cases, impossible without access to private network logs. This challenge motivates our approach of seeking consensus across multiple public datasets to identify trends in Starlink usage.

3 Datasets

APNIC. Our first signal comes from the APNIC dataset, which provides daily Starlink user estimates by ISP and ASN. APNIC derives these estimates from its ad-based measurement platform and combines the resulting observations with population and Internet-use estimates to infer how many users are visible on a given network over time [10]. This signal is relevant because it approximates the number of users who are active or visible in a given economy and network over time. At the same time, it is not a ground-truth subscriber count. The inference process depends on assumptions about how well the underlying measurement sample represents the broader user population and how users are mapped to networks, so the signal reflects both actual usage dynamics and dataset-specific uncertainty [10]. APNIC itself notes that these estimates carry inherent uncertainty and that smaller networks are less reliably represented [10].

M-Lab. Our second signal comes from M-Lab, from which we derive the daily number of Starlink-associated tests for a given country and ASN. This signal is relevant because public measurement activity can reflect connectivity disruptions, degraded performance, or changes in user demand. Unlike APNIC, however, M-Lab captures user-initiated measurement activity rather than estimated user population. In practice, this means that the signal provides information about how users interact with the network (for example, when they experience degraded service or decide to run performance tests) and thus provides a complementary view of Starlink usage rather than a direct count of users [7, 20].

Cloudflare Radar. Cloudflare Radar provides a public view of Internet traffic trends observed across Cloudflare's global network.

Step	Purpose
EWMA smoothing	Reduce daily noise while preserving multi-week changes.
Piecewise linear fit	Identify dates where the dominant trend changes.
Slope-based merging	Remove weak changes with normalized slope-angle difference below θ.
Duration-based merging	Remove segments shorter than W days by merging them with the least-cost neighbor.
Event labeling	Treat remaining breakpoints as candidate events and label their direction.

Table 1: Per-signal candidate event extraction pipeline.

We use its Radar API to collect a time series of normalized HTTP requests associated with Starlink's Autonomous System in each country. This signal captures relative changes in web activity from Starlink users whose requests reach Cloudflare-hosted services. Because the dataset reports normalized rather than raw request counts, we use it to track trends over time rather than to estimate absolute numbers of users or traffic volume. Its main advantage is that it reflects passive, large-scale web activity, but it must be interpreted in light of Cloudflare's network coverage and customer base.

4 Cross-Source Trend Shift Detection

Our methodology detects candidate Starlink usage events by looking for temporally aligned structural changes in three datasets. We first process each time series separately, smoothing daily observations and fitting piecewise linear trends to identify dates where the dominant trajectory changes. We then merge weak or short-lived segments to retain only changes that are both pronounced and sustained. Finally, we compare the resulting event sets across datasets and flag shifts in the usage trend that occur in the same direction within a fixed temporal window. This cross-source agreement requirement helps distinguish meaningful shifts in Starlink usage from source-specific noise or artifacts.

4.1 Per-Dataset Event Detection

We first extract candidate events from each signal independently. The same procedure is applied separately each dataset. Table 1 summarizes the per-dataset pipeline. The algorithm identifies sustained changes in the dominant linear trajectory of each smoothed time series, then filters out weak or short-lived changes before returning dated candidate events. Unlike M-Lab and APNIC which offer daily data points, requiring smoothing, Cloudflare Radar only offers one data point per week. Thus, we do not smooth that dataset.

EWMA smoothing. Both APNIC and M-Lab can exhibit day-to-day variation that obscures the slower shifts of interest. We therefore smooth each daily series using an exponentially weighted moving average (EWMA) with a 30-day span. This span balances two undesirable extremes: a much shorter span would allow daily noise to leak into the segmentation, while a much longer span would blur the onset of the events we aim to detect. A 30-day span also aligns with the timescale of the disruptions examined in our case studies, which typically unfold over several weeks. Let $\{(t_i, \tilde{y}_i)\}_{i=1}^{n}$ denote the resulting smoothed series, where t_i is the day index and $\tilde{y}_i$ is the smoothed value.

Piecewise linear fit. The significance of a change in a Starlink usage signal depends strongly on the surrounding scale and trajectory. For example, in a country where the smoothed APNIC signal is around 100K users, an increase of another 100K over a few weeks represents a substantial shift. The same absolute increase against a baseline of 2M users would be far less pronounced and could simply reflect the continuation of an existing trend. Thus, the same absolute change can have different meanings depending on the context in which it occurs.

To capture these contextual changes, we model each time series as a sequence of contiguous linear segments rather than searching for isolated point anomalies. A new segment begins when the captured usage behavior is better explained by a different linear trend than by continuing the previous one (i.e., lower error between the data and the linear model). For each segment j, the fitted values are given by

$$\hat{y}_i = a_j t_i + c_j,$$

where a_j and c_j are the segment-specific slope and intercept. The breakpoints $\{b_1, \ldots, b_{m-1}\}$ and the per-segment parameters are chosen to minimize the total sum of squared errors:

$$\min_{\{b_j, a_j, c_j\}} \sum_{j=1}^{m} \sum_{i \in I_j} \left(\tilde{y}_i - (a_j t_i + c_j)\right)^2,$$

where I_j denotes the set of indices assigned to segment j. Each resulting breakpoint marks a date at which the dominant linear trend of the signal changes.

Slope-based merging. The initial segmentation may identify small course corrections that are not large enough to represent meaningful changes in usage. We therefore merge adjacent segments whose fitted slopes are too similar. Specifically, for each adjacent segment pair, we normalize the time and value axes and compute the angular difference between the fitted slopes. If the angular difference is below θ, we merge the two segments. This threshold establishes a minimum sharpness for a trend change to be retained.

Duration-based merging. The segmentation may also produce short-lived segments that reflect transient fluctuations rather than sustained shifts. We therefore enforce a minimum segment duration of W days. Any segment shorter than this threshold is merged with the neighboring segment that yields the smallest increase in total squared error. This duration threshold is long enough to filter out brief oscillations, while still allowing us to detect disruptions that unfold over several weeks.

Event labeling. After slope-based and duration-based merging, the remaining breakpoints are treated as candidate events. Each event is assigned the date of the breakpoint and labeled as either an increase or a decrease according to the direction of the post-break trend. The output of the per-time series pipeline is therefore a set of dated candidate events.

4.2 Cross-Source Agreement

After applying the per-dataset pipeline, we compare the three resulting event sets. We define a *correlated candidate event* as an event

in one dataset that has a same-direction event in the other datasets within a matching window of $\pm k$ days. Temporal co-occurrence is the primary criterion for cross-source agreement. We do *not* require the signals to have identical slope magnitudes or perfect alignment. Our reasoning is that each dataset observes the network from its unique vantage point, making perfect agreement improbable. For example, a real-world disruption may cause M-Lab testing activity to react immediately as users troubleshoot their networks, while APNIC may reflect a slower change as users shift to another network.

In addition to temporal alignment, we require paired events to share the same direction of change. An increase in one signal coinciding with a decrease in the other is unlikely to reflect the same external cause. We therefore use same-direction temporal agreement as the main criterion for cross-source support, and we interpret the detailed shape of paired changes qualitatively in the evaluation.

The parameter k, controls how strictly the two signals must align in time. A smaller k requires tighter temporal agreement, while a larger k admits more paired events at the risk of spurious co-occurrence. The choice of k is also constrained by W. Events within a single signal are separated by at least this duration, choosing $2k$ to be no greater than the minimum segment duration limits the possibility that one event in a signal matches multiple events in the other.

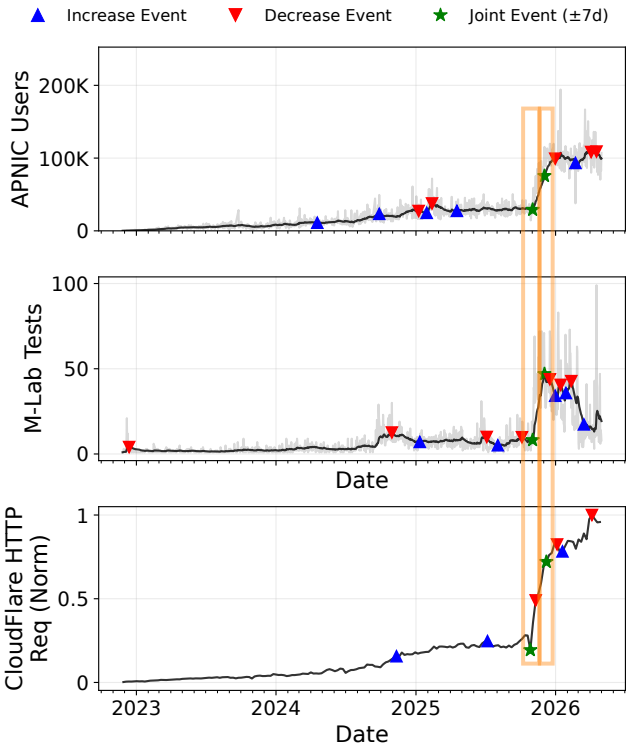

Figure 2: Jamaica: cross-source agreement around a hurricane-related disruption. All three public signals show a sustained increase in Starlink usage beginning in late October 2025, followed by a later slowdown in growth. The jointly detected events, highlighted in orange rectangles, coincide with the period when satellite connectivity became part of the Hurricane Melissa response and the subsequent recovery.

5 Case Studies

Methodology. We consider Starlink ASN 14593 and study Jamaica, Ukraine, and Germany. These countries were chosen to span both cases with known or suspected external triggers and cases where the public signals themselves appeared visually interesting. Jamaica was selected because of a known hurricane-related disruption; Ukraine was selected because of the known wartime expansion in Starlink usage; and Germany was selected because APNIC alone showed a prominent rise-and-fall pattern that appeared eventful. We use the same configurations across all datasets. In particular, we configure $\theta = 6°$, $W = 14$, and $k = 7$.

Case study 1: Jamaica. All datasets show the beginning of a large sustained increase in Starlink usage in late October 2025, despite measuring different aspects of usage and carrying largely independent noise. Moreover, they also detect the slowdown in that growth towards the end of the year. Our detection technique flags both events as shown in Figure 2. The context of the change, extracted from posts on social media and news articles, validates the significance of detected events. Hurricane Melissa made landfall on October 28, 2025, as a Category 5 storm. The Emergency Telecommunications Cluster reports that 77% of Jamaica lost power and internet access due to extensive network disruption [5]. Satellite connectivity quickly became part of the response. SpaceX contributed 1,000 Starlink kits, with 80% deployed in the most affected areas [5, 21]. Liberty Latin America separately reported working with Starlink Direct to Cell to provide emergency connectivity for FLOW Jamaica customers when local mobile infrastructure was unavailable [16]. The detected event is therefore consistent with a broader disruption of power and communications, in which satellite connectivity became part of the emergency response and early recovery effort.

Case study 2: Ukraine. All datasets show the onset of a substantial increase in Starlink usage in April 2025, which is flagged by our detection technique as shown in Figure 3. Several public reports help contextualize this shift. On April 3, 2025, *The Kyiv Independent* reported that Ukraine received 5,000 additional Starlink terminals from Poland. Ukrainian officials described these terminals as necessary for maintaining communications in de-occupied territories, where the destruction of cellular base stations had made regular service unavailable [23]. The IODA project also reported that repeated Russian attacks on Ukraine's energy grid from late 2024 through January 2025 had already caused measurable disruptions to Internet connectivity [17]. Taken together, these reports suggest that the April increase reflects both expanded Starlink availability and growing reliance on resilient communication paths amid degraded terrestrial infrastructure.

A second increase is detected in late August and early September 2025. This event coincides with a documented rise in attacks on Ukrainian infrastructure [22], suggesting renewed pressure on terrestrial power and communications networks. In this context, the observed increase in Starlink usage is consistent with a shift

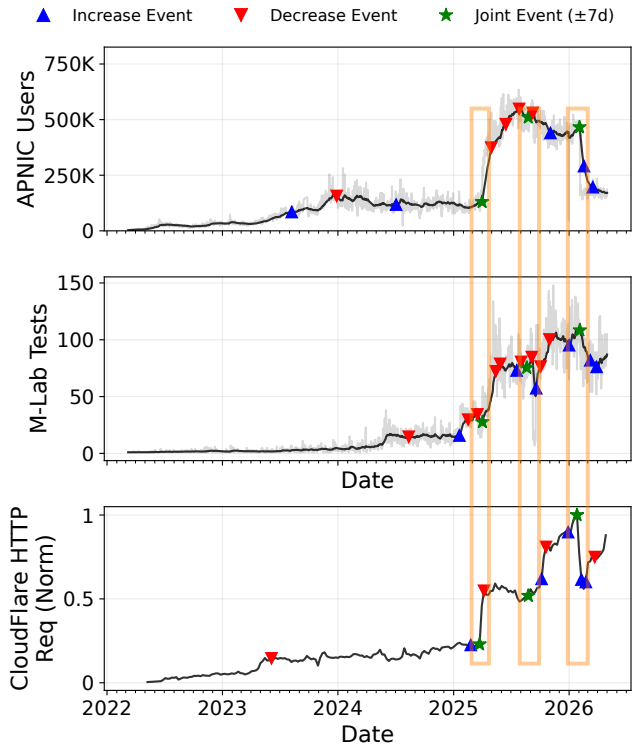

Figure 3: Ukraine: jointly detected shifts in Starlink usage. The public signals show a large increase in 2025, consistent with the reported expansion of Starlink availability and use, followed by later jointly detected changes, including a decline in early 2026. The orange rectangles mark same-direction events detected across datasets within the matching window.

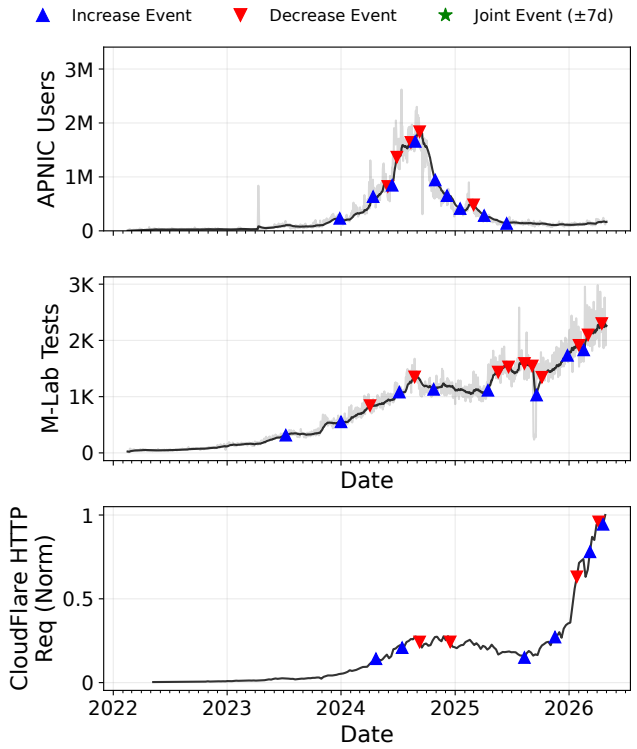

Figure 4: Germany: a prominent single-source pattern without cross-source support. APNIC shows a large rise and fall that appears eventful in isolation, but M-Lab and Cloudflare Radar do not show a corresponding country-wide pattern. This disagreement illustrates how the consensus rule filters out visually striking trends that are not supported across independent public vantage points.

toward more resilient satellite-based connectivity during periods of infrastructure disruption.

We also detected an event in early 2026 that corresponded to a downward shift in all datasets. This correlates with the Starlink shutdown and re-verification episode reported in early February 2026. In particular, *Al Jazeera* reported that SpaceX blocked the use of Starlink terminals geolocated on Ukrainian territory, after which only verified terminals added to frequently updated white lists could be restored [19]. *The Moscow Times* similarly reported a sweeping Starlink outage on February 5, 2026, citing Russian military bloggers [1].

Case study 3: Germany. As shown in Figure 4, APNIC exhibits a pronounced rise and fall that could easily be mistaken for an event. The other two datasets, in contrast, show a much steadier increase over the same broad period, with no corresponding rise and fall. This cross-source disagreement acts as an external sanity check on the APNIC pattern and is exactly the kind of filter our analysis is designed to provide.

6 Related Work

LEO Networks Measurement. A growing body of work characterizes the performance, behavior, and measurement visibility of Starlink and other LEO networks. Early end-user measurement studies [14, 18] provided some of the first public measurements of Starlink throughput, latency, temporal variation, and user-perceived performance. More recent work has broadened this view. For instance, the work in [20] combines large-scale crowdsourced measurements with targeted experiments to study Starlink across countries and applications, while WetLinks [15] contributes a longitudinal open dataset and analyzes weather sensitivity. Other work has focused on improving visibility into LEO networks or explaining specific behaviors; for example, the work in [12] introduces HitchHiking to lower the barrier to large-scale LEO measurement, and the work in [9] studies Starlink from a vantage point directly behind the dish to better understand periodic throughput drops.

Although this literature shows how public or semi-public measurements can reveal non-trivial aspects of Starlink behavior, their primary focus differs from ours: they mostly aim to characterize performance, routing, visibility, or measurement methodology, rather than to determine whether an apparent change in public usage signals corresponds to a real external event. More importantly, they typically rely on direct measurements, dedicated vantage points, or a single signal class and therefore do not address our central question: whether *agreement across independently noisy public aggregate datasets* can help distinguish plausible event-driven changes from source-specific noise.

Using APNIC and M-Lab Signals. Our work also connects to prior studies that use APNIC or M-Lab as public measurement signals.

APNIC Labs has published technical write-ups on the construction and interpretation of its ISP user-count estimates. Huston [10] describes how the per-ISP user signal is inferred and discusses its uncertainty. Huston's later work [11] shows that Starlink complicates economy-level interpretation through roaming, shipping, and geolocation anomalies that can inflate apparent user populations. M-Lab has been used to study Internet resilience under disruption [7]. Jain et al. [13] use M-Lab NDT data to analyze degradation and routing changes during the war in Ukraine. Closer to our setting, Mohan et al. [20] use M-Lab speed-test measurements to study global Starlink performance.

These studies use APNIC or M-Lab in isolation, or for goals other than cross-source event inference. Our work uses both signals jointly to surface Starlink-related events of interest. Temporal agreement between independently noisy sources provides a stronger public lens than either signal alone.

7 Limitations

Our method is intentionally conservative, identifying an event only when multiple public datasets show a same-direction structural change within a narrow temporal window. This high bar minimizes the risk of treating source-specific artifacts as real-world events, a critical safeguard since many Starlink usage shifts cannot be independently validated without private operator telemetry. Consequently, our detections represent high-confidence candidate events rather than an exhaustive catalog.

The primary tradeoff of this approach is that it favors precision over recall. By prioritizing the avoidance of false positives, the method inherently increases false negatives; requiring cross-source agreement discards genuine events that are uniquely visible to a single vantage point or expressed differently across measurements. Furthermore, our reliance on daily or weekly country-level aggregations smooths over short-lived outages, regional disruptions, and localized mobility effects. This limitation is particularly acute for LEO networks characterized by rapid routing shifts and dynamic user footprints.

Our analysis is also constrained by the inherent imperfections of IP geolocation. Grouping observations by country relies on databases that frequently lag behind operational changes, an issue compounded by Starlink's terminal roaming, cross-border activations, and shifting gateway allocations. Finally, cross-source agreement indicates synchronized behavioral changes rather than definitive causality. While external reports provide valuable context, they cannot prove a single real-world event caused the observed signal. Ultimately, our framework is designed to filter out isolated public noise and prioritize the most plausible candidate events for deeper investigation.

8 Conclusion

We presented a cross-source tool for detecting high-confidence shifts in Starlink usage from public measurements. The method combines APNIC, M-Lab, and Cloudflare Radar as complementary vantage points, identifying candidate events when multiple sources observe same-direction changes in usage trends within a narrow temporal window. Through case studies in Jamaica, Ukraine, and Germany, we show that this approach can surface plausible real-world changes while avoiding unsupported single-vantage-point interpretations. We hope this tool can help the measurement community track the evolving role of LEO networks in Internet access, resilience, and emergency connectivity.

Acknowledgments

We thank the anonymous reviewers and our shepherd Paulo Mendes for their constructive feedback that helped us improve this paper. This project was partially supported by NSF grant CNS 2345827.

References

[1] AFP. 2026. *Pro-War Bloggers Report Starlink Outage on Ukrainian Front Line.* The Moscow Times. https://www.themoscowtimes.com/2026/02/05/pro-war-bloggers-report-starlink-outage-on-ukrainian-front-line-a91875

[2] Vaibhav Bhosale, Ying Zhang, Sameer Kapoor, Robin Kim, Miguel Schlicht, Muskaan Gupta, Ekaterina Tumanova, Zachary S. Bischof, Fabián E. Bustamante, Alberto Dainotti, and Ahmed Saeed. 2025. Assessing LEO Satellite Networks for National Emergency Failover. In *Proceedings of the 2025 ACM Internet Measurement Conference (IMC '25)*. Association for Computing Machinery, New York, NY, USA, 274–293. doi:10.1145/3730567.3764482

[3] Zachary S. Bischof, Kennedy Pitcher, Esteban Carisimo, Amanda Meng, Rafael Bezerra Nunes, Ramakrishna Padmanabhan, Margaret E. Roberts, Alex C. Snoeren, and Alberto Dainotti. 2023. Destination Unreachable: Characterizing Internet Outages and Shutdowns. In *Proceedings of the ACM SIGCOMM 2023 Conference*. Association for Computing Machinery, New York, NY, USA, 608–621. doi:10.1145/3603269.3604883

[4] Eric Cheung. 2024. Developing Taiwan's own 'Starlink' crucial for island-wide emergency, space agency says. (Accessed on September 29th, 2025).

[5] Emergency Telecommunications Cluster. 2025. *Caribbean: Hurricane Melissa.* Emergency Telecommunications Cluster. https://www.etcluster.org/emergency/caribbean-hurricane-melissa

[6] The NATO Science for Peace and Security Programme. 2025. Hybrid Space-Submarine Architecture Ensuring Infosec of Telecomunications (HEIST) - A NATO Project. Available at https://natoheist.org/. Accessed on September 29th, 2025.

[7] Phillipa Gill, Christophe Diot, Lai Yi Ohlsen, Matt Mathis, and Stephen Soltesz. 2022. M-Lab: User Initiated Internet Data for the Research Community. *SIGCOMM Comput. Commun. Rev.* 52, 1 (March 2022), 34–37. doi:10.1145/3523230.3523236

[8] Jack Haddon. 2024. Ghana to licence Starlink in response to subsea cable cuts. (Accessed on September 29th, 2025).

[9] Sarah-Michelle Hammer, Vamsi Addanki, Max Franke, and Stefan Schmid. 2024. Starlink Performance through the Edge Router Lens. In *Proceedings of the 2nd International Workshop on LEO Networking and Communication (LEO-NET '24)*. Association for Computing Machinery, New York, NY, USA, 67–72. doi:10.1145/3697253.3697273

[10] Geoff Huston. 2024. *How we measure: ISP user counts.* APNIC Blog. https://blog.apnic.net/2024/11/11/how-we-measure-isp-user-counts/

[11] Geoff Huston. 2025. *Geolocation and Starlink.* APNIC Blog. https://blog.apnic.net/2025/09/30/geolocation-and-starlink/

[12] Liz Izhikevich, Manda Tran, Katherine Izhikevich, Gautam Akiwate, and Zakir Durumeric. 2024. Democratizing LEO Satellite Network Measurement. *Proceedings of the ACM on Measurement and Analysis of Computing Systems* 8, 1 (Feb. 2024), 1–26. doi:10.1145/3639039

[13] Akshath Jain, Deepayan Patra, Peijing Xu, Justine Sherry, and Phillipa Gill. 2022. The Ukrainian Internet under Attack: An NDT Perspective. In *Proceedings of the 22nd ACM Internet Measurement Conference (IMC '22)*. Association for Computing Machinery, New York, NY, USA, 166–178. doi:10.1145/3517745.3561449

[14] Mohamed M. Kassem, Aravindh Raman, Diego Perino, and Nishanth Sastry. 2022. A Browser-Side View of Starlink Connectivity. In *Proceedings of the 22nd ACM Internet Measurement Conference (IMC '22)*. Association for Computing Machinery, New York, NY, USA, 151–158. doi:10.1145/3517745.3561457

[15] Dominic Laniewski, Eric Lanfer, Bernd Meijerink, Roland van Rijswijk-Deij, and Nils Aschenbruck. 2024. WetLinks: A Large-Scale Longitudinal Starlink Dataset with Contiguous Weather Data. In *2024 8th Network Traffic Measurement and Analysis Conference (TMA)*. IEEE, Piscataway, NJ, USA, 1–9. doi:10.23919/TMA62044.2024.10558998

[16] Liberty Latin America. 2025. *Liberty Latin America Working with Starlink Direct to Cell to Provide Emergency Service in Aftermath of Hurricane Melissa in Jamaica.* Liberty Latin America. https://lla.com/blog/liberty-latin-america-working-starlink-direct-cell-provide-emergency-service-aftermath-hurricane-melissa-jamaica

[17] Amanda Meng and Tara Kelly. 2025. *How Russia's Recent Attacks on Ukraine's Energy Grid Impacted its Internet Connectivity.* IODA Reports. https://ioda.inetintel.cc.gatech.edu/reports/how-russias-recent-attacks-on-ukraines-energy-grid-impacted-its-internet-connectivity-2/

[18] François Michel, Martino Trevisan, Danilo Giordano, and Olivier Bonaventure. 2022. A First Look at Starlink Performance. In *Proceedings of the 22nd ACM Internet Measurement Conference (IMC '22)*. Association for Computing Machinery, New York, NY, USA, 130–136. doi:10.1145/3517745.3561416

[19] Mansur Mirovalev. 2026. *How does the cutoff of Starlink terminals affect Russia's moves in Ukraine?* Al Jazeera. https://www.aljazeera.com/news/2026/2/10/how-does-the-cutoff-of-starlink-terminals-affect-russias-moves-in-ukraine

[20] Nitinder Mohan, Andrew E. Ferguson, Hendrik Cech, Rohan Bose, Prakita Rayyan Renatin, Mahesh K. Marina, and Jörg Ott. 2024. A Multifaceted Look at Starlink Performance. In *Proceedings of the ACM Web Conference 2024 (WWW '24)*. Association for Computing Machinery, New York, NY, USA, 2723–2734. doi:10.1145/3589334.3645328

[21] Starlink. 2025. *Starlink tweet in response to Hurricane Melissa.* X (formerly Twitter). https://x.com/Starlink/status/1985494339796218202

[22] United Nations in Ukraine. 2025. Civilian Casualties Remain Alarmingly High as Short and Long-Range Weapons Devastate Lives Across Ukraine, UN Human Rights Monitors Say. https://ukraine.un.org/en/301284-civilian-casualties-remain-alarmingly-high-short-and-long-range-weapons-devastate-lives. Press release. Accessed 14 May 2026.

[23] Tim Zadorozhnyy. 2025. *Ukraine receives 5,000 more Starlink terminals from Poland, minister says.* The Kyiv Independent. https://kyivindependent.com/ukraine-receives-5-000-more-starlink-terminals-from-poland-minister-says/