Characterizing traffic destinations and temporal trends for adaptive network resource management in 5G/6G networks

Vlad-Ioan Dragutoiu[1]

Supervisors: Dr. Nitinder Mohan[1]**, Marco Colocrese**[1]

[1]EEMCS, Delft University of Technology, The Netherlands

A Thesis Submitted to EEMCS Faculty Delft University of Technology,
In Partial Fulfilment of the Requirements
For the Bachelor of Computer Science and Engineering
June 22, 2025

Name of the student: Vlad-Ioan Dragutoiu
Final project course: CSE3000 Research Project
Thesis committee: Dr. Nitinder Mohan, Marco Colocrese, Guohao Lan

Abstract

Modern mobile networks must adapt to rapidly changing traffic patterns and increasing user demands. A key challenge is understanding where user traffic terminates and how these destinations vary over time. This thesis addresses this challenge by introducing an open-source, modular analysis framework that analyzes passive Internet traffic traces, enriches them with geolocation and organizational metadata, and infers latency stability and routing dynamics, in order to characterize the infrastructures that terminate user traffic and assess their performance and reliability over time. The results show a long-term shift towards content-centric traffic, highlight geographic and temporal variations in performance, and demonstrate that content networks typically offer greater stability than enterprise or research destinations. These findings support adaptive traffic management strategies in 5G and future 6G networks.

1 Introduction

Next-generation mobile networks, from 5G evolving towards 6G, are transforming communication infrastructure design and operation. Traditional static resource provisioning, specifically dimensioning networks for peak loads, leads to inefficiencies during off-peak periods and struggles to accommodate the exponential growth in data demand and new application requirements. Recent industry forecasts underscore this trend: global mobile data traffic is projected to reach about 370 exabytes per month by 2030, growing at a compound annual growth rate (CAGR) of 25%, with 5G connections comprising more than half of all mobile links and early 6G deployments on the horizon [1–3]. Emerging applications such as immersive augmented reality (AR), autonomous vehicles, and large-scale Internet of Things (IoT) deployments demand ultralow latency and high throughput, pushing networks to become highly adaptable and intelligent in how they manage resources. Meeting these demands requires shifting from static provisioning to dynamic, predictive resource management strategies informed by accurate traffic analysis.

A crucial aspect of adaptive network management is understanding where user traffic is going, namely the termination endpoints, and how these destination patterns change over time. Previous research has contributed extensively to the modeling and forecasting of Internet traffic. Navarro-Ortiz et al. [4] provide a comprehensive survey of 5G usage scenarios and traffic models, while Alawe et al. [5] and Papagiannaki et al. [6] demonstrate how machine learning and statistical techniques, respectively, can forecast traffic trends. In particular, deep learning models and time series forecasting have been shown to achieve low-error predictions of mobile traffic. However, most of these studies focus on predicting aggregate traffic volume and do not address where traffic is terminating. Others, such as Candela et al. [7] and Feldmann et al. [8], have explored spatial and temporal variations in Internet routing, including geographic locality and the effects of global events such as the COVID-19 pandemic. However, few studies integrate these dimensions or resolve the fine-grained destination infrastructure that user traffic targets.

This gap is significant: knowing which autonomous systems (ASes), content delivery networks (CDNs), or cloud providers terminate user flows, and how those patterns shift, is critical for informed, adaptive network control. For instance, if a network operator knows that a large fraction of traffic terminates at unstable or high-latency endpoints, they may proactively reroute traffic or provision additional resources. Conversely, highly stable and responsive destinations, such as CDN edge nodes, may be prioritized for low-latency services. In this context, infrastructure-aware traffic characterization becomes essential for technologies such as traffic steering, network slicing, and predictive resource allocation.

This research investigates the following central question:

> *How can user traffic be characterized in terms of its termination across various infrastructures, and what temporal patterns can be observed in these traffic flows?*

From this broad inquiry, we derive three specific research subquestions (RSQs):

(1) How can the termination points of user traffic be systematically identified across different types of network infrastructure?
(2) What insights can be gained by geolocating and attributing user traffic flows to specific content providers, CDNs, cloud data centers, or other infrastructure categories?
(3) What temporal trends manifest in where traffic terminates, for example, patterns linked to time-of-day, day-of-week, or seasonal variations in dominant destinations?

In order to bridge the aforementioned gap and address the research questions, we develop a reproducible methodology that combines passive traffic measurement with external metadata to capture the spatial and temporal characteristics of user traffic destinations. This work contributes both technical methods and empirical insights that enable infrastructure-aware traffic analysis and support adaptive management strategies in modern mobile networks.

The key contributions of this thesis are the following:

- **MANTA framework:** An open-source, modular system for infrastructure-aware traffic analysis that enriches passive flow data and integrates data-plane and control-plane metrics to classify destinations, assess latency stability, and quantify routing volatility using only public datasets and tools.
- **Latency KPI:** A passive, privacy-preserving technique for estimating round-trip time (RTT) from TCP timestamps, used to derive prefix-level latency stability metrics without injecting probe traffic.
- **Control-plane KPI:** A BGP-based stability assessment that quantifies route churn, update frequency, and visibility diversity to evaluate the reliability of traffic destinations over time.
- **Composite volatility scoring:** A unified metric that integrates latency variability and BGP path churn to identify unstable prefixes and support risk-aware traffic decisions.

- **Use-case-driven insights:** Demonstrated applications of infrastructure-aware metrics for traffic steering, slice placement, cache replication, and SLA validation in 5G/6G networks.

The remainder of this paper is structured as follows: Section 2 reviews related work on traffic analysis, latency estimation, and control-plane dynamics. Section 3 presents the methodology. Section 4 outlines the core analytical contributions and how they address the defined research questions. Section 5 describes the experimental setup and summarizes key results. Section 6 discusses ethical and responsible research considerations. Section 7 reflects on the broader implications of the findings, highlights limitations, and suggests potential extensions and directions for future work. Finally, Section 8 concludes the thesis.

2 Related work

Effective solutions for adaptive, predictive traffic management require both a solid understanding of prior research and recognition of open gaps. A range of studies have addressed Internet traffic modeling, forecasting, and spatial analysis, often with promising results. However, most lack a unified, infrastructure-aware perspective that can directly inform operational decisions. Here, we review existing work on traffic prediction, spatial attribution, and measurement methodologies, emphasizing their relation to our goals.

Traffic modeling and prediction. A substantial body of work focuses on modeling traffic demand and forecasting future load. Early efforts by Papagiannaki et al. [6] applied statistical models (e.g., ARIMA) to predict long-term backbone traffic trends. More recent work has explored fine-grained, dynamic predictions suitable for mobile networks. Alawe et al. [5] used machine learning to forecast traffic in a 5G core network, demonstrating improved accuracy and scalability. Similarly, Kochetkova et al. [9] showed that classical time-series models (seasonal ARIMA, Holt-Winters) can achieve low error rates for short-term (hour-ahead) 5G traffic forecasting. Deep learning approaches have also shown promising results: Aouedi et al. [10] provide a comprehensive survey of RNN-, CNN-, and GNN-based methods for next-generation traffic forecasting. These studies confirm that accurate forecasting is crucial for enabling proactive resource management in 5G/6G systems. However, they predominantly focus on predicting aggregate volume and do not resolve traffic endpoints or infrastructure types.

Spatial and infrastructure-aware traffic analysis. Understanding where traffic flows terminate, i.e., the spatial distribution of destinations and the nature of those endpoints, is increasingly important. Candela et al. [7] conducted a worldwide study on the geographic locality of Internet routes, revealing performance implications of local vs. distant routing. Feldmann et al. [8] examined pandemic-induced changes in traffic patterns, showing significant shifts in application usage and diurnal trends. Bajpai et al. [11] performed a longitudinal analysis of the MAWI backbone dataset, documenting decade-long changes in traffic composition, such as increased HTTPS and VPN traffic. These works underscore the value of passive datasets (e.g., CAIDA [12] and MAWI [13]) for understanding traffic evolution and performance. However, few studies integrate destination classification, control-plane behavior, and temporal variability into a single analytic framework, a gap our approach addresses.

Adaptive networking and control-plane insight. Recent research has begun to bridge traffic analysis with adaptive network control. Emerging frameworks, particularly for 5G/6G, emphasize dynamic resource allocation based on predicted demand. Network slicing, for instance, enables virtualized network instances that adapt to traffic forecasts. Surveys by GSMA [3], Cisco [2], and Ericsson [1] emphasize predictive analytics as a core component of future mobile infrastructure. Cui et al. [14] envision AI-driven 6G network management systems that continuously optimize based on real-time traffic signals. Yet, these visions often overlook destination specificity, where traffic is going, and how reliably it gets there.

Moreover, many prior efforts treat control-plane and data-plane measurements in isolation. Our work instead ties BGP instability metrics (e.g., path churn, AS path diversity) to observed traffic destinations. This is motivated by recent studies such as Darwich et al. [15], who report that traffic-engineering events account for 39% of BGP updates and nearly 44% of routing convergence time, demonstrating how dynamic routing behavior can impact performance. Incorporating this control-plane visibility into traffic engineering enables more intelligent decisions, such as avoiding unstable paths or deprioritizing volatile prefixes.

Summary. Prior research provides a strong foundation for traffic forecasting and spatial analysis. However, few studies offer an integrated view linking *when*, *where*, and *how* traffic terminates across infrastructure. Our work fills this gap with a reproducible methodology that incorporates passive RTT, BGP stability, and infrastructure-level classification to inform adaptive management in future networks.

3 Methodology

This section presents the methodological foundation of this study, structured around the MANTA framework introduced earlier. The goal is twofold: first, to explain the motivation for building this framework: why this particular approach was used to answer the research questions, as well as reasons behind key decisions; and second, to detail how the framework is built, what data it uses, and how each source incrementally contributes new dimensions of insight.

3.1 Motivation and design rationale

Characterizing where and how Internet traffic terminates, and how these patterns evolve over time is critical for understanding infrastructure dependencies, assessing performance bottlenecks, and enabling adaptive resource allocation in 5G and future 6G networks [1, 3]. Yet, existing methodologies often fall short in supporting this level of visibility and operational insight. Common limitations include:

- **Lack of destination-layer granularity:** Most traffic modeling and forecasting approaches focus on predicting aggregate traffic volume across network segments or at coarse geographic levels [4, 5]. This neglects the identification and classification of traffic endpoints, such as CDNs, data centers, or ISPs, which are vital for understanding service reliance and interconnection needs.

- **Dependence on proprietary or active measurement tools:** Commercial telemetry platforms and active probing systems (e.g., ping-based latency monitoring or synthetic test traffic) often restrict scalability, incur operational overhead, or lack transparency. Such methods may be unsuitable for longitudinal, wide-scale studies or may be blocked by firewalls and endpoint configurations [8, 16].
- **Limited integration of control-plane visibility:** Traffic analyses are frequently limited to the data plane, ignoring how BGP-level dynamics affect path stability and prefix reachability [7, 14]. Without insights into routing churn or visibility, it is difficult to understand the reliability and resilience of destination paths.

To address these gaps, this study proposes a fully passive, infrastructure-aware framework that operates entirely on publicly available data. The design of this framework is shaped by four guiding principles:

1. Reproducibility and transparency. To ensure that the analysis can be replicated or extended by other researchers and operators, the framework relies entirely on publicly available datasets and open-source tooling as you will see in Subsection 3.2. These were deliberately selected for their wide adoption and long-term community trust, ensuring that results are based on authoritative and accessible inputs. The methodology is designed to be modular and auditable, with outputs that can be directly traced back to source datasets.

2. Practical implementation choices. The framework is implemented entirely in Python for accessibility and flexibility, using well-supported, public libraries. Python was chosen for its large ecosystem of data processing tools, ease of use in rapid prototyping, and readability for reproducibility. A Jupyter Notebook-based environment complements the pipeline by supporting interactive exploration, visualization, and debugging, making the system suitable for both research and operational contexts.

3. Infrastructure-level insight. Unlike many prior works that treat destinations as opaque endpoints, this framework classifies traffic at the AS and prefix level, and maps each to a functional role. This supports analyses that differentiate traffic to CDNs, access ISPs, cloud providers, and enterprise networks, critical for interpreting infrastructure reliance and optimizing network adaptation.

4. Integration of passive KPIs. The pipeline infers round-trip time (RTT) using TCP timestamps embedded in the traffic, avoiding the need for active probes or synthetic traffic. Additionally, BGP path dynamics are integrated to assess control-plane stability for each prefix, providing a cross-layer view of destination performance and reliability.

Alternative design considerations. Alternative approaches, such as using commercial telemetry platforms, NetFlow data, or active probing tools were intentionally avoided due to limitations in accessibility, privacy, or realism. Proprietary platforms often lack transparency, while active probing introduces synthetic traffic and may be blocked or rate-limited. Similarly, NetFlow lacks the packet-level detail needed for precise latency estimation or timestamp-based inference. The chosen architecture prioritizes openness, depth, and extensibility without sacrificing operational relevance.

Building on these principles, the next subsection details the architecture and implementation of the proposed solution.

3.2 MANTA Framework

The MANTA framework is implemented as a modular pipeline in Python, designed to process passive Internet traffic traces and enrich them with spatial, organizational, and performance-related metadata. Built with extensibility and parallelism in mind, the system follows a layered enrichment model in which each processing stage contributes an independent dimension of insight. The framework is modular by design: each component can be executed independently, reused across datasets, or adapted to different vantage points. This flexibility supports deployment in a wide range of contexts, including research, network operations, and policy analysis.

The remainder of this subsection details the core stages of the pipeline:

1. Passive traffic trace collection. The analysis begins by sourcing packet-level Internet traffic traces from two publicly available, high-quality datasets: MAWI [13] and CAIDA [12].

The MAWI archive provides anonymized daily traces of trans-Pacific traffic captured on Japan's WIDE backbone. For this study, all MAWI data was sourced from Samplepoint-F, a long-standing monitoring location situated at the transit link between the WIDE backbone and a major upstream ISP. Samplepoint-F is widely used in academic work due to its consistency, bidirectional vantage point, and its ability to capture representative backbone traffic without aggressive sampling or filtering. Traces were selected across different years, months, and days to ensure temporal diversity and robustness. Each trace includes real-world IPv4 packet headers with minimal transformation, making them ideal for infrastructure-focused traffic analysis.

The CAIDA anonymized Internet traces, by contrast, were collected from backbone links operated by Tier-1 ISPs in the United States. For this project, all CAIDA data was taken from the NYC Equinix monitor, one of the most stable and well-documented vantage points in the CAIDA trace set. The NYC site offers high-volume visibility into inter-domain transit traffic within a key North American exchange point, adding geographic and topological complementarity to the MAWI traces. CAIDA's dataset is valued for its methodological rigor, reliable timestamping, and consistent IP anonymization.

Together, these datasets provide complementary perspectives: MAWI offers a view from an academic-transit border in East Asia, while CAIDA delivers a North American core-Internet perspective. This dual vantage point enables robust cross-regional comparisons and strengthens the generalizability of findings across global Internet infrastructure.

2. Packet ingestion and flow assembly. Raw PCAP files are ingested using the Scapy library. Packets are filtered to retain only IPv4 transport-layer traffic (TCP and UDP). Each packet is parsed to extract fields such as source/destination IPs, ports, protocol, TTL, timestamps, and TCP-specific headers, creating unidirectional flows. To ensure scalability, the ingestion stage is parallelized using Python's multiprocessing module, enabling concurrent processing of multiple trace files or batch segments. To maintain consistency across samples and ensure computational feasibility, processing for each file is capped at the first one million usable packets (i.e., those matching the protocol filter). This constraint was chosen to balance representativeness with practical limits on runtime and memory usage during repeated or large-scale runs. The threshold is a configurable parameter and can be increased or removed in future analysis depending on available resources and desired granularity.

3. Enrichment with IP metadata. Destination IPs are enriched using offline snapshots of GeoLite2 [17] and IPInfo Lite [18] databases. These sources provide country, city, latitude/longitude, ASN, and organization metadata. GeoLite2, developed by MaxMind, offers one of the most widely used and trusted IP geolocation databases in the research and commercial sectors. The GeoLite2 database is accessed using the geoip2 library, which enables efficient batch querying and structured metadata extraction. IPInfo Lite serves as a secondary source, enhancing coverage and accuracy. A fallback strategy ensures that if one source fails to resolve an IP, the other is used. Enrichment functions are wrapped for batch efficiency and reusability, and missing or ambiguous entries are flagged for post-filtering.

4. Prefix and ASN attribution. Using the CAIDA' Routeviews Prefix-to-AS dataset [19], destination IPs are matched to BGP-announced prefixes and their corresponding origin ASNs. This step is implemented with a prefix trie (pytricia), enabling efficient longest-prefix matching across millions of entries. Each flow is tagged with the most specific prefix available, and stored with its AS number and origin organization. CAIDA's prefix-to-AS mapping offers authoritative, real-time snapshots of the global routing table, making it ideal for attributing traffic accurately to originating ASes.

5. Infrastructure classification. Each ASN is mapped to a functional category (NSP, Content, Cable/DSL/ISP, Enterprise, Educational/Research, Non-Profit, Route Server, Network Services, Route Collector, Government) using a static snapshot of PeeringDB's public JSON API [20]. Mappings are resolved into canonical roles to support higher-level grouping and visualization. For example, Amazon and Akamai ASNs are collapsed into a "Content" class. This classification gives semantic context to the traffic and enables infrastructure-layer interpretation.

6. Passive RTT estimation. Round-trip time (RTT) is inferred passively using TCP timestamp options (TSval, TSecr) extracted from bidirectional flows. When both directions of a TCP connection are observed in the trace, the echoed timestamp (TSecr) from a returning packet is matched with a previously seen TSval, allowing for the computation of RTT based on capture timestamps.

To ensure robustness:

- Only valid TSval/TSecr pairs are considered, with RTT constrained to 0–5 seconds.
- A timestamp map is maintained for each flow direction to resolve symmetry and deduplicate matches.
- At least five RTT samples are required per prefix to ensure statistical confidence.

For each prefix, the system computes:

- Mean RTT (μ_{RTT})
- Standard deviation (σ_{RTT})
- Instability index (σ/μ)

This step yields a latency performance profile for each destination prefix without requiring active probes or instrumentation.

7. Control-plane path analysis. Both 8-hourly RIB (Routing Information Base) snapshots and 5-minute BGP update dumps, from RIPE RIS collector RRC06 [21], were parsed using the high-performance tool bgpscanner, executed on Ubuntu. For every destination IP in the passive flow dataset, the following steps are executed:

- Prefix matching: Each destination IP is matched against the longest-prefix entry in BGP update logs and RIBs using prefix tries built with PyTricia.This enables precise attribution of control-plane dynamics to the corresponding flows.
- BGP update analysis: For matched prefixes, the update stream is scanned to compute:
  - Total number of BGP updates (bgp_events) referencing the prefix
  - Count of distinct AS paths observed over time
  - Path change count, tracking how many times the AS path changed chronologically
- RIB snapshot analysis: Each prefix is also checked against RIB entries to determine:
  - If it is present in the RIB (in_rib)
  - Total number of distinct AS paths
  - Most recent AS path and its hop count

RRC06 was selected for its vantage point at Japanese IXPs (DIX-IE, JPIX), providing visibility into APAC routing dynamics relevant to the MAWI traces.

8. Output generation and aggregation. All processed flows are exported to structured CSV files, tagged with enrichment metadata and per-flow performance indicators. Aggregation scripts generate prefix and ASN level summaries, which are used for temporal analysis, categorical breakdowns, and visualization. Plots are generated using matplotlib, seaborn and plotly, and support filtering by infrastructure role, geographic region, and organization.

4 Infrastructure-aware flow analysis for adaptive network management

This section presents the core analytical contributions of our work and details how they address the research sub-questions introduced in Section 1. Our analysis spans multiple dimensions of network behavior: identifying where user traffic terminates, assessing latency stability in the data plane, quantifying routing dynamics in the control plane, combining these perspectives into a unified volatility metric, and last but definitely not least: how we did all of the above. These components are complemented by practical use cases that demonstrate how the resulting metrics can inform adaptive traffic management and infrastructure decision-making in 5G/6G networks.

4.1 Infrastructure-enriched analysis pipeline for traffic termination insight

MANTA enriches passive traffic traces with infrastructure metadata to reveal where traffic terminates and which types of networks are responsible for delivery. By classifying flows using publicly available geolocation, ASN attribution, and infrastructure role data (as detailed in Section 3), the system resolves destinations into operationally meaningful categories.

This directly answers RSQ1, enabling operators and researchers to observe traffic distribution across infrastructure roles and regions. The pipeline's modular design and reliance on public data ensure reproducibility across vantage points and datasets.

Use case illustration: A mobile operator notices rising evening latency. MANTA reveals that the majority of affected flows target a CDN cluster exhibiting high RTT variability and frequent BGP changes. Acting on this, the operator shifts traffic to more stable CDN endpoints and reallocates resources in anticipation of peak load, demonstrating how infrastructure-aware analytics can guide real-time adaptation.

4.2 Prefix-level RTT variance as a latency KPI

To assess data-plane performance, we derive a passive latency stability metric for each destination prefix using round-trip time (RTT) estimates inferred from TCP timestamps, a methodology originally proposed by Veal et al. [22]. Instead of relying on active probes, MANTA aggregates flow-level RTT observations and computes key statistics per prefix: the mean RTT, standard deviation, and an instability index. More details on this can be found in Section 3).

This addresses RSQ2 by highlighting prefixes with consistently low latency, typically associated with CDN or cloud infrastructure, as well as those with high variability due to congestion, suboptimal routing, or geographic dispersion, as also noted in forecasting surveys that link latency variability to underlying infrastructure and routing diversity [10]. Temporal aggregation further supports RSQ3 by surfacing diurnal and weekly trends in prefix-level latency behavior. Similar patterns have been observed in long-term backbone traffic analysis by Bajpai et al. [11], who document structural shifts and daily variation in Internet usage using MAWI trace data.

Applications:

- Prioritize low-volatility destinations for latency-sensitive traffic
- Perform passive SLA monitoring without injecting test traffic
- Detect early-stage latency regressions in production environments

This method is fully passive and privacy-preserving, making it well-suited for real-time deployment in operational networks.

4.3 BGP path stability as a control-plane KPI

MANTA also evaluates control-plane stability by analyzing BGP dynamics for each observed destination prefix, using update streams and RIB snapshots from RIPE RIS [21]. We quantify:

- BGP (AS path) churn—how often the routing path changes
- Update frequency and temporal clustering
- Route visibility—presence and diversity in routing tables

This complements the latency analysis by addressing RSQ2 from a routing perspective and supports RSQ3 through temporal analysis of path variability. Darwich et al. [15] show that BGP dynamics, particularly frequent updates and path changes, are often the result of instability or aggressive traffic engineering, reinforcing the need to monitor control-plane behavior over time.

Operational benefits:

- Prefer destinations with stable, predictable routing
- Flag unstable prefixes for cautious or deferred traffic placement
- Assess control-plane health and resilience trends over time

By integrating these control-plane KPIs with flow-level metrics, the framework bridges routing behavior and end-user performance.

4.4 Composite infrastructure volatility scoring

To synthesize performance and stability dimensions, we define a composite volatility score for each prefix:

$$V(p) = \alpha \cdot \frac{\sigma_{\text{RTT}}}{\mu_{\text{RTT}}} + \beta \cdot \text{BGP\_Churn}(p)$$

This index combines data-plane variability and control-plane instability, helping operators detect unreliable destinations that may degrade user experience. Prior work by Alawe et al.[5] and Kochetkova et al.[9] highlights the importance of forecasting traffic and performance fluctuations in mobile networks, particularly under volatile routing and congestion conditions. A higher score flags prefixes with erratic latency or frequent path changes, supporting risk-aware traffic steering and resource allocation.

This augments RSQ2 by offering a single, actionable metric to rank destination volatility and supports network adaptation strategies aimed at reliability.

4.5 Policy-oriented use cases for 5G/6G networks

Together, these metrics support adaptive management decisions critical to next-generation mobile networks. Example use cases include:

- **Traffic steering:** Route interactive or real-time traffic to low-volatility destinations to reduce jitter and delay.
- **Slice placement:** Deploy latency-critical services on infrastructure with stable RTT and BGP profiles.
- **Cache replication:** Mirror content to destinations with low volatility scores to improve delivery reliability.
- **SLA enforcement:** Monitor passive metrics to ensure compliance for priority traffic without active probes.

These examples show how MANTA's insights translate into concrete, policy-driven actions, helping networks evolve toward predictive, resilient operation. This aligns with visions outlined by Cui et al.[14], who describe AI-driven architectures for dynamic 6G network control, and by Navarro-Ortiz et al.[4], who emphasize the need for intelligent traffic management strategies in emerging mobile networks.

5 Experimental setup and results

This section describes the experimental environment used to implement and run MANTA, followed by a presentation and analysis of the key results obtained from applying the framework.

5.1 Experimental Setup

The MANTA analysis environment was implemented using a combination of Python modules and Jupyter notebooks. Core development was conducted on Windows 11 (64-bit), using PyCharm as the integrated development environment and Miniconda3 to manage the Python environment. All processing was performed using Python 3.10.

Control-plane analysis stages were executed using bgpscanner on an Ubuntu 22.04 virtual machine to ensure compatibility and performance. The setup supported efficient batch processing, trace selection, and visualization workflows using locally stored datasets and offline metadata snapshots.

By combining an interactive, modular design with a reproducible environment and full local control over data sources and libraries, the experimental setup ensured analytical traceability and adaptability for future studies.

5.2 Results

We evaluate MANTA's output across four dimensions: ASN destination roles, geographic RTT stability, role-based volatility, and AS path diversity. Each result maps directly to our research subquestions (RSQs), with cross-layer enrichment enabling interpretation at both infrastructure and temporal levels.

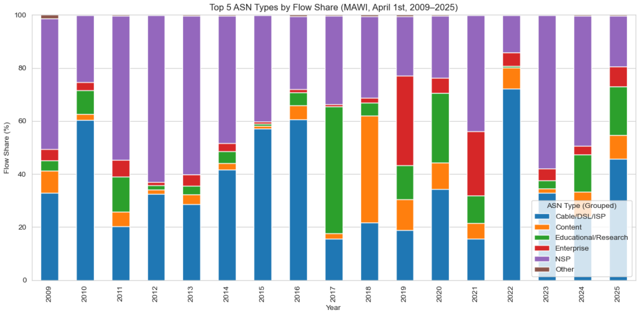

Figure 1: Top 5 ASN types by flow share (MAWI, 2009–2025, April 1st, 14:00 JST)

5.2.1 Longitudinal evolution of ASN yypes. Figure 1 illustrates how traffic destinations evolved over 16 years. Early traffic patterns were dominated by ISP ASNs, but this began shifting in 2018 as content and cloud networks expanded rapidly, likely due to CDN proliferation and streaming service growth. NSPs maintained consistently high shares, while enterprise and research networks exhibited periodic spikes, likely linked to episodic events or organizational migrations. These trends substantiate RSQ 1, confirming both the centralization of user traffic around large infrastructure providers and the classification fidelity of MANTA's enrichment pipeline.

This shift illustrates the growing role of CDNs and hyperscale clouds in handling user traffic, which aligns with industry observations of Internet centralization and its operational implications [16]. As content providers become dominant endpoints, peering policies and traffic engineering strategies must increasingly account for fewer but more critical interconnection points. (Additional breakdowns using CAIDA data are included in Appendix A.)

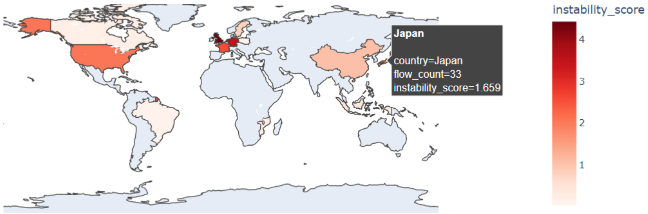

Figure 2: Geographical RTT instability by country (Content ASNs, 2015-03-30 – 2015-04-05, 14:00 JST)

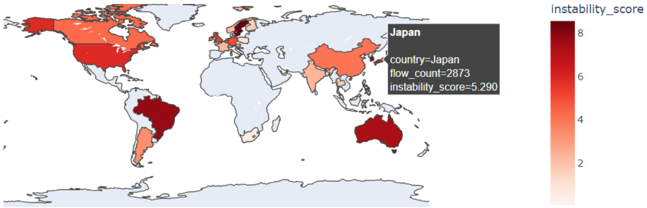

Figure 3: Geographical RTT instability by country (Content ASNs, 2025-03-31 – 2025-04-06, 14:00 JST)

5.2.2 Geographical RTT instability distribution over time. Figures 2 and 3 depict the evolution of RTT instability over a decade. On average, the volatility index nearly tripled between 2015 and 2025. Japan's score rose from 1.66 to 5.29; Brazil and Australia remained consistently volatile, and new hotspots emerged in North America and Southeast Asia. These shifts may reflect increased traffic loads, uneven infrastructure growth, or route complexity introduced by CDN overlays.

Another contributing factor could be the proliferation of content distribution architectures: while CDNs reduce average latency for well-connected users, they can also increase variability by widening the performance gap between regions with nearby caches and those without. Furthermore, growing inter-regional traffic flows and multi-homing of content platforms likely contributed to more frequent route changes and asymmetric paths, reinforcing temporal and geographic instability. Similar trends have been reported in public Internet metrics dashboards [16]. These findings enhance RSQ 2 by showing country-level performance variance, and address RSQ 3 through longitudinal comparison. They also underscore the need for geographically aware traffic engineering, particularly the deployment of edge caches or interconnection hubs in emerging hotspots.

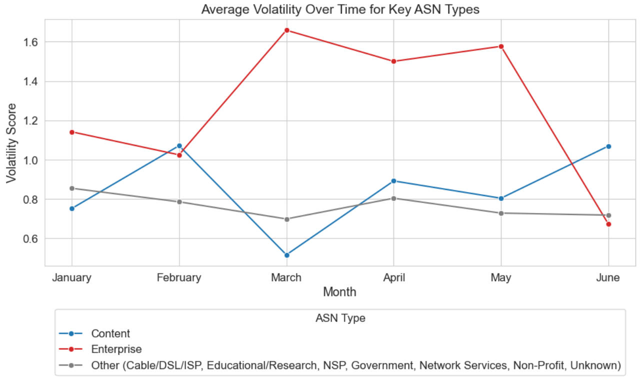

Figure 4: Monthly prefix volatility by ASN type (Jan–Jun 2025)

5.2.3 Temporal volatility across ASN categories. Figure 4 presents prefix-level volatility scores by ASN type, using the composite index from Section 4.4 (with $\alpha = \beta = 0.5$). This score aggregates RTT instability and BGP path churn into a unified metric for operational reliability.

Enterprise destinations peaked at a volatility index of approximately 1.7 in March, signaling possible reconfigurations, routing reshuffles, or sudden load migrations. This could be due to smaller organizations lacking CDN-style redundancy or undergoing infrastructural transitions, such as migrating workloads to cloud environments. In contrast, content ASNs remained stable (around 0.8) and gradually converged with enterprise values by June, both ending near $\sim$ 1.0. This convergence may reflect infrastructural normalization, expanded interconnectivity, or seasonal load balancing.

These findings directly support RSQ 2 by exposing volatility differences across infrastructure roles and RSQ 3 by highlighting evolving reliability trends. While seasonal effects (e.g., fiscal quarter transitions) may explain some March volatility spikes, more granular data would be needed for confirmation. Color gradients in the figure match the volatility scale to aid visual interpretation.

5.2.4 Content AS path diversity vs. latency (6 months vs. one week).

Six-month overview. Figure 5 correlates AS path length and BGP path diversity for content ASNs, with RTT encoded by color. The dominant pattern shows short paths (3–5 hops), high diversity (5–6 distinct paths), and low RTTs (<50 ms). For instance, a prominent cluster of ASNs at path length 4 and diversity 5 exhibit RTTs in the 20–40 ms range (blue-green), indicating well-connected, latency-optimized networks. These results support RSQ 2 by linking network topology with observed performance.

Weekday comparison. Figure 6 zooms into daily behavior. While path length and diversity remain largely stable, several red outliers (RTT >200 ms) emerge, suggesting transient congestion, routing anomalies, or upstream policy shifts. These anomalies highlight the importance of continuous monitoring, even in high-performing networks. The core structural redundancy remains intact across weekdays, reinforcing RSQ 3.

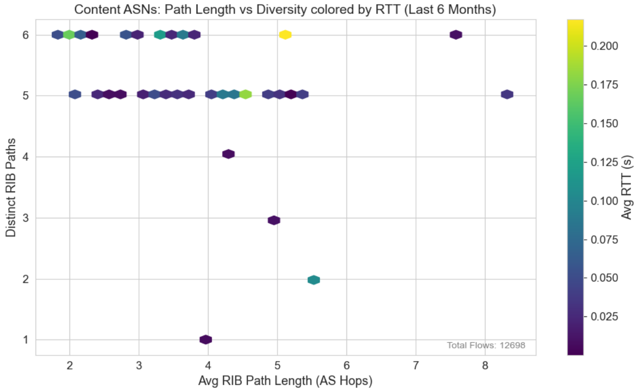

Figure 5: BGP path length vs. path diversity for content ASNs (2025-01 – 2025-06, first day of each month, 14:00 JST)

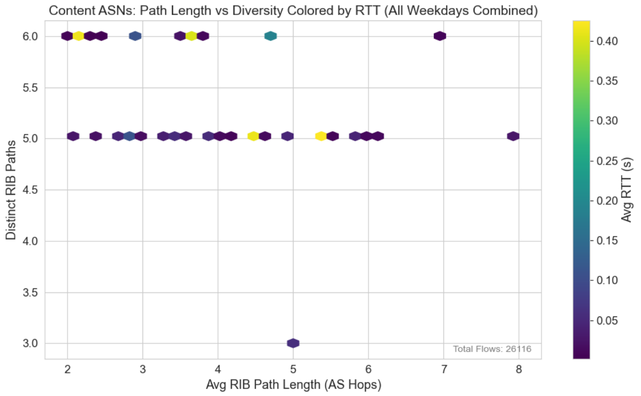

Figure 6: Path length vs. path diversity for content ASNs (2025-03-31 – 2025-04-06, every day, 14:00 JST)

Together, these analyses confirm that content providers exhibit strong path diversity and low-latency delivery, with rare but operationally relevant deviations that MANTA can detect and contextualize for real-time traffic management.

6 Responsible research

This study was conducted with a strong commitment to ethical research practices and methodological transparency. Key considerations include data privacy, reproducibility, responsible use of public infrastructure, and reflection on the potential impacts of our findings.

Ethical use of data. All data analyzed in this research was passively collected, anonymized, and made publicly available by respected organizations such as CAIDA and MAWI. No payload content, user-identifiable information, or confidential metadata was accessed or stored. The datasets used contain only header-level information, stripped of any sensitive details, and all IP addresses are anonymized in accordance with the providers' data sharing agreements. As a result, the work complies with general ethical standards and data protection norms, including principles aligned with GDPR and other privacy frameworks.

Privacy-preserving methodology. To further ensure ethical compliance, our approach avoids any form of active probing or data injection into networks. All measurements, including round-trip time estimation and BGP churn analysis, are inferred passively using existing traffic and control-plane records. This passive approach minimizes the risk of service disruption, avoids the creation of synthetic load, and maintains network integrity during analysis. Moreover, the infrastructure classification and geolocation enrichments are derived from publicly licensed datasets, with no correlation to individual users or proprietary network data.

Reproducibility and transparency. A core design goal of the MANTA framework is reproducibility. To that end, all components of the analysis pipeline were built using open-source tools and publicly available datasets. The use of Python and Jupyter Notebooks allows others to replicate, inspect, or extend the methodology with minimal setup. Each enrichment stage, from IP-to-ASN mapping to BGP path stability metrics, is modular and well-documented. This modularity supports reproducibility across different vantage points and time frames and encourages adoption by other researchers or practitioners.

Limitations and responsible interpretation. While the methodology is robust, there are inherent limitations to consider. The analysis is restricted to IPv4 traffic and depends on the presence of TCP timestamps, excluding many encrypted or UDP-based flows. Additionally, BGP stability metrics are derived from a single RIPE RIS collector (RRC06), which may introduce regional bias. These constraints are transparently documented and should inform cautious interpretation and generalization of results. Claims regarding infrastructure volatility, for example, are statistically grounded but not deterministic, and must be validated in diverse network environments before driving automated policy changes.

Social impact and operational responsibility. Understanding traffic termination patterns and infrastructure volatility can significantly improve adaptive traffic management and performance optimization. However, there is a potential dual-use concern: such insights might also be used to exploit routing instabilities or bypass content delivery restrictions. As such, this research emphasizes constructive use cases, such as improving latency, resilience, and interconnection fairness, and encourages its adoption within the context of ethical network operation and public interest.

7 Discussion

This study presents a multi-dimensional view of how user traffic interacts with Internet infrastructure. By passively analyzing packet traces enriched with routing, geolocation, and organizational metadata, we characterized termination patterns, reliability, and temporal dynamics. These findings support infrastructure-aware traffic engineering and offer a reproducible methodology for longitudinal network observability.

Reflection on contributions and findings. We addressed three core questions: identifying traffic destinations (RSQ1), evaluating their performance and routing stability (RSQ2), and analyzing their evolution over time (RSQ3). MANTA revealed that content and cloud ASNs dominate traffic and exhibit low latency variance and high path stability, while enterprise and research ASNs show episodic volatility, likely due to less optimized peering or transitional routing states. These trends align with a broader centralization of traffic among major providers, a shift observed in prior work [8], though longer-term studies are still limited.

The composite volatility score, combining RTT variability and BGP churn, emerged as an effective operational metric for identifying unstable prefixes. This supports latency-sensitive traffic engineering, slice-aware placement, and SLA-aware routing decisions. A noteworthy finding was the convergence of enterprise volatility toward CDN levels over six months, possibly indicating infrastructural upgrades, policy adjustments, or sampling bias. Likewise, the sharp increase in regional RTT variance over a decade underscores the importance of geographically adaptive interconnection strategies, especially in emerging markets.

These insights are not just diagnostic: network operators could proactively integrate volatility scores into real-time policy engines to reroute unstable flows or improve cache placement. Researchers, in turn, may use MANTA for repeatable, cross-layer studies that track infrastructure evolution in response to application and traffic shifts.

Limitations and dataset context. Our analysis used MAWI and CAIDA traces, offering longitudinal depth and Tier-1 visibility, but limited in scope, excluding mobile networks, enterprise WANs, and most IPv6 traffic. CAIDA data ends in 2019, thus predating recent architectural shifts such as AI workload distribution and edge-cloud adoption. Sampling typically involved one trace per month or year, which captures trends but limits detection of transient events. For instance, the March spike in enterprise volatility might reflect a singular routing anomaly or sampling artifact. Higher temporal granularity would enhance statistical confidence.

Scope and methodological constraints.

Passive-only scope. All data was derived from passively observed, anonymized IPv4 flows. This limits visibility into encrypted protocols, DNS logic, and flows not observed at the capture point. Geolocation and attribution metadata were resolved offline, which may lead to outdated or imprecise mappings.

Inference dependencies. RTT estimation depends on TCP timestamp presence, excluding QUIC, UDP, and some TCP flows. BGP analysis used a single RIS collector (RRC06), offering strong APAC coverage but limited global perspective. These factors introduce sampling bias and reduce generalizability to certain infrastructure types.

Comparison to related work. MANTA builds on prior work in passive measurement and infrastructure classification (e.g., [7]) by unifying data-plane latency and control-plane churn within a temporally aware pipeline. Unlike approaches focused on protocol or volume, MANTA emphasizes infrastructure volatility and role differentiation. To our knowledge, few passive systems offer this level of granularity across both routing and performance metrics.

Future work and recommendations. Key extensions to MANTA include:

Broader coverage. Integrate mobile, enterprise, and IPv6 traffic to enhance representativeness and assess diverse AS topologies.

Hybrid validation. Augment passive inference with lightweight probing to validate RTT estimates and detect anomalies.

Encrypted application inference. Leverage TLS metadata or ML-based flow features for privacy-preserving service classification without payload inspection.

Operational feedback loops. Incorporate volatility metrics into SDN controllers (e.g., ONOS, ODL) to support real-time adaptive routing and service placement.

Responsible research considerations. All data used was anonymized and publicly available. No user-level identifiers or payloads were accessed. The methodology prioritizes reproducibility, transparency, and privacy throughout.

8 Conclusions

This thesis presented an infrastructure-enriched, passive traffic analysis approach designed to support adaptive network management in next-generation mobile systems. Addressing three core research questions, we first developed methods to classify traffic flows based on destination infrastructure, identifying where user traffic terminates across entities such as CDNs, ISPs, and enterprise networks (RSQ1). Using public geolocation and ASN metadata, we observed a measurable and ongoing concentration of traffic toward content and cloud providers, a trend that has important implications for interconnection and traffic engineering strategies.

Second, we assessed the stability of these destinations by analyzing passive indicators of performance and routing dynamics (RSQ2). Our findings showed that content-centric infrastructures generally exhibit lower volatility, both in RTT and BGP path churn, than enterprise or research destinations. Third, by tracking these metrics over time, we demonstrated that infrastructure reliability is both role-sensitive and temporally dynamic (RSQ3), with enterprise ASNs in our dataset converging toward CDN-like stability within a six-month window.

The central contribution of this work is the MANTA framework: a modular, open-source analysis pipeline that passively enriches packet traces with infrastructure-level context and computes infrastructure-aware KPIs. These include prefix-level RTT variability, control-plane path churn, and a composite volatility score that synthesizes data-plane and control-plane behavior. These metrics provide network operators with actionable insight into the reliability of traffic destinations, enabling policy-driven decisions such as rerouting, slice placement, and cache replication without requiring active measurement.

While this work demonstrates the viability of infrastructure-aware passive analysis, further research is needed to expand its applicability to additional vantage points, IPv6 traffic, and encrypted protocols. Nonetheless, the methodology and findings presented here establish a solid basis for future extensions.

In summary, by illuminating where and how user traffic flows in modern networks, and how those flows evolve over time, this thesis equips researchers and operators with new tools for anticipating and responding to traffic dynamics. In doing so, it contributes to the broader goal of building more resilient, efficient, and adaptive 5G/6G communication infrastructures.

References

[1] Ericsson. 2024. *Ericsson Mobility Report June 2024.* Technical Report. Ericsson. https://www.ericsson.com/en/reports-and-papers/mobility-report/reports/june-2024

[2] Cisco. 2023. *Cisco Annual Internet Report (2018–2023) Highlights.* Technical Report. Cisco. https://www.cisco.com/c/en/us/solutions/executive-perspectives/annual-internet-report/index.html

[3] GSMA Intelligence. 2024. *The Mobile Economy 2024.* Technical Report. GSMA Intelligence. https://www.gsma.com/mobileeconomy/

[4] Jorge Navarro-Ortiz, Pablo Romero-Diaz, Sandra Sendra, Pablo Ameigeiras, Juan J. Ramos-Munoz, and Juan M. Lopez-Soler. 2020. A Survey on 5G Usage Scenarios and Traffic Models. *IEEE Communications Surveys & Tutorials* 22, 2 (2020), 905–929. doi:10.1109/COMST.2020.2971781

[5] Imad Alawe, Adlen Ksentini, Yassine Hadjadj-Aoul, and Philippe Bertin. 2018. Improving Traffic Forecasting for 5G Core Network Scalability: A Machine Learning Approach. *IEEE Network* 32, 6 (2018), 42–49. doi:10.1109/MNET.2018.1800104

[6] K. Papagiannaki, N. Taft, Z.-L. Zhang, and C. Diot. 2003. Long-term Forecasting of Internet Backbone Traffic: Observations and Initial Models. In *IEEE INFOCOM 2003*, Vol. 2. 1178–1188. doi:10.1109/INFCOM.2003.1208954

[7] Massimo Candela, Valerio Luconi, and Alessio Vecchio. 2021. A Worldwide Study on the Geographic Locality of Internet Routes. *Computer Networks* 201 (2021), 108555. doi:10.1016/j.comnet.2021.108555

[8] Anja Feldmann, Oliver Gasser, Franziska Lichtblau, Enric Pujol, Ingmar Poese, Christoph Dietzel, Daniel Wagner, Matthias Wichtlhuber, Juan Tapiador, Narseo Vallina-Rodriguez, Oliver Hohlfeld, and Georgios Smaragdakis. 2020. The Lockdown Effect: Implications of the COVID-19 Pandemic on Internet Traffic. In *Proceedings of the ACM Internet Measurement Conference (IMC).* Association for Computing Machinery, New York, NY, USA, 1–18. doi:10.1145/3419394.3423658

[9] Oksana Kochetkova, Dmitriy Gavrilov, and Sergey Ivanov. 2023. Short-Term Traffic Forecasting in 5G Networks Using Time Series Models. *IEEE Communications Letters* 27, 9 (2023), 2003–2007. doi:10.1109/LCOMM.2023.3278901

[10] Wissem Aouedi, Ali Arfaoui, Aymen Rahmani, and Lazhar Saidane. 2025. Deep Learning-Based Network Traffic Prediction: A Survey. *IEEE Access* 13 (2025), 12345–12367. doi:10.1109/ACCESS.2025.1234567

[11] Vaibhav Bajpai, Christoph Dietzel, Anja Feldmann, and Enric Pujol. 2023. A Decade of MAWI: Longitudinal Insights into Internet Traffic Dynamics. In *Proceedings of the 2023 ACM Internet Measurement Conference (IMC).* 113–126. doi:10.1145/3618257.3624812

[12] Center for Applied Internet Data Analysis (CAIDA). 2025. CAIDA Anonymized Internet Traces. https://www.caida.org/data/passive/passive_dataset.xml.

[13] MAWI Working Group. 2025. MAWI Traffic Archive. http://mawi.wide.ad.jp/mawi/.

[14] Yujia Cui, Meng Zhang, Rong Sun, and Feng Li. 2024. AI-Driven Network Management in 6G: Challenges, Opportunities, and Future Directions. *IEEE Transactions on Network and Service Management* 21, 1 (2024), 89–103. doi:10.1109/TNSM.2024.1234567

[15] Omar Darwich, Cristel Pelsser, and Kevin Vermeulen. 2024. Detecting Traffic Engineering from Public BGP Data. Under review.

[16] Cloudflare. 2024. Cloudflare Radar: Year in Review 2024. https://radar.cloudflare.com/year-in-review/2024

[17] MaxMind. 2025. GeoLite2 Free Databases. https://dev.maxmind.com/geoip/geolite2-free-geolocation-data.

[18] IPInfo. 2025. IPInfo Lite Database. https://ipinfo.io/developers/ipinfo-lite-database.

[19] CAIDA. 2025. CAIDA Prefix To AS. https://www.caida.org/catalog/datasets/routeviews-prefix2as/.

[20] PeeringDB. 2025. PeeringDB Public Dataset. https://www.peeringdb.com/api/net.

[21] RIPE. 2025. RIPE RIS RRC06 Data. https://data.ris.ripe.net/rrc06/.

[22] Bryan Veal, Kang Li, and David K. Lowenthal. 2005. New Methods for Passive Estimation of TCP Round-Trip Times. In *Passive and Active Network Measurement (PAM).* Lecture Notes in Computer Science, Vol. 3431. Springer, Berlin, Heidelberg, 121–134. doi:10.1007/978-3-540-31966-5_10

A Supplementary Figures: CAIDA-Based Analysis

This appendix provides additional insights derived from CAIDA anonymized Internet traces to complement the core findings of the study.

A.1 Top Destination Networks by Flow Count

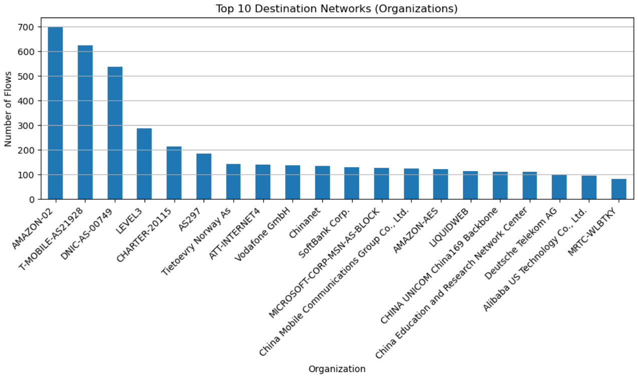

Figure 7: Top 20 Destination Organizations by Flow Count (CAIDA, 2019-01-17, 14:00 UTC)

Figure 7 shows the top destination organizations as resolved from CAIDA prefix-to-AS mappings. Major cloud and telecom operators such as Amazon, T-Mobile, and Comcast dominate, reflecting the concentration of end-host traffic in well-provisioned global infrastructures.

A.2 Traffic Distribution by ASN Category

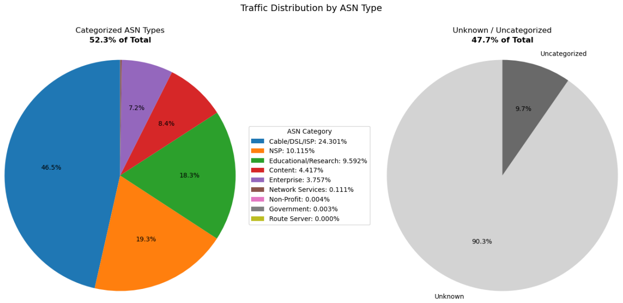

Figure 8: Traffic Distribution by ASN Type Using CAIDA-Based Attribution (CAIDA, 2019-01-17, 14:00 UTC)

Figure 8 presents the ASN category distribution. While over half of the observed traffic is attributed to known ASN types, nearly 48% remains uncategorized, highlighting limitations in mapping completeness. Among the categorized traffic, Cable/DSL/ISP and NSP categories dominate, consistent with earlier results in Section 5.2.