SoK: Mapping the Privacy Landscape of Geolocation Ecosystems

Augustin Laouar[1], Paul Lachat[2,3], Loïc Desgeorges[1], Mathieu Cunche[2,3], Vincent Roca[2], Pascale Vicat-Blanc[2], Francesco Bronzino[1,4]

[1]ENS de Lyon, CNRS, UCBL1, LIP [2]Inria [3]INSA-Lyon, Univ. of Lyon, CITI Lab [4]Institut Universitaire de France

Abstract

Modern geolocation ecosystems rely on diverse technical solutions and architectures, often proprietary, making it difficult to develop a global understanding of the privacy implications of location data production. This challenge is particularly critical given the ubiquity of geolocation in modern digital infrastructures and the central role of location data in privacy concerns. Yet, existing work largely focuses on isolated case studies, resulting in a fragmented understanding of the privacy risks associated with location data production. In this work, we introduce an abstract model of geolocation ecosystems together with a systematic methodology for analyzing their privacy implications, which we apply to nine representative case studies spanning a broad range of architectures, from OS-level geolocation services to object-tracking platforms. This comparative analysis identifies structural design choices that significantly impact users' privacy and reveals common structural privacy risks across heterogeneous ecosystems, which reflect architectural design decisions rather than security vulnerabilities or poor system design. These findings, together with gaps identified in existing defense mechanisms, motivate research directions aimed at strengthening privacy in future geolocation architectures.

Keywords

Geolocation, Location Data Privacy, Systematization of Knowledge

1 Introduction

Geolocation, the act of determining a device or user's physical location, has evolved from isolated positioning mechanisms to complex ecosystems of interacting entities that collectively produce and exploit location data. It now supports a wide spectrum of use cases, including navigation, asset tracking, and safety features. As a result, geolocation ecosystems differ across use cases, with architectures that significantly vary depending on the supported application.

Location data is widely recognized as highly sensitive [52, 58, 95], and its privacy implications have been extensively studied in prior work [67, 74, 96, 128, 129, 152, 156, 168]. However, far less attention has been devoted to the geolocation ecosystems that produce such data, including their architectures and actor interactions. Existing studies that examine such ecosystems typically focus on a single case study, for example [22, 85, 139, 163]. These works rarely analyze the ecosystem in its entirety, leading to a fragmented and non-systematic understanding. Yet these ecosystems, often operated by private companies at massive scale, are involved in the production of billions of location data points, making them critical structural components in privacy risk.

This paper addresses this gap by providing a systematic privacy analysis of a broad range of modern geolocation ecosystems, covering a wide spectrum of use cases, from OS-level geolocation services to object-tracking platforms and IP geolocation. Through this analysis, we aim to provide a unified understanding of how these complex ecosystems operate, highlighting their core architectural patterns and structural sources of privacy risks.

However, systematizing geolocation ecosystems raises two main challenges. First, their high heterogeneity makes it difficult to conduct a comprehensive analysis spanning a wide variety of architectures and use cases. To address this, we develop an abstract model that captures diverse ecosystems under a common representation, enabling transversal analysis. Second, these ecosystems are complex and often proprietary, limiting purely empirical investigation. We therefore adopt a systemic privacy analysis grounded in dataflow modeling, informed by prior work, official documentation, open-source code, and our own reverse-engineering analysis.

Contributions and Outline. We aim to provide a comprehensive assessment of privacy risks in modern location data production ecosystems and identify open research challenges in this space. Concretely, we make the following contributions:

- We develop a unified abstract model of geolocation ecosystems, capturing their structure, actors, and data flows, and enabling the identification of shared architectural patterns and cross-ecosystem privacy insights (§ 2).
- We propose a privacy analysis framework built on our abstract model and the LINDDUN[1] threat modeling framework [60] (§ 3).
- We provide a systematic characterization of nine representative modern ecosystems, grouped into three categories and spanning a wide spectrum of use cases, each introduced at the beginning of its respective case study section (§ 4, § 5, and § 6).
- We conduct a comparative privacy analysis of these ecosystems and identify recurring design patterns and structural sources of privacy risk. Our analysis shows that the core issues arise from architectural centralization, inherent properties of communication protocols, and cross-ecosystem inference effects (§ 7).
- We review existing defense mechanisms applicable to these ecosystems, identify structural gaps that current approaches fail to address, and derive concrete research directions to guide future work (§ 7).

2 Conceptual Model

Geolocation encompasses a broad set of mechanisms whose implementations differ substantially and often involve multiple actors. While these approaches appear highly diverse, they share common

[1]LINDDUN stands for Linking, Identifying, Non-repudiation, Detecting, Data Disclosure, Unawareness, and Non-compliance.

Proceedings on Privacy Enhancing Technologies 2026(4), 455–479

https://doi.org/10.56553/popets-2026-0130

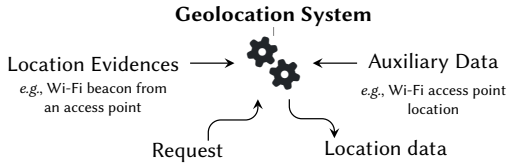

Figure 1: Illustration of a Geolocation System

underlying principles. This section introduces a conceptual model that captures these principles at an abstract level by defining the key entities involved, their functional roles, and the relationships and data flows that connect them. By reducing existing approaches to a set of core notions, relationships, and orthogonal distinctions, the model provides a unified framework for understanding the structure and functioning of geolocation processes across different technologies and architectures.

2.1 Data taxonomy

Systems that process spatial information manipulate different forms of data. For clarity, we distinguish four related notions:

Position. A position is a quantitative estimate of an entity's spatial state, typically expressed as coordinates within a defined reference system (*e.g.*, latitude and longitude). Unless stated otherwise, we use *position* by default to refer to an entity's spatial state.

Location. A location is a semantic description of an entity's position (*e.g.*, street address, city name).

Location Data. Location data denotes structured information representing either a position or a location (*e.g.*, a location record), potentially enriched with contextual metadata.

Location Evidence. Location evidence denotes any signal or observation that enables estimating a spatial relationship between an entity and another (*e.g.*, a GNSS signal).

2.2 Geolocation Techniques

Geolocation techniques are the methods by which an entity's spatial state is estimated. They rely on various forms of location evidences and auxiliary data to compute location data. We characterise these techniques along three orthogonal dimensions:

Active vs. Passive. The difference lies in the target's participation in the geolocation process. Geolocation is *active* when the target deliberately engages in the process, and *passive* when its position is estimated without its involvement or awareness.

Ego-geolocation vs. Allo-geolocation. This distinction is based on which entity performs the position estimation. In *ego-geolocation* approaches, the target estimates its own position by observing its environment. In *allo-geolocation* approaches, the target's position is inferred by other entities based on their observations of the target.

Autonomous vs. Assisted. Geolocation is *autonomous* when the location evidences contain all the necessary information to compute a position. It is *assisted* when the evidences must be mapped against databases to determine a position.

2.3 Geolocation System

A geolocation system, illustrated in Figure 1, constitutes an abstraction layer built on top of geolocation techniques. It orchestrates and combines these techniques to produce location data and provides standardized interfaces to expose them. Despite architectural variations, most geolocation systems offer a common set of functional capabilities to support these processes:

Evidence collection. Gathering location evidences either directly by the host entity or via other entities in the ecosystem.

Processing. Computing location data from evidences, optionally combined with auxiliary data.

Storage. Persistently storing evidences, auxiliary data, or computed location data for further use.

Access Control. Restricting access to location data according to privacy and security policies.

2.4 Geolocation Ecosystem

A geolocation system operates within a broader set of entities involved in the geolocation process. Together, these entities form what we refer to as a geolocation ecosystem, defined as a distributed set of entities that contribute to the production, processing, and consumption of location data.

At a high level, our model captures the structure of geolocation ecosystems by identifying the roles played by the different entities involved in the geolocation process. Each entity may assume one or more roles and may operate its own geolocation system. Our abstract model focuses on how these roles interact and compose to form end-to-end data flows across the ecosystem.

We define functional roles of entities as follows:

Target. A technical entity whose position is being determined (*e.g.*, a smartphone or an IoT device). The target's position is generally used to infer that of another entity (*e.g.*, a person or an object).

Consumer. An entity that receives and utilizes location data to perform an action, provide a service, or conduct further analysis (*e.g.*, a location-based service, an emergency response system).

Service Provider. An entity that provides location data to consumers. It may act as a relay, collecting and exposing location data produced by other entities, directly compute location data from location evidences, or store and manage access to location data.

Reference Point. A physical entity with a known position that serves as a spatial anchor for geolocation (*e.g.*, GNSS satellites, Wi-Fi access points). A reference point may be identified by a unique identifier, which we call *Reference Point Identifier* (*e.g.*, the *basic service set identifier (BSSID)* for a Wi-Fi access point). The spatial relationship between a target and a reference point is established by the observation of location evidence.

Reference Resolver. A subset of *Service Provider* that supplies dataset of spatial reference information associated with reference point identifiers (*e.g.*, BSSIDs positions databases).

Contributor. An entity (*e.g.*, crowdsourcing device, network probe) that collects location evidence and shares these observations, potentially enriched with contextual data, with another entity involved in the geolocation process.

Figure 2 provides a concrete illustration of how the introduced terminology applies to real-world systems and establishes the color scheme used consistently across the figures of this paper. In this

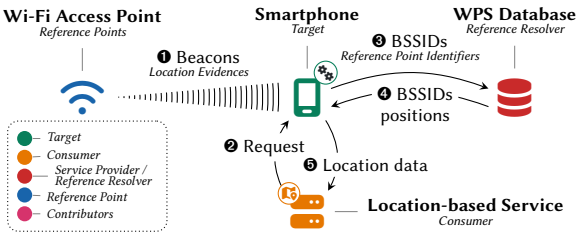

Figure 2: Illustration of the proposed terminology for a device computing its position via a Wi-Fi Positioning System (WPS).

case, a Wi-Fi device computes its position using a Wi-Fi Positioning System (WPS) [104, 106], which corresponds to an *active* and *assisted ego-geolocation* technique.

3 Methodology

To better understand how modern technologies are used for geolocation and their impact on user privacy, we select a representative set of consumer solutions. This section presents these case studies and their categorization, defines our adversarial model, and details the method used to identify the privacy threats arising within each ecosystem.

3.1 Case Studies

We selected case studies based on their prevalence in real-world deployments and the availability of sufficient information for rigorous privacy analysis. We categorize them along the *Active/Passive* and *Ego/Allo* dimensions of the geolocation techniques used, as these shape the target's involvement and data exposure. We exclude the *Autonomous/Assisted* dimension, which reflects an implementation choice rather than a structural distinction. This yields three groups (Table 1), aligned with their operational purpose and trust model:

Active Ego-Geolocation. The target actively determines its own position, typically through a built-in geolocation system, and retains primary control over the resulting location data.

Active Allo-Geolocation. The target actively acts to be localized by an external entity. These are opportunistic solutions where the target is not autonomous and delegates the management of its data to an external entity to enable tracking or recovery.

Passive Allo-Geolocation. External entities localize the target by leveraging signals not intended for geolocation. Here, localization is a side-effect of connectivity used for surveillance or administration, occurring without a specific request or explicit trust contract.

Passive Ego-Geolocation is omitted as these properties are fundamentally incompatible, as an entity cannot determine its own position without active processing or observation.

3.1.1 Active Ego-Geolocation: OS-level Geolocation Systems. In these case studies, we focus our analysis of active ego-geolocation on OS-level geolocation systems. Except for trivial architectures such as standalone microcontrollers, the geolocation system is typically implemented as a core service within the operating system, providing a unified interface to applications and enforcing security boundaries by preventing them from accessing hardware sensors (*e.g.*, GNSS, Wi-Fi, Bluetooth) directly. Consequently, the OS acts as the primary orchestrator of the geolocation ecosystem and the critical point for privacy enforcement.

To ensure the representativeness of our study, we selected the most prevalent operating systems: Android, Windows, and iOS/-macOS, which together account for the majority of the market share [72]. For Android, we specifically focus on devices running *Google Mobile Services (GMS)*, the standard configuration for most commercial smartphones [37]. We also included *GeoClue* [71], the standard geolocation framework for Linux distributions.

These ecosystems have been only partially explored in prior work. For iOS/macOS and Android (GMS), existing research has focused on isolated components of the geolocation process or specific ecosystem features [36, 39, 44, 99, 110, 111, 134, 139, 150]. However, no study has yet provided a comprehensive overview of the privacy impact of these architectures as a whole. Furthermore, to the best of our knowledge, Microsoft Location Services and GeoClue have not been subject to any systematic privacy study.

3.1.2 Active Allo-Geolocation: Offline Finding Networks. For this category, we focus on large-scale *Offline Finding Networks (OFNs)*, as they constitute the most pervasive and privacy-critical instantiations of this paradigm. OFNs leverage a global, dynamic infrastructure of billions of mobile devices and represent a growing market [130] with direct impacts on individuals' privacy. We therefore analyze two prominent examples, Apple Find My [5] and Google Find Hub [7]. These ecosystems offer an *offline finding* feature, which allows locating objects even when they are not connected to the Internet. This specific capability has been the primary focus of recent studies [25, 41, 48, 65, 85].

Other active allo-geolocation systems, such as emergency distress beacons or UWB-based industrial tracking, generally operate within simpler ecosystems and are restricted to specialized use cases. In contrast, OFNs provide the most compelling case for analyzing the systemic risks of modern active allo-geolocation.

3.1.3 Passive Allo-Geolocation: Network Location Services. In modern systems, passive allo-geolocation consists in diverting traditional networks from their original purpose to localize their members. It is considered *passive* because the underlying signals are not primarily emitted for geolocation. We therefore analyze three types of network infrastructures, namely the Internet, cellular networks, and low-power wide-area networks (LPWANs), and select for each a representative real-world geolocation service.

We studied two distinct popular solutions: IPinfo [9] for IP-based geolocation services [73] and 1NCE [4] for cellular-based localization [51]. Additionally, we analyzed Sigfox Atlas [10] as a representative of LPWAN-based systems. This choice was motivated by access to detailed technical information provided by Sigfox.

3.2 Privacy Threat Modeling

Given the diversity of actors involved in geolocation ecosystems, an attacker-driven analysis would be restrictive and potentially incomplete. We therefore favor a systemic approach grounded in LINDDUN [60], a framework for eliciting privacy threats. While most threat modeling frameworks focus primarily on security [141], privacy-oriented alternatives such as PANOPTIC [89] and MAP [62] have emerged; however, LINDDUN remains, to the best of our

Table 1: Classification of Case Studies According to the Conceptual Model

Case Studies	Target involvement	Mode	Supported Geolocation Techniques
OS-level Geolocation Systems			
iOS & macOS	Active	Ego-geolocation	GNSS[1], WPS, Cell-ID Location[1], Inertial Sensors, BLE beacons, UWB[1,2]
Android (Google Mobile Services)	Active	Ego-geolocation	GNSS, WPS, Cell-ID Location, Inertial Sensors, IP-Geolocation[2], UWB[2]
Windows	Active	Ego-geolocation	WPS, IP-Geolocation[2], GNSS, Cell-ID Location, Inertial Sensors
GeoClue (Linux)	Active	Ego-geolocation	IP-Geolocation[2], WPS, GNSS, Cell-ID Location, Inertial Sensors
Offline-Finding Networks			
Apple Find My	Active	Allo-geolocation	BLE Beacon, UWB[2]
Google Find Hub	Active	Allo-geolocation	BLE Beacon, UWB[2]
Network-Based Geolocation Services			
Sigfox Atlas	Passive	Allo-geolocation	Sigfox Atlas
IPinfo	Passive	Allo-geolocation	IP-Geolocation
1NCE	Passive	Allo-geolocation	Cellular Network Geolocation

[1] Only on iOS. [2] While these techniques can be inherently passive, they are considered active here as the target deliberately participates in the process, either by contacting a service (IP) or by engaging in a dedicated ranging protocol (UWB).
Acronyms: GNSS – Global Navigation Satellite System; WPS – Wi-Fi Positioning System; BLE – Bluetooth Low Energy; UWB – Ultra-Wideband.

Table 2: Adversary Model Index

Adversary	Behavior
A1 Target-based	Malicious
A2 Trusted External	Honest-but-curious / Malicious
A3 Untrusted External	Honest-but-curious / Malicious

knowledge, the only top-down approach providing a comprehensive taxonomy and methodology. It is widely adopted in academic research [45, 61] and recommended by regulatory authorities [16, 144].

We classify threats according to the seven LINDDUN types, summarized in Appendix A.

To systematically elicit privacy threats using LINDDUN, we model the studied systems using *Data Flow Diagrams (DFDs)*, representing four standard elements: *external entities*, *data stores*, *processes*, and *data flows*. These DFDs are constructed from scientific literature, official documentation, governmental reports, and our own reverse-engineering analyses. We present one representative DFD per group of case studies in Appendix E. For each data flow, we apply the LINDDUN framework to identify applicable threat types, then assess each through attacks documented in the literature or, in their absence, through our own analysis.

3.3 Adversarial Model

While our threat analysis relies on the LINDDUN methodology and is therefore data-driven rather than attacker-driven, we introduce a classification of adversarial actors to contextualize which entities may exploit the identified privacy threats. This classification provides a structured vocabulary rather than guiding threat elicitation. We distinguish three categories based on the actor's position and trust relationship relative to the target, summarized in Table 2.

A *target-based adversary* (**A1**) is an entity with direct access to the target, such as an installed application, malware, or an actor with physical control over the device. This adversary is inherently malicious and seeks to obtain the target's position for surveillance or commercial purposes, for instance through real-time bidding and the broader data brokerage ecosystem [52].

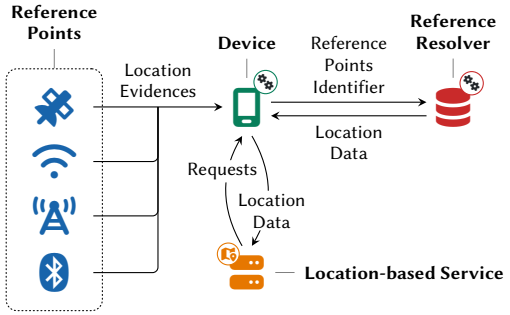

Figure 3: Generic architecture of ecosystem underlying OS-level geolocation systems.

A *trusted external adversary* (**A2**) is an entity that the target explicitly or implicitly trusts, typically through the acceptance of terms of service. This includes platform or service providers, which may behave maliciously or in an honest-but-curious manner. Their goals range from reconstructing mobility patterns to re-identifying anonymized contributions [156]. Prior cases have illustrated their involvement in controversial data practices [43], sharing with governmental actors [24], or commercial exploitation [52, 82, 120].

An *untrusted external adversary* (**A3**) is an entity with no formal trust relationship with the target, yet able to access, infer, or manipulate location-related information. This ranges from external attackers compromising a cloud provider, to web services inferring users' locations through IP-based geolocation without consent.

4 Case Study: OS-Level Geolocation Systems

We first describe the architecture underlying OS-level geolocation ecosystems and then analyze the design choices leading to privacy threats. These mainly stem from the ecosystem provider's ability to correlate location data, from security vulnerabilities exploitable by external entities, and from weaknesses in authorization design.

4.1 Ecosystem Architecture

OS-level geolocation systems are integrated into the operating system of a device to provide ego-centric geolocation services to its

user. These systems run with elevated privileges and act as an abstraction layer between applications and the underlying sensing infrastructure. They manage the entire geolocation pipeline, from permission handling and sensor fusion to position estimation and data delivery. Such systems typically rely on multiple geolocation techniques, which may be autonomous (*e.g.*, GNSS-based) or assisted (*e.g.*, Wi-Fi or cellular-based positioning), and combine these sources to improve accuracy and robustness [11, 12].

Figure 3 provides a general and deliberately simplified view of the typical ecosystem for OS-level geolocation systems. A complete data flow diagram for the iOS ecosystem is provided in Appendix E as a representative example. Such ecosystems involve multiple interacting actors:

Target. The device on which the OS-level geolocation system runs (*e.g.*, a smartphone or laptop).

Reference Resolvers. Targets often rely on multiple assisted geolocation techniques and query reference resolvers that possess large geolocation databases, typically operated by private entities (*e.g.*, Google, Apple, or Microsoft). These reference resolvers often operate their own geolocation systems to derive reference points' positions from contributors' observations [14, 29].

Reference Points. A wide range of communicating devices may serve as reference points (*e.g.*, Wi-Fi Access Points, GNSS satellites), depending on the target's capabilities to observe location evidence.

Contributors. Reference resolvers' clients often contribute to the maintenance and update of reference resolvers' databases by periodically uploading their own observations, typically coupled with a reliable position estimate derived from GNSS. Thus, targets frequently act as contributors as well. Depending on the reference resolver, such contributions may be mandatory or optional. Note that Figure 3 does not represent the crowdsourcing dataflow, which is often incorporated into reference point resolution queries.

Consumers. These are entities authorized to access location data produced by the target. While depicted in Figure 3 as external, consumers may also be services running locally on the target device.

4.2 Reference Point Identifier Resolution

Assisted geolocation techniques require resolving reference point identifiers through a remote reference resolver, as reference databases are large-scale and continuously updated, making local replication impractical. Apple [29] and Google [14] each operate their own proprietary reference resolver, whereas Microsoft Location Services relies on proprietary partners acting as resolvers [119]. In contrast, GeoClue does not impose a fixed backend and allows users to choose their reference resolver [71], or none at all, thereby enabling related privacy threats to be reduced or avoided altogether.

Resolving reference point identifiers reveals the requester's position to the resolver, which can infer it from the submitted identifiers and its knowledge of the corresponding reference points locations. Our experiments on iOS, reported in Appendix B, show that requests to reference resolvers occur frequently, even without any user interaction with the device, enabling fine-grained tracking.

Authentication choices. Authenticated access to crowdsourced reference databases can limit large-scale extraction and abuse. Google Location Services, enabled by default on most Android devices, and selectable as a backend in GeoClue, enforces such account-based authentication. In contrast, Apple and Microsoft location services expose reference point resolution interfaces without strict account-based authentication, for privacy purposes.

In the presence of account-based authentication, the service provider, acting as **A2**, can associate queries with user accounts and link them over time for location tracking. Unauthenticated queries may raise the same threat, as IP addresses can still be used as identifiers and allow linking queries over time [67, 156]. Additionally, Rye et al. [134] demonstrated how this design choice can be exploited by **A3** to expose the location of arbitrary access points at scale in Apple's Wi-Fi positioning service, which may allow inferring mobility patterns of identified devices over time, such as mobile Wi-Fi routers.

Crowdsourcing. As reference resolvers rely on crowdsourcing to maintain accurate reference databases, user devices act both as contributors and reference points. To support this process, devices typically include their GNSS-based position in crowdsourcing queries, and some ecosystems additionally collect contextual data (*e.g.*, velocity, direction, or barometric pressure [11]).

On macOS, iOS, Android with Google Mobile Services, and Windows, users cannot access their ecosystem's reference resolver without participating in crowdsourcing [14, 29], leaving them with no meaningful opt-out. In contrast, GeoClue supports open, community-maintained reference resolvers (*e.g.*, BeaconDB [153]), enabling geolocation without mandatory user contribution, albeit with reduced coverage compared to proprietary databases such as Google's.

Crowdsourcing mechanisms amplify the threats raised by queries to reference resolvers as they enable automatic and frequent data collection, often beyond what is necessary for the geolocation process (*e.g.*, barometric pressure [11]).

4.3 Extended Services

We use the term extended services to denote OS-level functionalities built on top of the core geolocation system, which go beyond the basic computation and delivery of location data.

Location History. Beyond producing and exposing instantaneous position estimates, OS-level geolocation systems may support extended location history services. These services serve different purposes across ecosystems. Apple's Significant Locations [32] records user positions to provide contextual assistance and system-level optimizations. Google Timeline [76] stores historical traces to enable long-term mobility visualization. Windows Find My Device provides device recovery capabilities by periodically transmitting the device's position to Microsoft's servers (Apple and Google provide similar services discussed in § 5.2).

Privacy implications stem from the persistence of location records, whose accumulation enables reconstructing mobility patterns. Significant Locations and Google Timeline store data locally and, depending on user configuration, synchronize it to the cloud. Local storage exposes data to **A1** (*e.g.*, malware, physical access), and trusted execution environments such as the Secure Enclave [31] only partially mitigate this threat [151, 154]. Cloud synchronization places data under the provider's control, exposing it to the provider when acting as **A2**, as well as to **A3** in case of infrastructure compromise. Significant Locations applies *end-to-end encryption (E2EE)*,

mitigating both threats by preventing provider access and reducing exposure in case of server compromise.

Windows Find My Device follows a more constrained design, storing remotely only the last reported position. By avoiding historical accumulation, it structurally reduces longitudinal linking and disclosure risks, assuming the cloud provider behaves honestly.

Enriched Location API. Certain ecosystems expose APIs that provide enriched, location-derived information to applications. Google's Awareness API [77] enables applications to query high-level contextual states inferred from on-device signals, including coarse-grained location contexts such as proximity to known places. Similarly, Windows' CivicAddress API [19] allows applications to retrieve the civic address associated with the device's current position.

In both cases, obtaining enriched information may require transmitting the device's position to a remote service operated by Google or Microsoft. Access to these services is restricted through application-level authentication, typically enforced via API keys and, in some cases, device-specific identifiers.

Such designs allow the service provider, in the role of **A2**, to associate successive queries with a given application and correlate them with quasi-identifiers (*e.g.*, IP addresses).

4.4 Lack of Meaningful Control

Although users are theoretically granted control over location-related permissions, interface design choices may undermine this control in practice. Prior work has shown that some interfaces embed *dark patterns*, *i.e.*, design strategies that steer users toward choices misaligned with their intentions or interests [81, 103]. In the context of geolocation, this may result in users being located without genuine consent or being misled into granting consent under false or incomplete assumptions.

Locating users without explicit consent. Control over location-related permissions may be missing or unclear, allowing applications to infer a user's location without explicit authorization. Prior work has shown that inertial sensors can be leveraged to compute accurate indoor locations without user awareness [59, 123, 163]. Similar shortcomings affect Bluetooth permissions. Toubiana et al. [150] demonstrated that, on Android, location and active location permissions were not properly enforced, enabling unauthorized *Bluetooth Low Energy (BLE)* batch scanning despite users' expectations. On iOS, Schmidt et al. [139] showed that the local network permission can be exploited to infer a user's position by collecting identifiers of nearby Wi-Fi access points. Collectively, these practices exploit users' lack of awareness about ongoing data flows, enabling **A1** to access location information without informed consent.

Misleading users into granting location permissions. Even when permissions formally exist, interface designs may fail to clearly communicate their implications. Applications may circumvent users' expectations around consent. Beer et al. [39] demonstrated that, on Android, an application can use low-opacity animations to trick users into unknowingly interacting with a secondary application, such as granting it location permission.

Platform-level opacity and policy deviations. Platform operators such as Google, Apple, and Microsoft, through their role as service providers (*e.g.*, reference resolvers or cloud providers), are able to access users' location data, as previously discussed. Access to such data is governed by privacy policies and terms of service. However, due to their complexity and insufficient transparency regarding secondary uses of location data [23, 115, 131, 158], users may consent without fully understanding how their data are processed. More critically, deviations from declared policies and undisclosed data-sharing practices have been documented in prior cases [54], questioning the trust users can place in platform operators.

4.5 External Location Evidence Tampering

Assisted geolocation techniques rely on broadcast identifiers (*e.g.*, Wi-Fi BSSIDs or cellular Cell-IDs) resolved against external databases. As these identifiers are typically neither confidential nor authenticated, location evidence can be manipulated by an **A3**, causing the system to return incorrect positions.

Wi-Fi positioning systems are vulnerable to BSSID reuse and manipulation [112, 148]. Cellular system information block type 1 signals [17] can be spoofed to advertise false Cell-IDs [36, 143], and BLE-based systems are particularly exposed due to static beacon identifiers [146]. Although GNSS is autonomous and does not rely on external resolvers, GNSS signals themselves remain susceptible to spoofing under certain conditions [26, 117].

By spoofing identifiers or broadcast signals, **A3** can cause incorrect but credible location records, undermining the integrity of attributed location data. This risk is particularly sensitive when location data are used as legal or forensic evidence [138], as manipulated records may lead to false attribution or wrongful suspicion. Moreover, tampering with location evidence may refine the identity of a device's owner [112].

4.6 Takeaways

T4.1: Service Provider Inference Power. The primary privacy risk is structural, as a central entity, the service provider, accumulates disproportionate inference power by operating multiple services within the ecosystem (*e.g.*, reference point resolution, cloud services). This functional concentration grants repeated access to location data, enabling longitudinal and cross-context inference. Past cases [54] illustrate that such trust assumptions may not always align with actual data practices. More distributed deployments such as those enabled by GeoClue can mitigate these privacy risks.

T4.2: Limited Leverage For Non-Central Adversaries. In contrast to structural inference risks, other adversaries face limited leverage absent security flaws. They may bypass consent mechanisms to obtain location data or manipulate location evidence to produce false yet credible records; broader access to stored location data typically requires exploiting a security vulnerability. These threats are thus primarily mitigated through technical hardening.

5 Case Study: Offline Finding Networks (OFNs)

Offline Finding Networks [101] have rapidly become widely deployed geolocation infrastructures, with hundreds of millions of devices participating in ecosystems such as Apple Find My and Google Find Hub. This section analyzes their operation and highlights how service providers, due to their control over the underlying technologies, retain significant inference capabilities despite

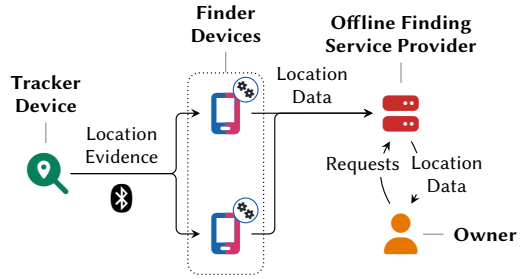

Figure 4: Generic architecture of Offline Finding Networks

strong technical privacy guarantees. We also examine vulnerabilities and architectural weaknesses that can be exploited by malicious actors.

Although legitimate OFN tracking devices can be misused by individuals to track other people (*e.g.*, for stalking) [48, 85, 114, 157], such downstream misuses by end users lie outside the primary scope of this paper, which focuses on the privacy implications of the OFN ecosystem itself.

5.1 Ecosystem Architecture

OFNs are large-scale, crowdsourced systems designed to locate offline devices, typically lost devices, by leveraging the proximity of third-party participants. They are engineered to operate with minimal energy consumption. To achieve this objective, offline devices periodically broadcast short-lived identifiers, typically via BLE, which are detected by nearby participant devices. These *finder devices* relay the observed identifiers together with their computed position (using OS-level geolocation systems; see § 4) to a central service, allowing the device owner to retrieve the corresponding location reports. In the OFNs studied in this work, location reports are E2EE using cryptographic material derived from the broadcasted BLE identifiers.

Figure 4 is a high view of how OFNs work. As a representative example, a complete data flow diagram for the Find My ecosystem is provided in Appendix E. While the architectures of these ecosystems are similar, they differ in the way their protocols are implemented. Thus, the overview provided by this figure is valid for other networks such as Samsung SmartThings Find [15]. Multiple interacting actors are involved:

Target. In OFNs, the target is referred to as a *tracker device*. This category includes both dedicated tracking devices (*e.g.*, AirTags) and personal devices that have lost Internet connectivity (*e.g.*, a lost smartphone). Although these devices may be equipped with a geolocation system, without an Internet connection, this capability cannot be leveraged for positioning purposes. Thus, we consider them as not having one.

Reference Points / Contributors. The *finder devices* are connected to the Internet and can upload reports when they detect tracker device advertisements to service provider servers. Each report includes the finder device's current position, encrypted using the cryptographic secret from the detected BLE advertisement in the case of E2EE, along with additional metadata such as timestamps and device attestation data.

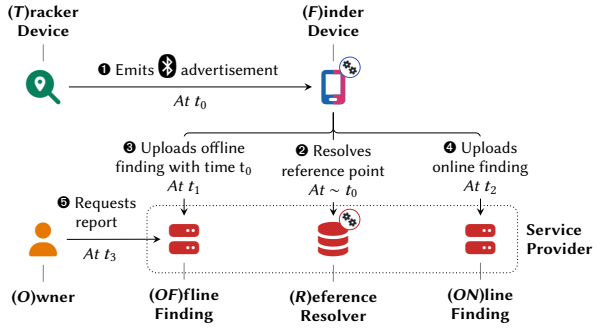

Figure 5: Service Provider monopoly over different services enables inference to bypass reports E2EE

Service Provider. The *offline finding service provider* collects reports from finder devices and makes them available to the owners of a tracker. In OFNs studied in this work, the reported positions are protected by E2EE, such that the service provider cannot access the plaintext of the reported location data.

Consumer. The *owner* of a tracker device accesses its associated location reports through a dedicated application running on their device (*e.g.*, the Find My application on iOS) that fetches reports from the service provider. In case of E2EE, the owner device deciphers these reports to recover the last known position of the tracker device, using the encryption keys established during the pairing process with this tracker.

5.2 Online Finding

Apple's and Google's offline finding networks are widely advertised as being E2EE. However, enabling offline finding on devices with Internet connectivity, such as smartphones, also activates an associated service referred to as *online finding*. It periodically reports the device's location directly to the service provider when the device is connected to the Internet, using an authenticated channel tied to the user's iCloud or Google account. In contrast to offline finding reports, these location updates are not E2EE and are therefore accessible to the service provider [6, 13]. We provide an example of such an online finding request in Appendix C.

To the best of our knowledge, on both Apple Find My and Google Find Hub, offline finding cannot be enabled without also activating the associated online finding service. Users are therefore unable to benefit from the stronger privacy guarantees of offline finding in isolation. Furthermore, because online location reports are stored remotely and tied to a user account, they expose users to threats from **A2** similar to those discussed for location history in § 4.3.

5.3 Service Provider Inference

Even if offline finding network providers (*i.e.*, Apple and Google) claim that they have no access to location data reported through their OFN [6, 13], several forms of inference remain possible. By correlating metadata and information obtained across multiple services they operate, the service provider can infer sensitive information such as reported locations or social relationships between users.

Bypassing reports' E2EE. Despite the use of E2EE for offline finding reports, service providers may still infer the location associated

with a report. This stems from the fact that the same entity operates all services involved in the reporting process on the finder device. To the best of our knowledge, this inference channel has not been explicitly documented or analyzed in the existing literature.

Figure 5 illustrates the sequence of interactions that follows the detection of an offline finding advertisement and highlights the data that can be observed and correlated by the service provider.

❶ **Advertisement.** A tracker device (T) broadcasts a BLE advertisement that is received by a finder device (F) at time t_0.
❷ **Reference point resolution.** To compute its position, F invokes the OS-level geolocation service immediately after the detection. As shown in § 4.2, these geolocation systems rely on assisted positioning techniques and query a reference point resolver (R) operated by the same service provider. This interaction reveals the position of F, together with its IP address and the timestamp of the query to R.
❸ **Offline report upload.** Once its position is computed, F uploads an offline finding report to the offline finding service (OF). While the reported location is protected by end-to-end encryption, metadata such as the finder's IP address and the report timestamp remain visible to OF.
❹ **Online finding upload.** When online finding is enabled on F, the device periodically reports its own location to the online finding service (ON). We denote by t_2 the time of one such upload. These reports are authenticated and can be directly associated with the finder's account, together with its IP address, the report timestamp, and reported location, as described in § 5.2.
❺ **Report download.** To locate a tracker, its owner (O) retrieves offline finding reports from OF through an authenticated request [41, 85]. This allows OF to associate reports with both an owner account and a specific tracker device.

By the end of this process, OF can associate an offline report with the owner's account, the finder's IP address, and the report timestamp. Simultaneously, R associates an IP address and timestamp with the finder's position, while ON may link an IP address and timestamp to both a position and a user account. As discussed in § 4.2, requests to R occur frequently, making linking based on IP address and timestamp highly effective. The correlation is further strengthened when the online finding upload at time t_2 is temporally close to the offline report at time t_0. Because all services are operated by the same entity, the provider is able to combine this information to infer the finder's position at advertisement detection time and, potentially, their identity. Most importantly, given the limited BLE range, this also reveals the position of the tracker itself, despite Apple's and Google's claims that this information cannot be retrieved.

Location inference from report metadata. Even without correlating data across multiple services, offline finding systems leak information through report metadata alone. Böttger et al. [41] show that, in Google Find Hub, location reports are directly associated with the device owner's Google account. Unencrypted metadata such as report frequency, timestamps, and advertised accuracy can therefore be leveraged to infer coarse-grained location information and user activity patterns (*e.g.*, being at home versus moving in the city). Over extended periods, such inferences enable building longitudinal profiles of users' routines and behavioral habits.

In Apple Find My, reports are not linked to the owner's account at upload time, but this link is established when the owner authenticates to retrieve them. Although location accuracy is protected by E2EE, metadata such as report frequency and timing remain visible to Apple. As a result, similar profiling threats arise, albeit mitigated by Apple's design choices compared to Google's.

Social Graphs. Prior work on Google Find Hub [41] and Apple Find My [85] has shown that both OFNs enable service providers to construct social graphs by linking location reports over time. Reports from the same finder can be correlated using authentication data (*e.g.*, device identity certificates) or IP addresses. When a finder uploads reports for multiple trackers within a short time window, the provider can infer their co-location. Likewise, owners authenticate to retrieve reports, associating downloaded reports with specific user accounts and allowing the provider to infer which users own trackers that were co-located at the same time.

These correlations enable **A2** to construct co-location graphs of tracker owners. The inference mechanisms described earlier further amplify this risk by allowing the geographic position of nodes within these graphs to be approximated.

5.4 Unauthorized Access to Reports

Beyond structural inference risks, specific design choices and implementation vulnerabilities may enable unauthorized access to location reports.

Location sharing. Google Find Hub permits users to share the location of their devices with other users. When sharing a tracker, the owner and the invited user establish a shared secret that enables access to the tracker's location reports. This secret is derived through a key exchange between the two parties, while the service provider relays the exchange messages [41].

As a result, Google could perform a man-in-the-middle attack by replacing the invitee's public key with its own during the key exchange, thereby deriving the shared secret and accessing location reports [41]. Additionally, a security vulnerability allows an attacker who has lost access to a device's location sharing to continue receiving reports under certain conditions [41]. On Apple's side, the location-sharing protocol is not publicly documented, preventing firm conclusions about similar threats.

These weaknesses enable **A2**, or **A3** exploiting the disclosed vulnerability, to obtain unauthorized access to shared location reports, while undermining users' ability to monitor and revoke access.

Local Storage of Cryptographic Secrets. An attacker who gains access to cryptographic secrets, held by the owner, can use them to retrieve and decipher the device's location reports, without requiring access to the corresponding account. For instance, Heinrich et al. [85] identified a vulnerability that allowed a macOS application, with user-granted permissions, to access such secrets and thereby obtain unauthorized access to location reports; this issue has since been patched by Apple. In Google Find Hub, exploiting a comparable threat would additionally require access to the owner's Google account, raising the attack complexity.

A3 can exploit this vulnerability to obtain unauthorized access to location report contents. Access to multiple reports may further enable the attacker to correlate location updates over time and reconstruct mobility patterns.

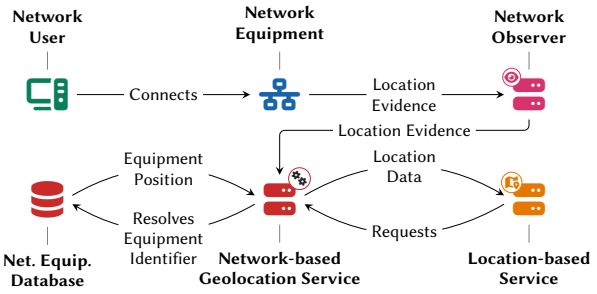

Figure 6: Generic Ecosystem of network-based geolocation services

5.5 Takeaways

T5.1: OFN Providers Hold Structural Inference Power. Even when designed with privacy mechanisms such as end-to-end encryption, the proprietary OFNs studied in this work still allow the service provider to retain significant inference capabilities due to its central, multi-role position in the ecosystem. End-to-end encryption alone is therefore insufficient to prevent such inference.

T5.2: Limited Leverage For External Adversaries. Robust cryptographic protocols and end-to-end encryption significantly limit the capabilities of external adversaries. Remaining threats are rare and typically depend on specific security vulnerabilities rather than structural design flaws.

6 Case Study: Network-based Geolocation Services

Network-based geolocation refers to the class of geolocation techniques that locate a target by observing how it attaches to a communication network. Network-based geolocation services provide access to these techniques, enabling their customers to obtain location data about targeted network users. We first characterise how entities interact within such ecosystems, and then analyze the privacy threats arising from their deployment, which primarily stem from the structural properties of the underlying network rather than from specific implementation choices.

6.1 Ecosystem Architecture

Network-based geolocation relies on three conditions: (i) the network entry points of network users can be observed; (ii) the position of these entry points is known; (iii) an estimation of the distance between the network user and its entry point is known.

Network-based geolocation services fulfill these conditions, either through their position within the network, their interactions with other entities, or the data they hold.

Figure 6 presents a generic view of the network-based geolocation ecosystem and its actors:

Target. The target is any network user whose position is to be inferred. To communicate, it may reveal certain identifiers required for connectivity (*e.g.*, an IP address or an *International Mobile Subscriber Identity (IMSI)*). These identifiers may subsequently be used to identify the target during the geolocation process.

Reference Point. A *network equipment* is an identifiable network entity through which the target accesses the network, whose position is known and used to infer the target's position.

Contributor. The *network observer* is able to observe the target's presence on the network and to identify which network equipment the target is attached to. This information is used as location evidence and transmitted to the network-based geolocation service, sometimes together with additional data that can improve geolocation accuracy.

Service Provider. The *network-based geolocation service* infers the target's position based on location evidence provided by network observers. It resolves network equipment identifiers by querying the *network equipment database* and uses the resulting positions to compute the target's location data.

Reference Resolver. The *network equipment database* maps network equipment identifiers to associated attributes, including location data. These databases may be built and maintained by network owners (*e.g.*, a mobile network operator storing its base station positions), in which case access is typically restricted. They may also be constructed by third-party entities through network measurements and inference (*e.g.*, IP geolocation providers such as IPinfo) and offered through paid or public interfaces.

Consumer. The entity acting as a consumer of the produced location data may be (i) a *customer*, often with a business relationship with the service, (ii) the *service provider* itself, for purposes such as analytics or network management, or (iii) an *external entity* that gains access to the data for other reasons, including regulatory or governmental requests.

Even if they rely on the same underlying principles, the practical implementation of network-based geolocation services may differ significantly depending on the underlying network on which they are built. The characteristics of this network determine which entities can observe location evidence and which have access to network equipment databases. Consequently, services built on the same type of network tend to share similar architectural properties, and the privacy threats they expose primarily stem from the structural properties of the underlying network rather than from specific implementation choices.

6.2 Studied Network-Based Geolocation Services

We select three representative types of networks: IP networks, cellular networks, and LPWAN. For each network type, we choose a representative network-based geolocation service operating on that infrastructure.

Table 3 summarizes how each service instantiates the actors defined in Figure 6. We present the data flow diagrams for each service in Appendix E.

IPinfo. One of the IP intelligence services provided by IPinfo is IP-based geolocation. This service is offered either through an API that subscribers can query or through the direct download of the full IP-to-location database. The database is built using a combination of active measurements performed through IPinfo's probes network [8] and passive measurements, often obtained from other data providers (*e.g.*, WHOIS records).

Table 3: Network-Based Geolocation Actor Mapping.

	IPinfo	1NCE	Sigfox Atlas
Target	Any Internet-connected device	Subscribed and equipped IoT device	Subscribed and equipped IoT device
Network Equipment	IP network interfaces	Cellular towers	Sigfox antennas
Network Observer	IPinfo customers	Partner MNO[1] core networks	Sigfox core network
Network Equipment Database	IPinfo database	Partner MNO[1] databases	Sigfox private database

[1] Mobile Network Operator.

1NCE. The company operates as a *mobile virtual network operator (MVNO)* through agreements with multiple physical ISPs worldwide, allowing its customers to locate IoT devices through cellular networks and thereby provide IoT tracking services.

The geolocation mechanisms used by 1NCE do not fundamentally differ from those employed by physical ISPs. The primary distinction is that 1NCE does not own the underlying network infrastructure and instead relies on partner operators to collect the required location evidences.

Sigfox Atlas. Similarly to 1NCE, the purpose of Sigfox Atlas is to locate IoT devices, but it relies on its own LPWAN infrastructure. This system enables Sigfox to locate IoT devices that use its network by collecting Sigfox signals received by its antennas. Antennas report measurements such as the *received signal strength indicator (RSSI)*, the *signal-to-noise ratio (SNR)*, and the Sigfox transmitter identifier to the Sigfox core network, which computes a position estimate for the device and stores it to make it accessible to clients. Other LPWAN offer services very similar to Sigfox Atlas (*e.g.*, [1]).

6.3 Structural Absence of Opt-Out

When using a communication network, users inherently expose unintended location evidences stemming from the underlying communication protocols. These evidences can be exploited to infer location data. Network-based geolocation services operationalize this property by relying on such unintended evidences, thereby enabling passive geolocation without meaningful opt-out.

The Internet. In the current Internet architecture, packets emitted by a network user carry a public source IP address that reflects its topological attachment to an IP network. IP address prefixes are allocated to network operators that deploy physical infrastructure within specific geographic regions. Observing such packets therefore provides location evidence by linking the user to geographically localizable infrastructure. This assumption, however, is increasingly weakened by the growing use of overlay networks [97].

IP-based geolocation providers, such as IPinfo, are thus based on this assumption to provide an *assisted* geolocation technique to their customers (that acts as network observer), as they provide a database mapping IP blocks to location, typically at city-level granularity. Their services are widely used for purposes such as access control, logging, or commercial profiling [168].

To build and maintain their databases, such providers also rely on *autonomous* techniques, such as latency triangulation [90]. In that case, network delay measurements act as location evidences by providing distance estimates between measurement probes and the targeted infrastructure.

Cellular Networks. In cellular networks, mobile devices periodically transmit control-plane signaling messages to establish and maintain connectivity with the infrastructure. Observing these signals reveals both the identifier of the serving cell and persistent subscription identifiers (*e.g.*, IMSI, TMSI) that can be directly associated with a customer account. These signals therefore constitute location evidence by establishing spatial proximity between an identified entity and the radio cell serving it, whose geographic coordinates are known to the operator.

Operators thus maintain *Cell Site Location Information (CSLI)* records, which associate subscriber identifiers with serving cell sites and timestamps. These records are persistently stored to comply with governmental and regulatory obligations (*e.g.*, judicial investigations [145]), as well as for security monitoring and commercial purposes [68, 145]. This mechanism also enables passive allo-geolocation services such as those provided by 1NCE.

LPWAN. Devices operating on LPWANs are assigned unique identifiers (*e.g.*, a Sigfox Device ID) that are typically stable over time and revealed during communication. Network operators can observe which antenna receives a given transmission together with the associated identifier, thereby establishing location evidence. As operators know the geographic position of their antennas, they can infer the approximate location of the transmitting device.

This mechanism enables services such as Sigfox Atlas, which provide geolocation services to Sigfox network customers.

Other Communication Protocols. While not explicitly studied in these case studies, most communication protocols, particularly wireless ones, enable passive geolocation techniques. For example, Wi-Fi and BLE signaling can act as unintended location evidences [22]. Although *Randomized and Changing MAC address (RCM)* techniques [69] aim to prevent device tracking, both passive and active de-randomization attacks have been demonstrated [38]. Even hardware-induced imperfections in radio transmitters can be exploited to derive stable signatures from emitted signals using *Radio Frequency Fingerprinting (RFF)* [20, 87], enabling attackers to recognize and track devices across locations.

These unintended location evidences can be exploited by network operators acting as **A2** to infer location data from communications they legitimately observe, or by **A3** who derive location information from publicly observable signals. In both cases, persistent network identifiers enable correlating location data over time and re-identifying targets, while users typically have no way to opt out without forgoing network access altogether.

6.4 Location Data Retention

Beyond real-time correlation, the retention of historical location records significantly amplifies longitudinal tracking risks. Sigfox Atlas retains location data associated with device identifiers for up to one year, encrypted at rest. Similarly, 1NCE provides customers

with access to a seven-day location history [2], while partner *Mobile Network Operators (MNOs)* may retain attachment or location records for operational or regulatory purposes.

Access to retained records enables longitudinal correlation even without continuous observation. In this case, the threat may stem from **A2** through legitimate access, or from **A3** through unauthorized access to the stored records.

6.5 Cross-Entity Inference

Interactions between the network observer and the geolocation service may enable additional inferences about the target.

In the case of IP-based geolocation, when an Internet service queries IPinfo through its online API, IPinfo learns that users behind the queried IP address interact with that service. This creates additional visibility into user activity patterns. This threat affects IP-based geolocation services that expose online APIs, but not those where the database is downloaded and queried locally.

For 1NCE, the situation illustrates a specific inference risk linked to its MVNO model. Because 1NCE does not operate the underlying infrastructure, partner MNOs act as network observers and compute the location of attached IMSIs. By identifying IMSI ranges allocated to 1NCE, they can infer that these identifiers correspond to IoT devices managed by 1NCE rather than to ordinary mobile subscribers. This reveals contextual information beyond pure location, such as the nature of the service associated with the subscriber. Additionally, 1NCE relies on an external reference resolver, which may thus learn that 1NCE-managed devices are present near specific reference points, raising similar inference risks.

More generally, similar inferences arise in traditional operator settings. IMSIs embed operator and country identifiers (MCC/MNC), which allow network operators to determine whether a subscriber belongs to a foreign network when roaming. From this information alone, operators can infer the presence and distribution of foreign subscribers within a geographic area, revealing population patterns that go beyond individual location data.

A common pattern across these cases is that contextual information is derived by entities outside the user's trust scope, placing these inferences under the control of **A3**. Sigfox Atlas structurally avoids this threat by operating both roles within a single entity, eliminating such interactions.

6.6 Takeaways

T6.1: No Meaningful Opt-Out. Network-based geolocation services operationalize communication mechanisms that inherently provide no meaningful opt-out for the target, which may not even be aware that location data can be inferred. Such systems enable large-scale inference that can significantly impact users' privacy.

T6.2: Network Providers Hold Inference Power. Despite the absence of service-level centralization, network operators occupy a structurally privileged position that enables large-scale location inference. Only highly distributed infrastructures such as the Internet partially mitigate this capability.

T6.3: Data Retention Amplifies Privacy Risks. The storage of location data increases longitudinal inference capabilities and creates additional exposure in the event of external compromise.

7 Discussion

This section summarizes takeaways of our analysis and presents applicable existing defense mechanisms. We further discuss their limitations and outline future research directions for the community. An overview of the privacy threats identified throughout the paper is provided in Appendix D.

7.1 Global Takeaways

Centralization Creates Inference Power. Across the studied systems, inference power concentrates in central actors that impose structural trust dependencies on users (**T4.1, T5.1, T6.2**). Whether proprietary reference resolvers, service providers, or network operators, these entities observe large volumes of location-related data, enabling longitudinal correlation, cross-service linking, and large-scale profiling without deviating from the protocol design. A key driver of this concentration is the widespread reliance on assisted geolocation techniques, which externalize trust to reference resolvers that observe queries and identifiers, creating inference opportunities (**T4.1, T5.1**). Although autonomous techniques confine position estimation locally, fully autonomous approaches remain rare in practice due to accuracy limitations in certain environments and higher energy or latency costs. Contemporary systems therefore combine autonomous and assisted mechanisms, trading privacy for performance and scalability, and revealing a fundamental architectural tension between efficiency and privacy.

Privacy Risks Inherent to Communication Protocols. A second structural source of privacy risk stems from the pervasive production of unintended location evidences. Many communication protocols, particularly wireless protocols, emit signals that can be exploited for spatial inference, even when geolocation is not their primary purpose. These signals often carry persistent identifiers or quasi-identifiers, enabling not only position estimation but also target identification and cross-session linking. Such evidences are routinely leveraged by network-based geolocation services and other external actors to infer spatial information without the explicit involvement of the target (**T6.1**).

Cross-Ecosystem Inference Amplification. The mechanisms discussed above do not remain confined to a single ecosystem. Network-based geolocation services, for instance, are often integrated as auxiliary components within other ecosystems, such as OS-level geolocation systems or backend service infrastructures. As a result, inference capabilities can compound across systems: evidences exposed in one context reinforce the inference power of another, as reflected by the interaction between **T4.1** and **T5.1**. This compositional effect enlarges the overall inference surface and complicates mitigation strategies, as reducing exposure within one ecosystem does not eliminate the broader privacy risk.

Structural Dependencies Outlast Vulnerabilities. The patterns identified above suggest that threats arising from security vulnerabilities, as highlighted in **T4.2**, **T5.2**, and **T6.3**, typically manifest as discrete security incidents once a flaw is discovered and exploited, often attracting significant public and regulatory attention. In contrast, some geolocation ecosystems structurally require users to rely on centralized actors for core functions, creating a persistent

Table 4: Taxonomy of Defense Mechanisms for Geolocation Ecosystems

	Defense	Principle	Existing mechanisms	Targeted threat
Centralization	Statistical obfuscation	Reduce precision of location data revealed, for point queries or mobility traces.	Spatial cloaking [80], k-anonymity [70], Casper [121], dummy queries [105], Mixzones [40], GLOVE [79], KLT [88].	Linking and identifying users via queries or location history (§ 4.2, § 4.3, § 6.4, § 6.5).
	Oblivious query primitives	Decouple identity from query, or hide query content from the provider.	Oblivious HTTP [147], PIR [50, 155], Oblivious Transfer for LBS [83, 125, 162], Private Nearest Neighbor Search [100, 133].	Linking and identifying users via queries (§ 4.2, § 4.3, § 5.3, § 6.5).
	Query limitation	Reduce query volume by reusing past results locally.	Caching and sampling [42, 102].	Linking of resolver queries over time (§ 4.2, § 5.3).
	Interoperability	Standardize interfaces so users can substitute, distribute, or self-host providers.	User-selectable backends (e.g., GeoClue [71], Solid [160]). Regulations (e.g., EU *Digital Markets Act* [3]).	Provider lock-in (§ 4.2, § 4.3, § 5.3).
	Interaction unlinkability	Use of blind signature schemes to limit the provider's ability to link successive service requests from the same entity.	Blind signatures [46], partial blind signatures [21] (*e.g.*, Blind My [113]).	Linking of successive interactions by the provider (§ 4.2, § 4.3, § 5.3).
	Auditability and Transparency	Make provider actions verifiable a posteriori via public logs, or constrain them a priori via attested execution.	Transparency logs and verifiable audits (CT [98], gossip-based audits [116], VAMS [86], TAP [132]); confidential computing via TEEs [53, 135].	Unauthorized use of collected data by the provider (§ 4.3, § 4.4, § 5.2, § 5.3, § 5.4, § 6.4, § 6.5).
Communication Protocols	Hide the evidence	Limit the ability to link protocol identifiers to a specific entity, or to correlate multiple observations over time, through identifier concealment or rotation.	*IP:* mixnets [47, 126], VPNs [92], Tor [63], browser-integrated overlays [30, 78, 122], ODoH [94], privacy proxies [93]. *Cellular:* subscriber identifier concealment [18], temporary identifier rotation [18], location-precision reduction (e.g., Apple's Limit Precise Location [35]). *Wi-Fi:* MAC address randomization [38, 69], frame anonymization [69]. *BLE:* Resolvable Private Addresses [159], Eddystone-EID [57]. *PHY:* random noise injection [124], adversarial perturbation [108], IQ compensation with DPD [164, 165].	Passive identification and tracking via unintended evidences (§ 6.3).
	Evidence authentication	Verify the authenticity of broadcast signals and reference identifiers, through cryptographic authentication or anomaly-based detection.	*GNSS:* Galileo OSNMA [66]. *Wi-Fi:* AP spoofing detection [109]. *Cellular:* CellGuard [36], SnoopSnitch [140].	Falsification of evidences (GNSS, BSSID, Cell-ID spoofing) (§ 4.5).
	Cross-validation	Detect falsified evidences by checking consistency across independent positioning sources.	Location Validation System [136], multi-source fusion [107], WLAN-attack detection [148].	Falsification of evidences (§ 4.5).
User Control	Disclosure standardization	Mandate a standardized, machine-readable disclosure of data practices.	Privacy Nutrition Labels [91], Privacy Manifests [34], App Tracking Transparency [34].	User unawareness of data collection (§ 4.4).
	Policy verification	Automatically compare declared policies with software behavior.	Polisis [84], PolicyLint [27], PoliCheck [28], PoliGraph [56].	Provider deviations from declared practices (§ 4.4, § 5.2, § 5.3).
	Runtime visibility	Surface data accesses to the user as they occur.	App Privacy Report (iOS) [33], Privacy Dashboard (Android) [75].	User unawareness of ongoing data access by apps (§ 4.4).

privacy risk that remains less visible yet continuously present, and cannot be addressed through conventional security patches.

7.2 Defense Mechanisms

A range of defense mechanisms may be applicable to mitigate the risks identified above. This section aims to identify which of these mechanisms are applicable and which gaps remain. Table 4 presents a structured taxonomy, organized by the layer at which each mechanism acts.

Centralization. Several mechanisms aim to limit the information a provider can infer about users, either by obfuscating queries (*e.g.*, statistical obfuscation, oblivious queries) or by reducing linkability across requests (*e.g.*, blind signatures). However, their effectiveness remains limited in practice.

Spatial obfuscation techniques suffer from well-known weaknesses [88, 142], limiting their effectiveness in most of our scenarios. Oblivious query protocols such as Oblivious HTTP rely on non-collusion assumptions that are difficult to guarantee in increasingly centralized infrastructures, as illustrated by iCloud Private Relay [137]. Blind signatures reduce cryptographic linkability between requests but do not prevent cross-service or metadata-based correlation, as discussed in § 5.3. PIR remains promising but faces significant scalability challenges.

In practice, simpler architectural mitigations may provide more deployable privacy improvements. Interoperability and caching can reduce data exposure without eliminating inference risks, while auditability and transparency mechanisms improve accountability but still rely on trust assumptions.

Communication Protocols. Defenses against unintended location evidences mainly target the linkage between network identifiers and specific users. On the Internet, such mechanisms remain limited by deployment centralization and trust assumptions. Anonymous communication systems such as Tor are vulnerable to traffic correlation attacks, while commercial proxies and VPN providers introduce additional centralization. In cellular and wireless networks, privacy-enhancing mechanisms may still suffer from practical or security limitations [49, 127].

Evidence falsification can be mitigated through mechanisms such as evidence authentication and cross-validation across multiple sources. However, these protections remain only partially deployed in practice.

User Control. Various mechanisms aim to improve users' understanding and control over location privacy. Disclosure standardization mechanisms such as Privacy Nutrition Labels [91] may help users interpret opaque Terms of Conditions, while runtime visibility mechanisms increase awareness of ongoing data access. Policy verification techniques may further help identify deviations between declared and actual provider behaviors.

However, such mechanisms often provide only partial visibility and are typically limited to applications rather than the operating system itself. For instance, iOS App Privacy Report does not expose all location accesses performed by OS services.

Identified Gaps. While a wide spectrum of defense mechanisms exists, some gaps remain. Many mechanisms listed in Table 4 are simply not deployed in practice. This stems from unfavorable cost-benefit ratios for providers, technical limitations such as the scalability of PIR, deployment barriers (*e.g.*, anti-RFF mechanisms requiring hardware modifications), and known weaknesses that undermine some mechanisms in realistic settings, such as the vulnerability of statistical obfuscation to longitudinal aggregation [88, 142]. However, even simple mechanisms such as caching or interoperability remain unimplemented despite raising no significant constraints.

A more fundamental gap is the absence of mechanisms addressing inference by structurally privileged actors. Network providers (**T6.2**) retain the ability to observe which subscriber connects from which location, since identification is a precondition for network access. No existing technique prevents cross-service inference by a provider operating multiple roles within the same ecosystem, as illustrated by the bypass of report E2EE in OFNs (§ 5.3). Finally, no real opt-out exists for unintended location evidences against the network provider itself. Current approaches conceal identifiers or alter signals to limit external observers [92, 127], but cannot prevent the provider from observing the connection, since identification at the network entry point is structurally required for connectivity.

A different kind of gap concerns mechanisms aimed at user awareness and control. Existing work focuses almost exclusively on third-party applications, leaving OS-level services largely unaddressed. Yet, as shown throughout the paper, these services are themselves significant sources of privacy risk.

7.3 Future Directions

Building on the gaps identified above and our broader analysis, we outline the following research directions.

Practical Privacy Mechanisms. As discussed in § 7.2, several privacy mechanisms remain impractical to deploy at scale, which weakens the incentive for providers to adopt them. Addressing these limitations is therefore a prerequisite for broader adoption.

PIR is a representative example: it could mitigate privacy risks in queries that reveal users' positions, such as reference resolver queries or the extended services queries discussed in § 4.3, but its scalability on large databases remains an open challenge. Recent work on PIR for DNS resolution [64, 161] provides a relevant starting point, as reference resolver queries share a similar key-lookup structure. Spatial queries based on a position, however, remain underexplored [118] and would require further dedicated research.

Beyond making individual mechanisms practical, a second challenge is extending existing anti-linkability mechanisms to realistic settings where a provider observes requests to multiple services it operates and their associated metadata.

Auditability. Since fully eliminating centralized intermediaries is unrealistic and cryptographic protections alone cannot prevent inference (*e.g.*, § 5.3), making structurally central actors auditable is a necessary complement. While existing mechanisms provide relevant foundations, their adaptation to location data remains largely unexplored. Two challenges are particularly specific to this context: audit traces of location accesses may themselves reveal sensitive information, and inference-based threats cannot be captured by mechanisms designed to log explicit accesses.

Extending User Control to OS-Level Services. Existing work on user control, policy verification, and runtime visibility predominantly targets third-party applications, leaving OS-level services largely unaddressed despite their direct impact on users' privacy. Extending these mechanisms to system components raises specific challenges, as OS services operate in a privileged context with no external audit surface comparable to application stores.

Extending the Scope of Privacy Analysis. Several classes of geolocation ecosystems remain underexplored, in particular collaborative environments such as vehicular networks [167], UAV swarms [149], or peer-to-peer localization [166], whose more distributed architectures have not yet been analyzed through the lens proposed in this work. A natural direction is to investigate whether similar structural patterns arise in such settings. Beyond ecosystem-level studies, a more systematic characterization of location evidences and their composition could also clarify the structural limits of privacy-preserving location data production.

8 Conclusion

This work introduced an abstract model and a systematic methodology for analyzing privacy risks in modern geolocation ecosystems. By structuring heterogeneous systems under a common set of roles and data flows, we provided a unified framework for reasoning about privacy beyond isolated case-specific analyses.

We applied this methodology to a broad range of modern geolocation ecosystems, covering diverse architectural designs. Across these case studies, we systematically analyzed the privacy risks arising from each ecosystem's architecture and design choices and identified common structural patterns, including excessive centralization, unintended geolocation enabled by communication protocols, and the amplification of risks across interacting ecosystems, that shape privacy exposure despite their apparent heterogeneity.

Building on this analysis, we reviewed existing defense mechanisms applicable to these ecosystems and identified structural gaps that current approaches fail to address. These findings suggest that privacy risks in modern geolocation ecosystems cannot be addressed solely through cryptographic protections or security hardening. Instead, they call for rethinking architectural assumptions and motivate concrete research directions toward privacy-oriented design.

Acknowledgments

We thank the reviewers and our revision editor for their constructive feedback. We warmly thank Lionel Molinier from Sentiens for his contribution to our understanding of the Sigfox ecosystem.

This work has been supported by a grant from the Agence Nationale de la Recherche (project no. ANR-24-CE25-1133 [GTTP]).

The authors used generative AI-based tools to revise the text, improve flow and correct any typos, grammatical errors, and awkward phrasing.

References

[1] 2020. LoRaWAN Geolocation White Paper. https://lora-alliance.org/wp-content/uploads/2020/11/geolocation_whitepaper.pdf Accessed January 2026.
[2] 1NCE 2024. *1NCE OS Device Locator*. 1NCE. https://help.1nce.com/dev-hub/docs/1nce-os-device-locator Accessed: 2026-02-12.
[3] 2024. *Interoperability - Digital Markets Act (DMA) - European Commission*. https://digital-markets-act.ec.europa.eu/developer-portal/interoperability_en
[4] 2025. 1NCE: The Only Provider of Flat Rate IoT Connectivity. https://1nce.com/. Accessed: 2025-01-17.
[5] 2025. Apple Find My. https://www.apple.com/icloud/find-my/
[6] 2025. Find My & Privacy. https://www.apple.com/legal/privacy/data/en/find-my/ Accessed January 2026.
[7] 2025. Google Find Hub. https://www.android.com/learn-find-hub/
[8] 2025. IPinfo Probe Network. https://ipinfo.io/probe-network Accessed January 2026.
[9] 2025. IPinfo: The Trusted Source for IP Address Data. https://ipinfo.io/. Accessed: 2025-01-17.
[10] 2025. Sigfox Atlas: IoT Location Service without GPS. https://www.sigfox.com/services/atlas/. Accessed: 2025-01-17.
[11] Apple Developer Documentation 2026. *Core Location*. Apple Developer Documentation. https://developer.apple.com/documentation/corelocation
[12] 2026. *Fused Location Provider API | Google for Developers*. https://developers.google.com/location-context/fused-location-provider
[13] 2026. How Find Hub protects your data. https://support.google.com/product-documentation/answer/14796936 Accessed January 2026.
[14] 2026. *Manage Your Android Device's Location Settings - Google Account Help*. https://support.google.com/accounts/answer/3467281
[15] 2026. Samsung SmartThings Find. https://smartthingsfind.samsung.com Accessed January 2026.
[16] ISO/IEC JTC 1/SC 27. 2019. *Information technology — Security techniques — Privacy engineering for system life cycle processes*. https://www.iso.org/standard/72024.html
[17] 3GPP. 2019. *TS 38.331: Radio Resource Control (RRC) protocol specification*. Technical Report TS 38.331. 3rd Generation Partnership Project (3GPP). https://www.etsi.org/deliver/etsi_ts/138300_138399/138331/15.05.01_60/ts_138331v150501p.pdf
[18] 3GPP. 2024. *TS 33.501: Security architecture and procedures for 5G System*. Technical Report. 3rd Generation Partnership Project (3GPP). https://www.etsi.org/deliver/etsi_ts/133500_133599/133501/
[19] .NET API 4.8.1. 2022. *CivicAddress Class (System.Device.Location)*. https://learn.microsoft.com/en-us/dotnet/api/system.device.location.civicaddress?view=netframework-4.8.1
[20] Sohail Abbas, Manar Abu Talib, Qassim Nasir, Sally Idhis, Mariam Alaboudi, and Ali Mohamed. 2024. Radio frequency fingerprinting techniques for device identification: a survey. *International Journal of Information Security* 23, 2 (2024), 1389–1427.
[21] Masayuki Abe and Tatsuaki Okamoto. 2000. Provably secure partially blind signatures. In *Annual international cryptology conference*. Springer, 271–286.
[22] Ali Abedi and Deepak Vasisht. 2022. Non-cooperative wi-fi localization & its privacy implications. In *Proceedings of the 28th Annual International Conference On Mobile Computing And Networking*. 570–582.
[23] Mark Ackerman, Trevor Darrell, and Daniel J Weitzner. 2001. Privacy in context. *Human–Computer Interaction* 16, 2-4 (2001), 167–176.
[24] ACLU. 2013. *Cotraveler Overview*. American Civil Liberties Union. https://www.aclu.org/documents/cotraveler-overview
[25] Hosam Alamleh, Ali Abdullah S AlQahtani, Michael Gogarty, and David Ruddel. 2025. Securing the invisible thread: A comprehensive analysis of ble tracker security in apple airtags and samsung smarttags. *Journal of Umm Al-Qura University for Engineering and Architecture* (2025), 1–19.
[26] Daniel Do Prado Allão, Isadora G Ferrão, Vitor Marçal, Lucas Gonzalez Da Silva, and Kalinka C Branco. 2025. A Systematic Review of GPS Spoofing: Methods, Tools, Tests, and Techniques in the State of the Art. In *2025 International Conference on Unmanned Aircraft Systems (ICUAS)*. IEEE, 1019–1026.
[27] Benjamin Andow, Samin Yaseer Mahmud, Wenyu Wang, Justin Whitaker, William Enck, Bradley Reaves, Kapil Singh, and Tao Xie. 2019. PolicyLint: investigating internal privacy policy contradictions on google play. In *28th USENIX security symposium (USENIX security 19)*. 585–602.
[28] Benjamin Andow, Samin Yaseer Mahmud, Justin Whitaker, William Enck, Bradley Reaves, Kapil Singh, and Serge Egelman. 2020. Actions speak louder than words:Entity-Sensitive privacy policy and data flow analysis with PoliCheck. In *29th USENIX Security Symposium (USENIX Security 20)*. 985–1002.
[29] Apple. 2019. Location Services Privacy Overview. (2019). https://www.apple.com/privacy/docs/Location_Services_White_Paper_Nov_2019.pdf
[30] Apple. 2021. iCloud Private Relay Overview. https://www.apple.com/privacy/docs/iCloud_Private_Relay_Overview_Dec2021.PDF Accessed: 2026-05-08.
[31] Apple. 2024. Apple Platform Security: Secure Enclave. https://support.apple.com/fr-fr/guide/security/sec59b0b31ff/web. Accessed: 2026-01-04.
[32] Apple. 2025. Location Services & Privacy. https://www.apple.com/legal/privacy/data/en/location-services/. Accessed: 2026-01-04.
[33] Apple. n.d.. About App Privacy Report. Apple Support Documentation. https://support.apple.com/en-us/102188
[34] Apple. n.d.. Privacy Manifest Files. Apple Developer Documentation. https://developer.apple.com/documentation/bundleresources/privacy-manifest-files
[35] Apple Inc. 2026. *Limit precise location from cellular networks*. Apple Inc. https://support.apple.com/en-us/126101 Accessed: 2026-02-11.
[36] Lukas Arnold, Matthias Hollick, and Jiska Classen. 2024. Catch you cause i can: busting rogue base stations using cellguard and the apple cell location database. In *Proceedings of the 27th International Symposium on Research in Attacks, Intrusions and Defenses*. 613–629.
[37] David Bassali, Adam Kinkley, Katie Ning, and Jackson Skeen. 2020. Google's Anticompetitive Practices in Mobile: Creating Monopolies to Sustain a Monopoly. (2020). https://som.yale.edu/sites/default/files/2022-01/DTH-GoogleMobile.pdf Working Paper Series No. 5.
[38] Johannes K Becker, David Li, and David Starobinski. 2019. Tracking anonymized bluetooth devices. *Proceedings on Privacy Enhancing Technologies* (2019).
[39] Philipp Beer, Marco Squarcina, Sebastian Roth, and Martina Lindorfer. 2025. TapTrap: Animation-Driven Tapjacking on Android. In *34th USENIX Security Symposium (USENIX Security 25)*. 3317–3335.
[40] Alastair R. Beresford and Frank Stajano. 2003. Location Privacy in Pervasive Computing. *IEEE Pervasive Computing* 2, 1 (2003), 46–55. https://doi.org/10.1109/MPRV.2003.1186725
[41] Leon Böttger, Alexander Matern, Dennis Arndt, and Matthias Hollick. 2025. Okay Google, Where's My Tracker? Security, Privacy, and Performance Evaluation of Google's Find My Device Network. *Proceedings on Privacy Enhancing Technologies* (2025).
[42] Antoine Boutet and Mathieu Cunche. 2021. Privacy protection for Wi-Fi location positioning systems. *Journal of information security and applications* 58 (2021), 102635.
[43] Mark Brnovich. 2022. *Attorney General Mark Brnovich Achieves Historic $85 Million Settlement with Google | Attorney General's Office*. https://www.azag.gov/press-release/attorney-general-mark-brnovich-achieves-historic-85-million-settlement-google
[44] Guillaume Celosia and Mathieu Cunche. 2020. Discontinued privacy: Personal data leaks in apple bluetooth-low-energy continuity protocols. *Proceedings on Privacy Enhancing Technologies* 2020 (2020), 26–46.
[45] Badreddine Chah, Alexandre Lombard, Anis Bkakria, Reda Yaich, Abdeljalil Abbas-Turki, and Stéphane Galland. 2022. Privacy threat analysis for connected and autonomous vehicles. *Procedia Computer Science* 210 (2022), 36–44.
[46] David Chaum. 1983. Blind Signature System.. In *Crypto*, Vol. 83. Springer, 153.
[47] David L Chaum. 1981. Untraceable electronic mail, return addresses, and digital pseudonyms. *Commun. ACM* 24, 2 (1981), 84–90.
[48] Junming Chen, Xiaoyue Ma, Lannan Luo, and Qiang Zeng. 2025. Tracking you from a thousand miles away! turning a bluetooth device into an apple AirTag without root privileges. In *34th USENIX Security Symposium (USENIX Security 25)*. 4345–4362.
[49] Merlin Chlosta, David Rupprecht, Christina Pöpper, and Thorsten Holz. 2021. 5G SUCI-Catchers: Still catching them all?. In *Proceedings of the 14th ACM Conference on Security and Privacy in Wireless and Mobile Networks*. 359–364.
[50] Benny Chor, Eyal Kushilevitz, Oded Goldreich, and Madhu Sudan. 1998. Private information retrieval. *Journal of the ACM (JACM)* 45, 6 (1998), 965–981.
[51] Rebecca Clement. 2024. *27 Cellular IoT MVNOs Managed 25% of Connections Outside China in 2023*. Kaleido Intelligence. https://kaleidointelligence.com/27-cellular-iot-mvnos-managed-25-of-connections-outside-china-in-2023/
[52] Lena Cohen. 2025. *Online Behavioral Ads Fuel the Surveillance Industry—Here's How*. Electronic Frontier Foundation. https://www.eff.org/deeplinks/2025/01/online-behavioral-ads-fuel-surveillance-industry-heres-how
[53] Confidential Computing Consortium. 2022. *A Technical Analysis of Confidential Computing*. Technical Report. Confidential Computing Consortium (Linux Foundation). https://confidentialcomputing.io/wp-content/uploads/sites/10/2023/03/CCC-A-Technical-Analysis-of-Confidential-Computing-v1.3_unlocked.pdf

[54] Connecticut Attorney General's Office. 2022. Attorney General Tong Announces Historic Google Settlement Over Location Tracking Practices. https://portal.ct.gov/AG/Press-Releases/2022-Press-Releases/Attorney-General-Tong-Announces-Historic-Google-Settlement-Over-Location-Tracking-Practices. Settlement with Google and 39 other U.S. state attorneys general.

[55] Aldo Cortesi, Maximilian Hils, Thomas Kriechbaumer, and contributors. 2010–2026. mitmproxy: A free and open source interactive HTTPS proxy. https://mitmproxy.org. Version 12.x, accessed 2026.

[56] Hao Cui, Rahmadi Trimananda, Athina Markopoulou, and Scott Jordan. 2023. PoliGraph: Automated privacy policy analysis using knowledge graphs. In *32nd USENIX Security Symposium (USENIX Security 23)*. 1037–1054.

[57] Liron David, Avinatan Hassidim, Yossi Matias, Moti Yung, and Alon Ziv. 2022. Eddystone-EID: Secure and private infrastructural protocol for BLE beacons. *IEEE Transactions on Information Forensics and Security* 17 (2022), 3877–3889.

[58] Yves-Alexandre De Montjoye, César A Hidalgo, Michel Verleysen, and Vincent D Blondel. 2013. Unique in the crowd: The privacy bounds of human mobility. *Scientific reports* 3, 1 (2013), 1376.

[59] Paula Delgado-Santos, Giuseppe Stragapede, Ruben Tolosana, Richard Guest, Farzin Deravi, and Ruben Vera-Rodriguez. 2022. A survey of privacy vulnerabilities of mobile device sensors. *ACM Computing Surveys (CSUR)* 54, 11s (2022), 1–30.

[60] Mina Deng, Kim Wuyts, Riccardo Scandariato, Bart Preneel, and Wouter Joosen. 2011. A privacy threat analysis framework: supporting the elicitation and fulfillment of privacy requirements. *Requirements Engineering* 16, 1 (2011), 3–32.

[61] Nathanael Denis, Maryline Laurent, and Sophie Chabridon. 2024. A decentralized model for usage and information flow control in distributed systems. *Computers & Security* 144 (2024), 103975.

[62] Jayati Dev, Bahman Rashidi, and Vaibhav Garg. 2023. Models of applied privacy (map): A persona based approach to threat modeling. In *Proceedings of the 2023 CHI Conference on human factors in computing systems*. 1–15.

[63] Roger Dingledine, Nick Mathewson, and Paul Syverson. 2004. Tor: The second-generation onion router. (2004).

[64] Natnatee Dokmai, L Jean Camp, and Ryan Henry. 2022. Assisted Private Information Retrieval. *Cryptology ePrint Archive* (2022).

[65] Harry Eldridge, Gabrielle Beck, Matthew Green, Nadia Heninger, and Abhishek Jain. 2024. Abuse-Resistant Location Tracking: Balancing Privacy and Safety in the Offline Finding Ecosystem. In *33rd USENIX Security Symposium (USENIX Security 24)*. 5431–5448.

[66] European Union Agency for the Space Programme. 2023. *Galileo Open Service Navigation Message Authentication (OSNMA) Signal-In-Space Interface Control Document*. Technical Document Issue 1.1. European Union Agency for the Space Programme (EUSPA). https://www.gsc-europa.eu/sites/default/files/sites/all/files/Galileo_OSNMA_SIS_ICD_v1.1.pdf

[67] Ali Farzanehfar, Florimond Houssiau, and Yves-Alexandre de Montjoye. 2021. The risk of re-identification remains high even in country-scale location datasets. *Patterns* 2, 3 (2021).

[68] Federal Communications Commission. 2024. *FCC Fines Largest Wireless Carriers for Sharing Location Data*. Technical Report. FCC. https://docs.fcc.gov/public/attachments/DOC-402213A1.pdf

[69] Domenico Ficara, Rosario G Garroppo, and Jerome Henry. 2023. A tutorial on privacy, RCM and its implications in WLAN. *IEEE Communications Surveys & Tutorials* 26, 2 (2023), 1003–1040.

[70] Bugra Gedik and Ling Liu. 2008. Protecting location privacy with personalized k-anonymity: Architecture and algorithms. *IEEE Transactions on Mobile Computing* 7, 1 (2008), 1–18.

[71] GeoClue Developers. 2025. *GeoClue: A D-Bus Geolocation Service*. https://gitlab.freedesktop.org/geoclue/geoclue Freedesktop.org project.

[72] GlobalStats. 2025. *Operating System Market Share Worldwide*. StatCounter Global Stats. https://gs.statcounter.com/os-market-share

[73] GlobeNewswire. 2025. *IPinfo Recognized as Leader in Snowflake's 2026 Modern Marketing Data Stack Report*. The Manila Times. https://www.manilatimes.net/2025/10/01/tmt-newswire/globenewswire/ipinfo-recognized-as-leader-in-snowflakes-2026-modern-marketing-data-stack-report/2192676

[74] Liliana Gonzalez, Pedro Wightman Rojas, M Labrador, et al. 2014. A survey on privacy in location-based services. *Ingeniería y Desarrollo* 32, 2 (2014), 314–343.

[75] Google. 2021. What's New in Android Privacy. Android Developers Blog. https://android-developers.googleblog.com/2021/05/android-security-and-privacy-recap.html

[76] Google. 2023. *Google Account Help: Manage your Timeline data*. Google. https://support.google.com/accounts/answer/3118687

[77] Google. 2024. *Google Awareness API*. Google. https://developers.google.com/awareness

[78] Google. 2025. IP Protection. Privacy Sandbox, Google. https://privacysandbox.google.com/protections/ip-protection Accessed: 2026-05-08.

[79] Marco Gramaglia and Marco Fiore. 2015. Hiding Mobile Traffic Fingerprints with GLOVE. In *Proceedings of the 11th ACM Conference on Emerging Networking Experiments and Technologies (CoNEXT '15)*. ACM. https://doi.org/10.1145/2716281.2836111

[80] Marco Gruteser and Dirk Grunwald. 2003. Anonymous usage of location-based services through spatial and temporal cloaking. In *Proceedings of the 1st international conference on Mobile systems, applications and services*. 31–42.

[81] Johanna Gunawan, Amogh Pradeep, David Choffnes, Woodrow Hartzog, and Christo Wilson. 2021. A comparative study of dark patterns across web and mobile modalities. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW2 (2021), 1–29.

[82] Avril D. Haines. 2023. *DNI Haines Statement on Declassified Report on Commercially Available Information*. https://web.archive.org/web/20230614224938/https://www.dni.gov/index.php/newsroom/press-releases/press-releases-2023/item/2390-dni-haines-statement-on-declassified-report-on-commercially-available-information

[83] Jinguang Han. 2019. Oblivious Location-Based Service Query. *arXiv preprint arXiv:1907.03045* (2019).

[84] Hamza Harkous, Kassem Fawaz, Rémi Lebret, Florian Schaub, Kang G Shin, and Karl Aberer. 2018. Polisis: Automated analysis and presentation of privacy policies using deep learning. In *27th USENIX Security Symposium (USENIX Security 18)*. 531–548.

[85] Alexander Heinrich, Milan Stute, Tim Kornhuber, and Matthias Hollick. 2021. Who can find my devices? security and privacy of apple's crowd-sourced bluetooth location tracking system. *arXiv preprint arXiv:2103.02282* (2021).

[86] Alexander Hicks, Vasilios Mavroudis, Mustafa Al-Bassam, Sarah Meiklejohn, and Steven J Murdoch. 2018. VAMS: verifiable auditing of access to confidential data. *arXiv preprint arXiv:1805.04772* (2018).

[87] Anu Jagannath, Jithin Jagannath, and Prem Sagar Pattanshetty Vasanth Kumar. 2022. A comprehensive survey on radio frequency (RF) fingerprinting: Traditional approaches, deep learning, and open challenges. *Computer Networks* 219 (2022), 109455.

[88] Fengmei Jin, Wen Hua, Matteo Francia, Pingfu Chao, Maria E Orlowska, and Xiaofang Zhou. 2023. A Survey and Experimental Study on Privacy-Preserving Trajectory Data Publishing. *IEEE Transactions on Knowledge and Data Engineering* 35, 6 (2023), 5577–5596. https://doi.org/10.1109/TKDE.2022.3174204

[89] Samantha Katcher, Ben Ballard, Cara Bloom, Katie Isaacson, Julie McEwen, Stuart Shapiro, Shelby Slotter, Mark Paes, and Ryan Xu. 2024. The PANOPTIC™ Privacy Threat Model. In *Twentieth Symposium on Usable Privacy and Security (SOUPS)*.

[90] Ethan Katz-Bassett, John P John, Arvind Krishnamurthy, David Wetherall, Thomas Anderson, and Yatin Chawathe. 2006. Towards IP geolocation using delay and topology measurements. In *Proceedings of the 6th ACM SIGCOMM conference on Internet measurement*. 71–84.

[91] Patrick Gage Kelley, Joanna Bresee, Lorrie Faith Cranor, and Robert W Reeder. 2009. A" nutrition label" for privacy. In *Proceedings of the 5th Symposium on Usable Privacy and Security*. 1–12.

[92] Mohammad Taha Khan, Joe DeBlasio, Geoffrey M. Voelker, Alex C. Snoeren, Chris Kanich, and Narseo Vallina-Rodriguez. 2018. An Empirical Analysis of the Commercial VPN Ecosystem. In *Proceedings of the Internet Measurement Conference 2018 (IMC '18)*. ACM, 443–456. https://doi.org/10.1145/3278532.3278570

[93] Heiko Kiesel and Jiska Classen. 2025. Privacy Promises Unmasked: A Comprehensive Study of Privacy Proxies for the Masses. In *Proceedings of the Twenty-sixth International Symposium on Theory, Algorithmic Foundations, and Protocol Design for Mobile Networks and Mobile Computing*. 351–360.

[94] Eric Kinnear, Patrick McManus, Tommy Pauly, Tanya Verma, and Christopher A Wood. 2022. Oblivious dns over https. *RFC 9230* (2022).

[95] John Krumm. 2007. Inference attacks on location tracks. In *International Conference on Pervasive Computing*. Springer, 127–143.

[96] John Krumm. 2009. A survey of computational location privacy. *Personal and Ubiquitous Computing* 13, 6 (2009), 391–399.

[97] Augustin Laouar, Loïc Desgeorges, Paul Schmitt, and Franceso Bronzino. 2025. Rethinking Geolocalization on the Internet. In *Proceedings of the 24th ACM Workshop on Hot Topics in Networks*. 236–244.

[98] Ben Laurie. 2014. Certificate transparency. *Commun. ACM* 57, 10 (2014), 40–46.

[99] Douglas J Leith. 2021. Mobile handset privacy: Measuring the data ios and android send to apple and google. In *International Conference on Security and Privacy in Communication Systems*. Springer, 231–251.

[100] Jingyu Li, Zhicong Huang, Min Zhang, Cheng Hong, Jian Liu, Tao Wei, and Wenguang Chen. 2025. Panther: Private approximate nearest neighbor search in the single server setting. In *Proceedings of the 2025 ACM SIGSAC Conference on Computer and Communications Security*. 365–379.

[101] Tong Li, Jiaxin Liang, Yukuan Ding, Kai Zheng, Xu Zhang, and Ke Xu. 2023. On design and performance of offline finding network. In *IEEE INFOCOM 2023-IEEE Conference on Computer Communications*. IEEE, 1–10.

[102] Yufeng Li, Bo Wang, Qi Liu, Xiangyu Zheng, Jiangtao Li, Yiwei Wang, Jiajun Xi, and Wutao Qin. 2025. LPPS-IKHC: Location Privacy-Preserving Scheme Using Improved k-Anonymity and Hybrid Cache for IoV. *IEEE Transactions on Vehicular Technology* 74, 8 (2025), 12864–12878. https://doi.org/10.1109/TVT.2025.3553550

[103] LINC. 2019. *IP Report: Shaping Choices in the Digital World | Linc.* https://www.cnil.fr/sites/cnil/files/2023-06/cnil_ip_report_06_shaping_choices_in_the_digital_world.pdf

[104] Thomas Lindner, Lothar Fritsch, Kilian Plank, and Kai Rannenberg. 2004. Exploitation of public and private WiFi coverage for new business models. In *Building the E-Service Society: E-Commerce, E-Business, and E-Government*. Springer, 131–148.

[105] Hai Liu, Shengqing Hu, and NaNa Zhou. 2026. Location Privacy Protection in Continuous LBSs: Enhancing Anonymity via Fake Queries. 38, 2 (2026), 53. https://doi.org/10.1007/s44443-025-00438-z

[106] Wen Liu, Qianqian Cheng, Zhongliang Deng, Hong Chen, Xiao Fu, Xinyu Zheng, Shixuan Zheng, Cunzhe Chen, and Shuo Wang. 2019. Survey on CSI-based indoor positioning systems and recent advances. In *2019 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*. IEEE, 1–8.

[107] Wenjie Liu and Panos Papadimitratos. 2025. Coordinated Position Falsification Attacks and Countermeasures for Location-Based Services. *IEEE Open Journal of the Communications Society* (2025).

[108] Zhaoyi Lu, Ming Tu, Xin Xie, Wenchao Xu, Cunqing Hua, and Weihua Zhuang. 2025. Concealing radio frequency fingerprints via active adversarial perturbation. *IEEE Transactions on Network Science and Engineering* (2025).

[109] Juraj Machaj, Clément Safon, Slavomír Matúška, and Peter Brída. 2024. Detection of access point spoofing in the Wi-Fi fingerprinting based positioning. *Sensors* 24, 23 (2024), 7624.

[110] Jeremy Martin, Douglas Alpuche, Kristina Bodeman, Lamont Brown, Ellis Fenske, Lucas Foppe, Travis Mayberry, Erik Rye, Brandon Sipes, and Sam Teplov. 2019. Handoff All Your Privacy–A Review of Apple's Bluetooth Low Energy Continuity Protocol. *Proceedings on Privacy Enhancing Technologies* (2019).

[111] Célestin Matte. 2017. *Wi-Fi tracking: Fingerprinting attacks and counter-measures*. Ph. D. Dissertation. Université de Lyon.

[112] Célestin Matte, Jagdish Prasad Achara, and Mathieu Cunche. 2015. Device-to-identity linking attack using targeted wi-fi geolocation spoofing. In *Proceedings of the 8th ACM Conference on Security & Privacy in Wireless and Mobile Networks*. 1–6.

[113] Travis Mayberry, Erik-Oliver Blass, and Ellis Fenske. 2023. Blind my-an improved cryptographic protocol to prevent stalking in apple's find my network. *Proceedings on Privacy Enhancing Technologies* (2023).

[114] Travis Mayberry, Ellis Fenske, Dane Brown, Jeremy Martin, Christine Fossaceca, Erik C Rye, Sam Teplov, and Lucas Foppe. 2021. Who tracks the trackers? circumventing apple's anti-tracking alerts in the find my network. In *Proceedings of the 20th Workshop on Workshop on Privacy in the Electronic Society*. 181–186.

[115] Aleecia M McDonald and Lorrie Faith Cranor. 2008. The cost of reading privacy policies. *Isjlp* 4 (2008), 543.

[116] Sarah Meiklejohn, Pavel Kalinnikov, Cindy S Lin, Martin Hutchinson, Gary Belvin, Mariana Raykova, and Al Cutter. 2020. Think global, act local: Gossip and client audits in verifiable data structures. *arXiv preprint arXiv:2011.04551* (2020).

[117] Lianxiao Meng, Lin Yang, Wu Yang, and Long Zhang. 2022. A survey of GNSS spoofing and anti-spoofing technology. *Remote sensing* 14, 19 (2022), 4826.

[118] Yinbin Miao, Jiaqi Yu, Jiliang Li, Xinghua Li, Jun Feng, Zhiquan Liu, and Robert H Deng. 2025. Security-Enhanced Spatial Range Query Over Large-Scale Encrypted Mobile Cloud Datasets. *IEEE Transactions on Mobile Computing* (2025).

[119] Microsoft. 2026. *Windows Location Service and Privacy - Microsoft Support.* https://support.microsoft.com/en-us/windows/windows-location-service-and-privacy-3a8eee0a-5b0b-dc07-eede-2a5ca1c49088

[120] mjlst. 2025. *Caught in the Digital Dragnet: The Controversy Over Geofence Warrants and Privacy Rights | LawSci Forum.* https://mjlst.lib.umn.edu/2025/01/20/caught-in-the-digital-dragnet-the-controversy-over-geofence-warrants-and-privacy-rights/

[121] Mohamed F Mokbel, Chi-Yin Chow, and Walid G Aref. 2006. The new casper: Query processing for location services without compromising privacy. In *Proceedings of the 32nd international conference on Very large data bases*. 763–774.

[122] Mozilla. 2025. Use built-in VPN in Firefox. Mozilla Support. https://support.mozilla.org/en-US/kb/built-in-vpn Accessed: 2026-05-08.

[123] Khuong An Nguyen, Raja Naeem Akram, Konstantinos Markantonakis, Zhiyuan Luo, and Chris Watkins. 2019. Location tracking using smartphone accelerometer and magnetometer traces. In *Proceedings of the 14th International Conference on Availability, Reliability and Security*. 1–9.

[124] Gabriele Oligeri and Savio Sciancalepore. 2024. Hideprint: Hiding the radio fingerprint via random noise. *arXiv preprint* (2024).

[125] Russell Paulet, Md Golam Kaosar, Xun Yi, and Elisa Bertino. 2013. Privacy-preserving and content-protecting location based queries. *IEEE transactions on knowledge and data engineering* 26, 5 (2013), 1200–1210.

[126] Ania M Piotrowska, Jamie Hayes, Tariq Elahi, Sebastian Meiser, and George Danezis. 2017. The loopix anonymity system. In *26th usenix security symposium (usenix security 17)*. 1199–1216.

[127] John Preuß Mattsson and Prajwol Kumar Nakarmi. 2021. Nori: Concealing the concealed identifier in 5g. In *Proceedings of the 16th International Conference on Availability, Reliability and Security*. 1–7.

[128] Vincent Primault, Antoine Boutet, Sonia Ben Mokhtar, and Lionel Brunie. 2018. The long road to computational location privacy: A survey. *IEEE Communications Surveys & Tutorials* 21, 3 (2018), 2772–2793.

[129] Guoying Qiu, Guoming Tang, Chuandong Li, Lailong Luo, Deke Guo, and Yulong Shen. 2023. Differentiated location privacy protection in mobile communication services: A survey from the semantic perception perspective. *Comput. Surveys* 56, 3 (2023), 1–36.

[130] Shriya Raban. 2022. *The Number of Active Consumer Asset Tracking Devices in Europe and North America to Reach 28 Million in 2026 | IoT Now News & Reports.* IoT Now News - How to run an IoT enabled business. https://www.iot-now.com/2022/05/25/121446-the-number-of-active-consumer-asset-tracking-devices-in-europe-and-north-america-to-reach-28-million-in-2026/

[131] Joel R Reidenberg, Travis Breaux, Lorrie Faith Cranor, Brian French, Amanda Grannis, James T Graves, Fei Liu, Aleecia McDonald, Thomas B Norton, Rohan Ramanath, et al. 2015. Disagreeable privacy policies: Mismatches between meaning and users' understanding. *Berkeley Tech. LJ* 30 (2015), 39.

[132] Daniël Reijsbergen, Aung Maw, Zheng Yang, Tien Tuan Anh Dinh, and Jianying Zhou. 2023. TAP: Transparent and Privacy-Preserving data services. In *32nd USENIX Security Symposium (USENIX Security 23)*. 6489–6506.

[133] M Sadegh Riazi, Beidi Chen, Anshumali Shrivastava, Dan Wallach, and Farinaz Koushanfar. 2016. Sub-linear privacy-preserving near-neighbor search. *arXiv preprint arXiv:1612.01835* (2016).

[134] Erik Rye and Dave Levin. 2024. Surveilling the masses with wi-fi-based positioning systems. In *2024 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2831–2846.

[135] Mohamed Sabt, Mohammed Achemlal, and Abdelmadjid Bouabdallah. 2015. Trusted execution environment: What it is, and what it is not. In *2015 IEEE Trustcom/BigDataSE/Ispa*, Vol. 1. IEEE, 57–64.

[136] Andrea Saracino, Francesco Restuccia, and Fabio Martinelli. 2018. Practical location validation in participatory sensing through mobile wifi hotspots. In *2018 17th IEEE International Conference On Trust, Security And Privacy In Computing And Communications/12th IEEE International Conference On Big Data Science And Engineering (TrustCom/BigDataSE)*. IEEE, 596–607.

[137] Patrick Sattler, Juliane Aulbach, Johannes Zirngibl, and Georg Carle. 2022. Towards a tectonic traffic shift? investigating apple's new relay network. In *Proceedings of the 22nd ACM Internet Measurement Conference*. 449–457.

[138] Victoria Saxe. 2020. Junk Evidence: A Call to Scrutinize Historical Cell Site Location Evidence. *UNHL Rev.* 19 (2020), 133.

[139] David Schmidt, Alexander Ponticello, Magdalena Steinböck, Katharina Krombholz, and Martina Lindorfer. 2025. Analyzing the iOS Local Network Permission from a Technical and User Perspective. In *2025 IEEE Symposium on Security and Privacy (SP)*. IEEE, 4229–4247.

[140] Security Research Labs. n.d.. SnoopSnitch. Open source project. https://github.com/srlabs/snoopsnitch

[141] Nataliya Shevchenko, Timothy A Chick, Paige O'riordan, Thomas P Scanlon, and Carol Woody. 2018. *Threat modeling: a summary of available methods*. Technical Report.

[142] Reza Shokri, George Theodorakopoulos, Jean-Yves Le Boudec, and Jean-Pierre Hubaux. 2011. Quantifying location privacy. In *2011 IEEE symposium on security and privacy*. IEEE, 247–262.

[143] Daehyun Strobel. 2007. IMSI catcher. *Chair for Communication Security, Ruhr-Universität Bochum* 14 (2007).

[144] European Data Protection Supervisor. 2018. *Preliminary Opinion on Privacy by Design.* https://www.edps.europa.eu/data-protection/our-work/publications/opinions/privacy-design

[145] Supreme Court of the United States. 2018. Carpenter v. United States, 585 U.S. 1. https://www.supremecourt.gov/opinions/17pdf/16-402_h315.pdf U.S. Supreme Court decision on historical cell site location information.

[146] Hui Jun Tay, Jiaqi Tan, and Priya Narasimhan. 2016. A survey of security vulnerabilities in bluetooth low energy beacons. *Carnegie Mellon University Parallel Data Lab Technical Report CMU-PDL-16-109* (2016).

[147] Martin Thomson and Christopher A Wood. 2024. RFC 9458: Oblivious HTTP.

[148] Nils Ole Tippenhauer, Kasper Bonne Rasmussen, Christina Pöpper, and Srdjan Čapkun. 2009. Attacks on public WLAN-based positioning systems. In *Proceedings of the 7th international conference on Mobile systems, applications, and services*. 29–40.

[149] Pengfei Tong, Xuerong Yang, Yajun Yang, Wei Liu, and Peiyi Wu. 2023. Multi-UAV collaborative absolute vision positioning and navigation: A survey and discussion. *Drones* 7, 4 (2023), 261.

[150] Vincent Toubiana and Mathieu Cunche. 2021. No need to ask the Android: Bluetooth-Low-Energy scanning without the location permission. In *Proceedings of the 14th ACM Conference on Security and Privacy in Wireless and Mobile Networks*. 147–152.

[151] Daniel Townley, Kerem Arıkan, Yu David Liu, Dmitry Ponomarev, and Oğuz Ergin. 2022. Composable cachelets: Protecting enclaves from cache Side-Channel attacks. In *31st USENIX Security Symposium (USENIX Security 22)*. 2839–2856.

[152] Janice Y Tsai, Patrick Gage Kelley, Lorrie Faith Cranor, and Norman Sadeh. 2010. Location-sharing technologies: Privacy risks and controls. *Isjlp* 6 (2010), 119.
[153] Unwired Labs. 2023. BeaconDB: An Open WiFi Location Database. https://beacondb.net. Accessed: 2026-02-16.
[154] Marno Herman Josephus van der Maas. 2021. *Protecting enclaves from side-channel attacks through physical isolation.* Ph. D. Dissertation. University of Cambridge.
[155] Xiaopang Wang, Lihui Wang, and Limin Guo. 2025. Privacy-Preserving Location-Based Services with a New Private Information Retrieval Scheme. In *2025 IEEE 24th International Conference on Trust, Security and Privacy in Computing and Communications (TrustCom).* 1649–1660. https://doi.org/10.1109/Trustcom66490.2025.00192
[156] Marius Wernke, Pavel Skvortsov, Frank Dürr, and Kurt Rothermel. 2014. A classification of location privacy attacks and approaches. *Personal and ubiquitous computing* 18, 1 (2014), 163–175.
[157] DULT WG. 2025. *Detecting Unwanted Location Trackers (Dult).* https://datatracker.ietf.org/wg/dult/about/
[158] Primal Wijesekera, Arjun Baokar, Ashkan Hosseini, Serge Egelman, David Wagner, and Konstantin Beznosov. 2015. Android permissions remystified: A field study on contextual integrity. In *24th USENIX Security Symposium (USENIX Security 15).* 499–514.
[159] Martin Woolley and A Negryn. 2025. Bluetooth core 5.2 feature overview. *Bluetooth Special Interest Group (SIG), Kirkland, WA, USA, Tech. Rep., Version 1.0. 2* (2025).
[160] Maxim Van de Wynckel and Beat Signer. 2022. A Solid-based Architecture for Decentralised Interoperable Location Data. In *WiP Proceedings of the Twelfth International Conference on Indoor Positioning and Indoor Navigation (IPIN-WiP 2022), Beijing, China, September 2022 (CEUR Workshop Proceedings)*, Hong Yuan, Dongyan Wei, Wen Li, and Antoni Pérez-Navarro (Eds.). CEUR-WS.org. https://ceur-ws.org/Vol-3248/paper11.pdf
[161] Yunming Xiao, Peizhi Liu, Ruijie Yu, Chenkai Weng, Matteo Varvello, and Aleksandar Kuzmanovic. 2025. Collusion Resistant DNS With Private Information Retrieval. *arXiv preprint arXiv:2507.20806* (2025).
[162] Vijay Kumar Yadav, Nitish Andola, Shekhar Verma, and S Venkatesan. 2022. EP2LBS: efficient privacy-preserving scheme for location-based services: VK Yadav et al. *The Journal of Supercomputing* 78, 13 (2022), 14991–15013.
[163] Zheng Yang, Chenshu Wu, Zimu Zhou, Xinglin Zhang, Xu Wang, and Yunhao Liu. 2015. Mobility increases localizability: A survey on wireless indoor localization using inertial sensors. *ACM Computing Surveys (Csur)* 47, 3 (2015), 1–34.
[164] Zhisheng Yao, Yang Peng, Yu Wang, Congan Xu, Juzhen Wang, Yun Lin, and Guan Gui. 2024. A novel radio frequency fingerprint concealment method based on IQ imbalance compensation and digital pre-distortion. *IEEE Transactions on Information Forensics and Security* 19 (2024), 7349–7361.
[165] Zhisheng Yao, Yu Wang, Guan Gui, Shiwen Mao, and Xianbin Wang. 2024. New frontier of communication security on radio frequency fingerprints concealment. *IEEE Wireless Communications* 31, 5 (2024), 156–163.
[166] Xianan Zhang, Wei Wang, Xuedou Xiao, Hang Yang, Xinyu Zhang, and Tao Jiang. 2020. Peer-to-peer localization for single-antenna devices. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 4, 3 (2020), 1–25.
[167] Chen Zhuang, Hongbo Zhao, Shan Hu, Wenquan Feng, and Rongke Liu. 2021. Cooperative positioning for V2X applications using GNSS carrier phase and UWB ranging. *IEEE Communications Letters* 25, 6 (2021), 1876–1880.
[168] Aviram Zilberman, Adi Offer, Bar Pincu, Yoni Glickshtein, Roi Kant, Oleg Brodt, Andikan Otung, Rami Puzis, Asaf Shabtai, and Yuval Elovici. 2024. A Survey on Geolocation on the Internet. *IEEE Communications Surveys & Tutorials* (2024).

A LINDDUN Privacy Threat Types

Table 5 summarizes the seven LINDDUN privacy threat types used throughout our analysis.

Table 5: LINDDUN privacy threat types.

Threat type	Description
Linking	Data or actions can be correlated to reveal additional information about one or multiple entities.
Identifying	Data or actions can be associated with a specific entity.
Non-repudiation	Data or actions can be attributed to an entity in a way that prevents plausible denial of involvement.
Detectability	It can be determined whether a particular entity is involved in an activity or present within a system.
Data disclosure	Personal data is excessively collected, stored, processed, shared, or exposed to unauthorized entities.
Unawareness / Unintervenability	An entity is not properly informed about data collection or lacks effective means to control its data.
Non-compliance	Data is processed in a manner that violates applicable regulations, risk management requirements, or privacy standards.

B OS-Level Location Disclosure Frequency on iOS

B.1 Experimental Setup

To estimate the frequency at which an iPhone's operating system contacts remote services that may reveal the user's location, we intercepted the device's network traffic using mitmproxy [55]. The mitmproxy CA certificate was installed and trusted on the iPhone to allow TLS decryption, and mitmproxy's VPN mode was enabled so that all traffic was captured regardless of the underlying network. Apple's certificate pinning prevents the interception of a subset of OS-level flows; the captured traffic therefore provides a lower bound on the frequency at which Apple could observe its users' location.

We ran two 45-minute scenarios: an outdoor walk in an urban environment and a stationary session at home. In both cases, the phone was not interacted with during the experiment, no application was granted the *Always* location authorization, and system service authorizations were left at their default values.

B.2 Results

Across both scenarios, we observed traffic to three Apple endpoints directly related to geolocation. The first, `gs-loc.apple.com`, corresponds to Apple's reference resolver, queried as part of the assisted geolocation process described in § B.2. The second, `gspe1-ssl.ls.apple.com`, is queried by the device to determine the country in which it operates based on its public IP address. The third, `gsp-ssl.ls.apple.com`, exposes a *reverse geocoding* service, which translates geographic coordinates into a human-readable location (e.g., a street address or place name), and was contacted by several system daemons, including the Wallet/Apple Pay daemon, the routine daemon underlying Siri Suggestions and Apple Intelligence, and the geolocation daemon itself.

Table 6: Number of location-related requests observed during each 45-minute scenario.

Service	Walking	Home
Reference resolution (`gs-loc.apple.com`)	34	15
Reverse geocoding (`gsp-ssl.ls.apple.com`)	2	5
IP to location (`gspe1-ssl.ls.apple.com`)	6	3
Total	**42**	**23**

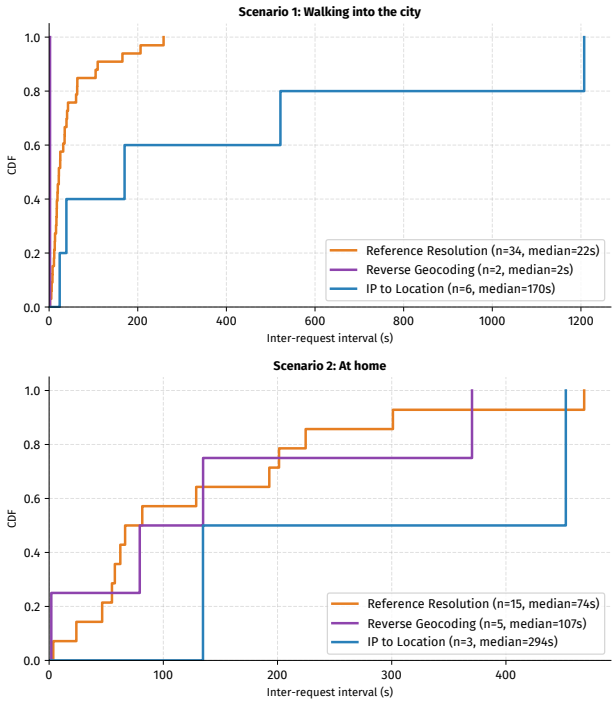

Figure 7: CDF of inter-request time by service and by scenario.

The results, summarized in Table 6 and complemented by the inter-request interval distributions in Figure 7, show that location-related queries are issued continuously throughout both scenarios, even without any user interaction, and more intensively in a non-static scenario.

These results support the claim that reference resolver queries constitute a significant and continuous channel of location data exposure, independently of any user activity.

C Example of an Apple Find My Online-Finding Request (iOS)

This appendix shows a representative intercepted exchange between an iPhone and Apple's Find My service, captured using the same set up as presented in the experiments described in Appendix B. The listing illustrates two key claims from our analysis. First, the device transmits its precise GPS coordinates (`latitude`, `longitude`, `horizontalAccuracy`, `verticalAccuracy`) in plaintext within the JSON request body. Second, the request is fully authenticated. The `Authorization` header carries a Basic token

tied to the user's Apple ID (prsId), and the device uniquely identifies itself via deviceUDID and apsToken. This means Apple can systematically link each location disclosure to a specific user account and a specific device. The response further reveals that Apple performs server-side reverse geocoding, returning a full street address, indicating that Apple's knowledge of user location extends beyond raw coordinates to human-readable place information.

C.1 Request

```
POST /fmipservice/friends/<PRS_ID>/<DEVICE_UDID>/myLocationChanged
     HTTP/1.1
Host: p<PARTITION>-fmfmobile.icloud.com
User-Agent: FindMy/<APP_VERSION> CFNetwork/<CFNETWORK_VERSION>
     Darwin/<KERNEL_VERSION>
x-apple-baa-sig-alt: <BAA_SIGNATURE_ALT>
X-Apple-I-MD: <APPLE_I_MD>
x-apple-baa-version: <VERSION>
x-apple-baa-time: <TIMESTAMP_MS>
X-Apple-Realm-Support: <VERSION>
X-MME-CLIENT-INFO: <CLIENT_INFO>
Content-Length: <CONTENT_LENGTH>
X-Apple-I-TimeZone: <TIMEZONE>
X-FMF-Model-Version: <VERSION>
X-Apple-I-Client-Time: <TIMESTAMP_UTC>
X-Apple-Find-API-Ver: <VERSION>
X-Apple-AuthScheme: <AUTH_SCHEME>
Authorization: Basic <BASIC_AUTH_TOKEN>
Accept-Language: <LOCALE>
Connection: keep-alive
X-Apple-I-MD-RINFO: <RINFO>
x-apple-baa-cert: <BAA_CERTIFICATE>
x-apple-i-device-type: <DEVICE_TYPE>
Accept: application/json
Content-Type: application/json
x-apple-baa-sig: <BAA_SIGNATURE>
Accept-Encoding: gzip, deflate, br
X-Apple-I-MD-M: <APPLE_I_MD_M>
X-Apple-I-Locale: <LOCALE>

{
  "serverContext": {
    "minCallbackIntervalInMS": <MIN_CB_MS>,
    "heartbeatIntervalInSec": <HEARTBEAT_S>,
    "authToken": "<AUTH_TOKEN>",
    "notificationToken": "<NOTIFICATION_TOKEN>",
    "minOfferLocationDurationInSec": <MIN_OFFER_S>,
    "liveAnimationInterval": <LIVE_ANIM_S>,
    "fallbackToLegacyIntervalInSec": <FALLBACK_S>,
    "reverseGeocodingThrottleDistance": <RG_THROTTLE_DIST_M>,
    "reverseGeocodingThrottle": <RG_THROTTLE_S>,
    "showAirDropImportViewOniCloudAlert": <BOOL>,
    "prsId": <PRS_ID>,
    "inaccuracyRadiusThreshold": <INACC_RADIUS_M>,
    "maxOfferLocationDurationInSec": <MAX_OFFER_S>,
    "liveTimeoutThreshold": <LIVE_TIMEOUT_S>,
    "maxCallbackIntervalInMS": <MAX_CB_MS>,
    "sendMyLocation": <BOOL>,
    "iterationNumber": <ITERATION>,
    "transientDataContext": { "<numeric_keys>": "<omitted>" },
    "callbackTimeoutIntervalInMS": <CALLBACK_TIMEOUT_MS>,
    "clientId": "<CLIENT_ID>"
  },
  "dataContext": { "<numeric_keys>": "<omitted>" },
  "clientContext": {
    "appPushModeAllowed": <BOOL>,
    "osBuild": "<OS_BUILD>",
    "currentTime": <CLIENT_TIME_MS>,
    "deviceClass": "<DEVICE_CLASS>",
    "pushMode": <BOOL>,
    "appVersion": "<APP_VERSION_SHORT>",
    "regionCode": "<REGION_CODE>",
    "apsToken": "<APS_TOKEN>",
    "countryCode": "<COUNTRY_CODE>",
    "limitedPrecision": <BOOL>,
    "osVersion": "<OS_VERSION>",
    "deviceSKU": "<DEVICE_SKU>",
    "userInactivityTimeInMS": <INACTIVITY_MS>,
    "deviceUDID": "<DEVICE_UDID>",
    "productType": "<PRODUCT_TYPE>",
    "location": {
      "longitude": <LONGITUDE>,
      "horizontalAccuracy": <HACC_METERS>,
      "verticalAccuracy": <VACC_METERS>,
      "latitude": <LATITUDE>
    }
  }
}
```

C.2 Response

```
HTTP/1.1 <STATUS_CODE> <STATUS_MSG>
Server: AppleHttpServer/<SERVER_ID>
Date: <DATE>
Content-Type: application/json; charset=UTF-8
Transfer-Encoding: chunked
Connection: keep-alive
Set-Cookie: <SESSION_COOKIE>
X-Responding-Instance: <SERVER_INSTANCE>
X-Robots-Tag: noindex, noarchive, nosnippet, nofollow
X-Responding-Server: <SERVER_NAME>
X-Responding-Partition: <PARTITION>
X-Apple-Loc-Src: <INT>
Content-Encoding: gzip
Strict-Transport-Security: <HSTS_POLICY>
X-Apple-User-Partition: <USER_PARTITION>
Via: <VIA>
X-Apple-Request-UUID: <REQUEST_UUID>
Access-Control-Expose-Headers: X-Apple-Request-UUID,Via
Alt-Svc: <ALT_SVC>
X-Apple-Edge-Response-Time: <RESPONSE_TIME_MS>

{
  "location": {
    "altitude": <FLOAT>,
    "isInaccurate": <BOOL>,
    "address": {
      "formattedAddressLines": ["<STREET_ADDRESS>", "<POSTAL_CODE> <
      CITY>", "<COUNTRY>"],
      "country": "<COUNTRY>",
      "streetName": "<STREET_NAME>",
      "streetAddress": "<STREET_ADDRESS>",
      "countryCode": "<COUNTRY_CODE>",
      "locality": "<CITY>",
      "administrativeArea": "<REGION>"
    },
    "locSource": <NULL>,
    "secureLocation": <NULL>,
    "secureLocationTs": <INT>,
    "latitude": <LATITUDE>,
    "shortAddress": <NULL>,
    "floorLevel": <INT>,
    "horizontalAccuracy": <HACC_METERS>,
    "labels": <ARRAY>,
    "verticalAccuracy": <VACC_METERS>,
    "batteryStatus": <NULL>,
    "locationId": "<LOCATION_ID>",
    "fullAddress": <NULL>,
    "locationTimestamp": <INT>,
    "locationMode": <NULL>,
    "longitude": <LONGITUDE>,
    "timestamp": <TIMESTAMP_MS>
  },
  "accuracyThreshold": <INT>,
  "distanceThreshold": <INT>,
  "statusCode": "<STATUS_CODE>"
}
```

D Summary of the Privacy Threats

Table 7: Taxonomy of Privacy Attacks in Geolocation Ecosystems

Attack	Adv.	Affected ecosystem	Consequences	LINDDUN
Ecosystem Architecture				
Forced use of a specific reference resolver and crowdsourcing exposes user position to the provider, linkable to a user account or IP address (§ 4.2).	A2	OS-Level: Apple, Android, Windows	Live location over time.	Li, In, Dd, Un
APIs (Google Awareness, Windows CivicAddress) transmit user position to the provider, linkable to IP and application identifier (§ 4.3).	A2	OS-level: Android, Windows	Live location over time; linking client location to used application.	Li, In
Online finding services are not E2EE and cannot be disabled while OFN is active (§ 5.2).	A2	OFN: Apple Find My, Google Find Hub	Live location accessible to the provider.	In, Li, Un
Cross-service inference by a single provider operating multiple subsystems bypasses report E2EE (§ 5.3).	A2	OFN: Apple Find My, Google Find Hub	Live location over time.	Li, In, Un
MITM during key exchange for location sharing (§ 5.4).	A2	OFN: Google Find Hub	Live location over time.	Dd, Li
Online IP-geolocation queries reveal which services are used by which IP (§ 6.5).	A3	Network-based: any IP-geolocation provider	Linking used services to IP addresses.	Li
Ecosystem Architecture and Communication Protocol				
Roaming agreements let foreign operators observe and locate roaming subscribers across their coverage (§ 6.5).	A3	Network-based: any MNO	Linking MNO users to location data.	Li, In
The network provider (ISP, MNO, Sigfox) passively infers customer locations from connectivity signals (§ 6.3).	A2	Networks: any managed network	Location inference.	Li, In, Un
Communication Protocols				
Crafted broadcast identifiers (BSSID, Cell-ID) or GNSS signals falsify the user's computed position (§ 4.5).	A3	OS-level: any OS using GNSS, Wi-Fi, or cellular	Incorrect but credible location records.	Id, Un
IP addresses enable passive location inference of clients (§ 6.3).	A3	Networks: IP networks	Coarse-grained location inference.	Li, Un
Wireless protocols (Wi-Fi, LPWAN, cellular) expose persistent network identifiers (e.g., MAC) to nearby observers (§ 6.3).	A3	Networks: any wireless network	Presence detection within an area.	Li, In, Un
RFF lets attackers track wireless devices via transmitter imperfections (§ 6.3).	A3	Networks: any wireless network	Presence detection within an area.	Li, In, Un
Design Vulnerabilities				
Local retention of location history exposes records on device compromise (§ 4.3).	A1	OS-level: Apple, Android	Access to location records.	Li, Dd
Unauthenticated access to Apple's WPS database enables large-scale extraction (§ 4.2).	A3	OS-level: Apple	Public exposure of millions of Wi-Fi access point locations.	In, Li, Dd
UI design flaws mislead users into granting location permissions (§ 4.4).	A1	OS-level: Android	Live location over time.	Un
Cloud retention of location history exposes records on provider compromise (§ 4.3, § 6.4).	A3	OS-level: Apple, Android; Network-based: Sigfox, 1NCE	Access to location records.	Li, Dd
Unencrypted report metadata (timestamp, accuracy, frequency) reveals coarse-grained location and activity patterns (§ 5.3).	A2	OFN: Apple Find My, Google Find Hub	Location and activity inference.	In, Li
Reports from the same finder can be linked via report timestamps (§ 5.3).	A2	OFN: Apple Find My, Google Find Hub	Tracker co-location inference; social graph construction.	Li, In
Exploit allows continued receipt of reports after location sharing is revoked (§ 5.4).	A3	OFN: Find Hub	Live location over time.	Dd, Li, Un
Unauthorized access to cryptographic secrets stored on owner devices grants access to reports (§ 5.4).	A3	OFN: Apple Find My (patched)	Live location over time.	Dd, Li
Consent and Transparency Failures				
Applications bypass location permissions by exploiting other signals such as sensors, BLE, or local network (§ 4.4).	A1	OS-level: Android, iOS	Live location over time.	Un
Complex privacy policies and ToS prevent users from understanding how their data are processed (§ 4.4).	A2	OS-level: Android, Apple, Windows	User unawareness.	Un
Provider deviates from declared policies on location data usage (§ 4.4).	A2	OS-level: Android	User unawareness.	Nc, Un

Adversary: A1 — Target-based; A2 — Trusted External; A3 — Untrusted External. LINDDUN: *Li* — Linking; *In* — Identifying; *Nr* — Non-repudiation; *De* — Detection; *Dd* — Data Disclosure; *Un* — Unawareness/Unintervenability; *Nc* — Non-compliance.

E Example of our Data Flow Diagrams

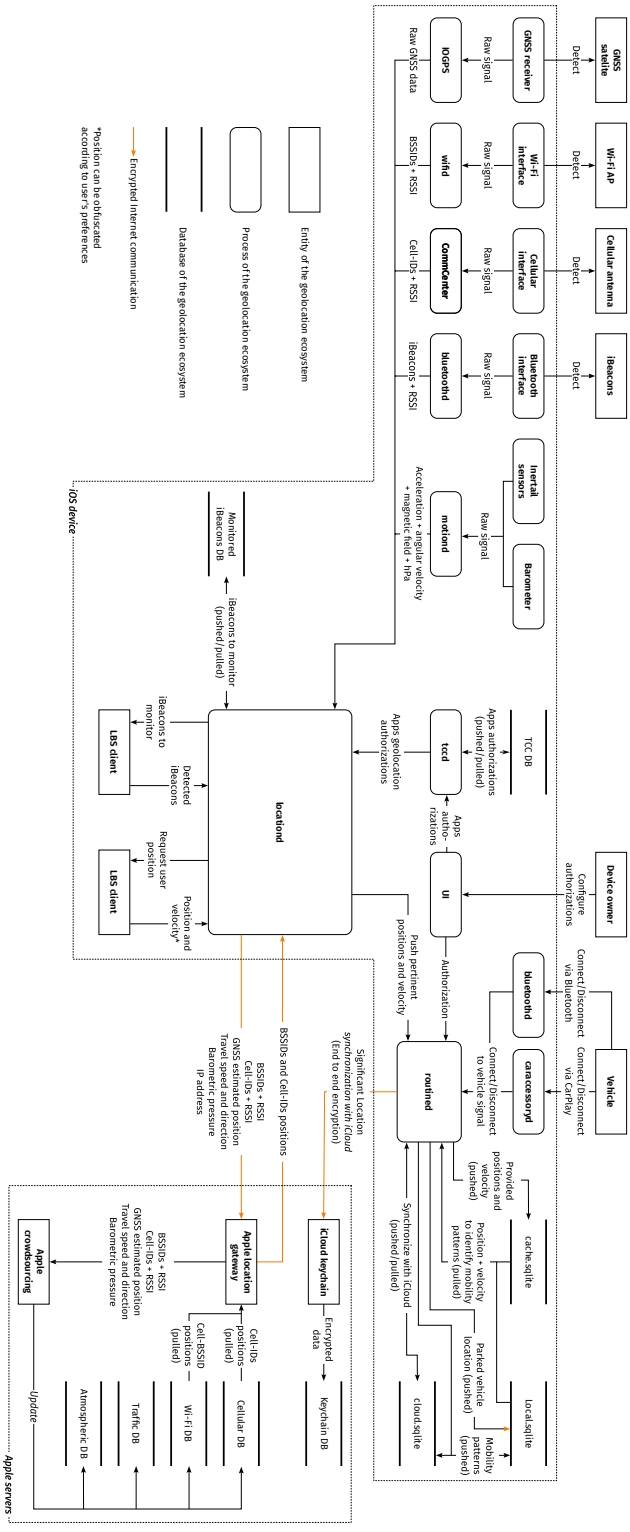

Figure 8: Data Flow Diagrams for iOS Geolocation Ecosystem

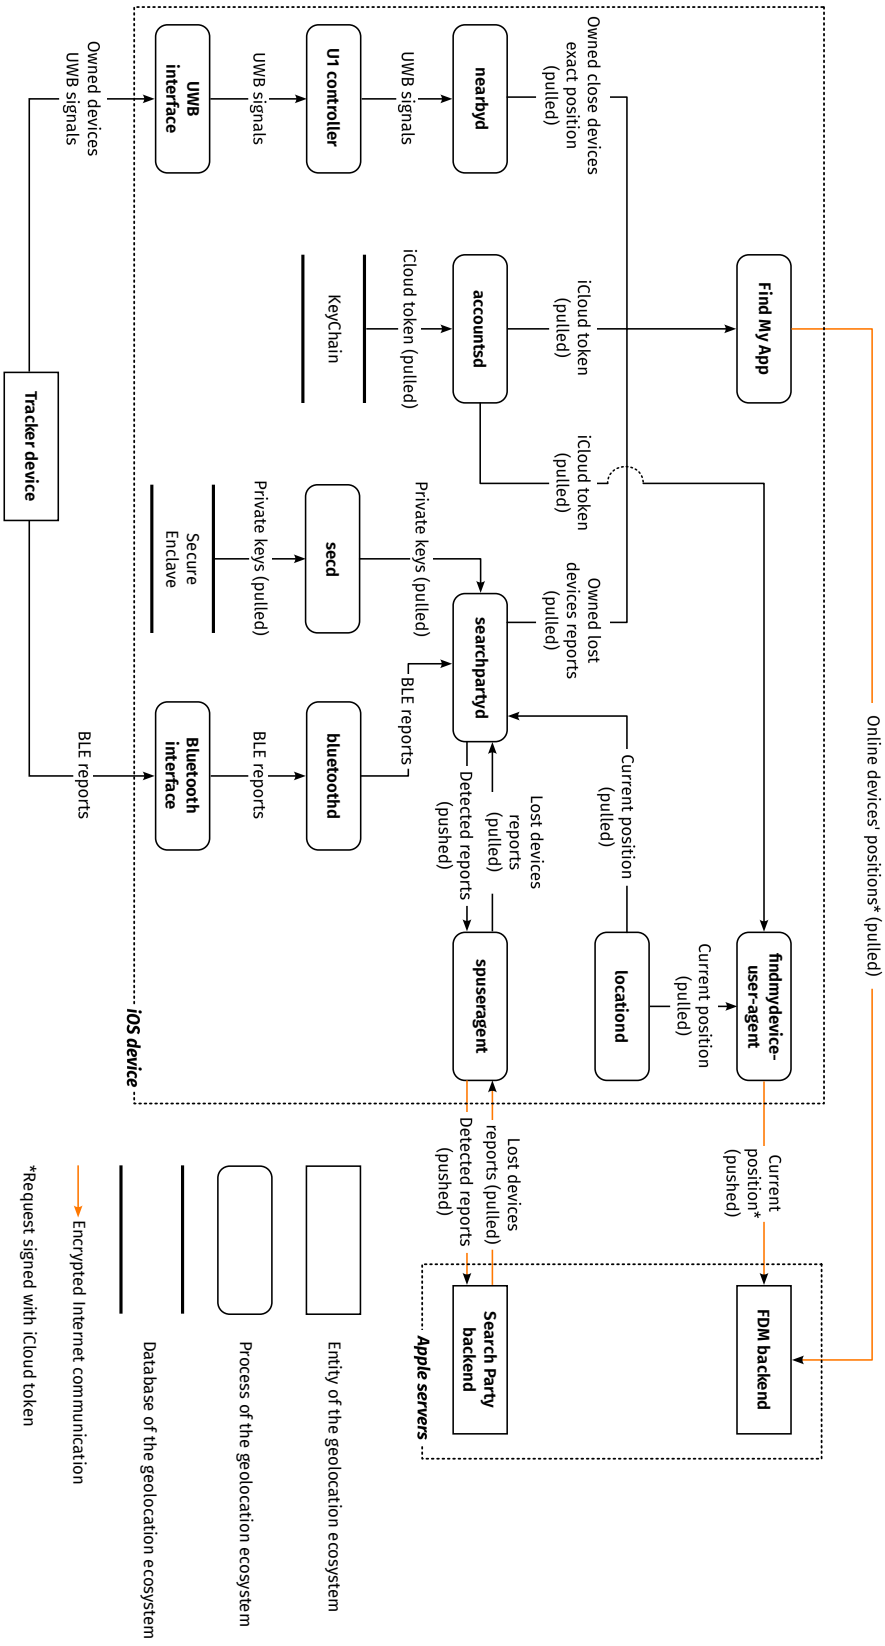

Figure 9: Data Flow Diagrams for Find My Ecosystem

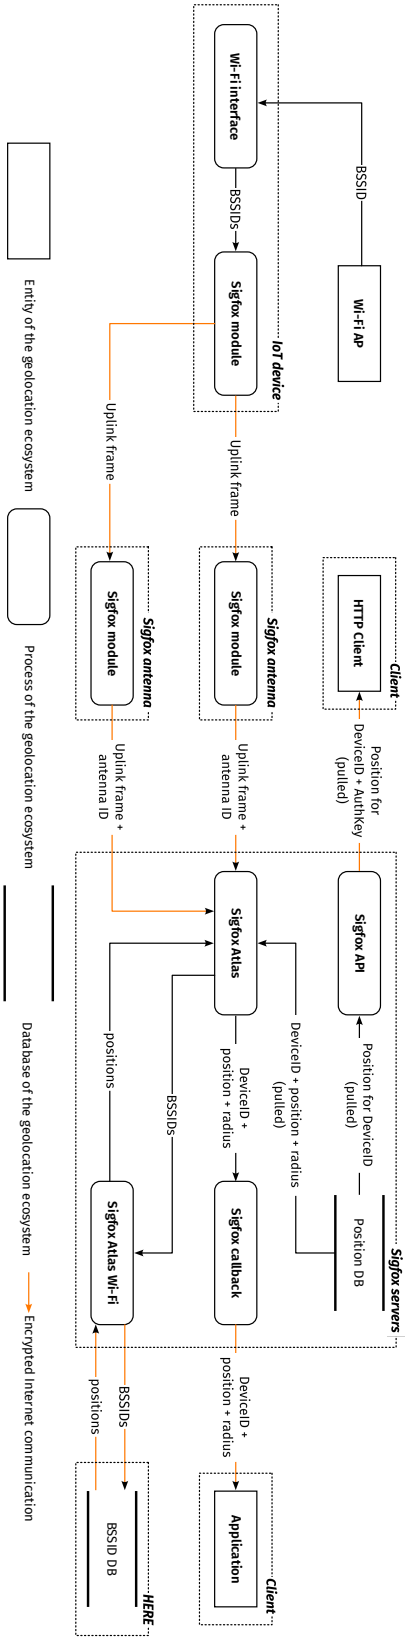

Figure 10: Data Flow Diagrams for Sigfox Atlas Ecosystem

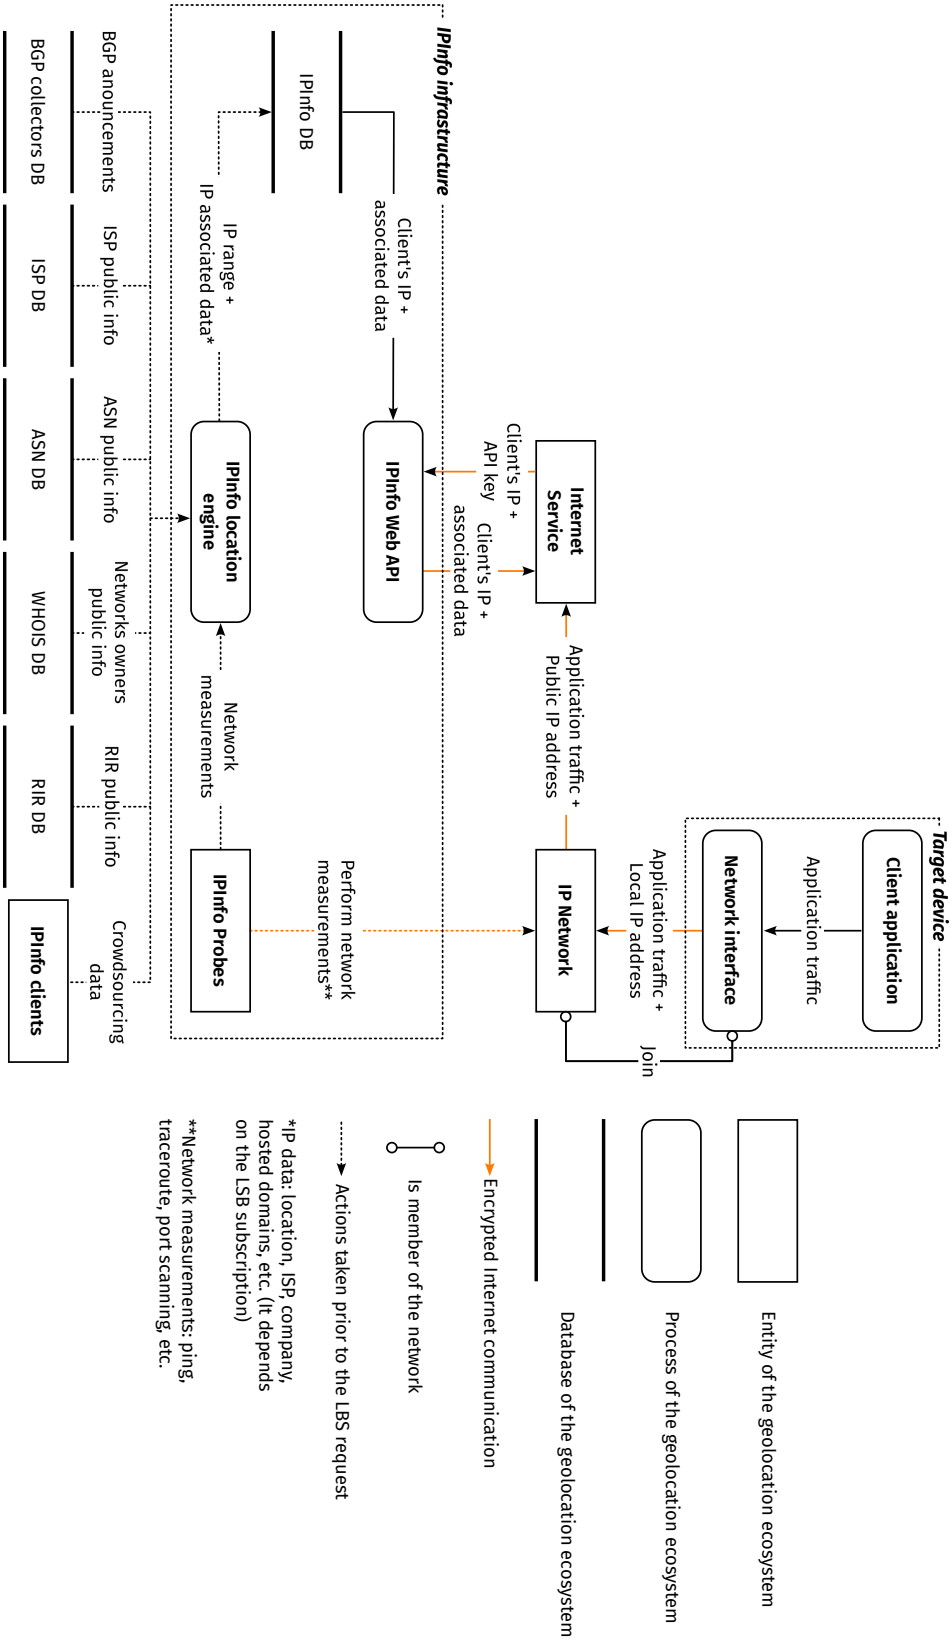

Figure 11: Data Flow Diagrams for IPinfo Ecosystem

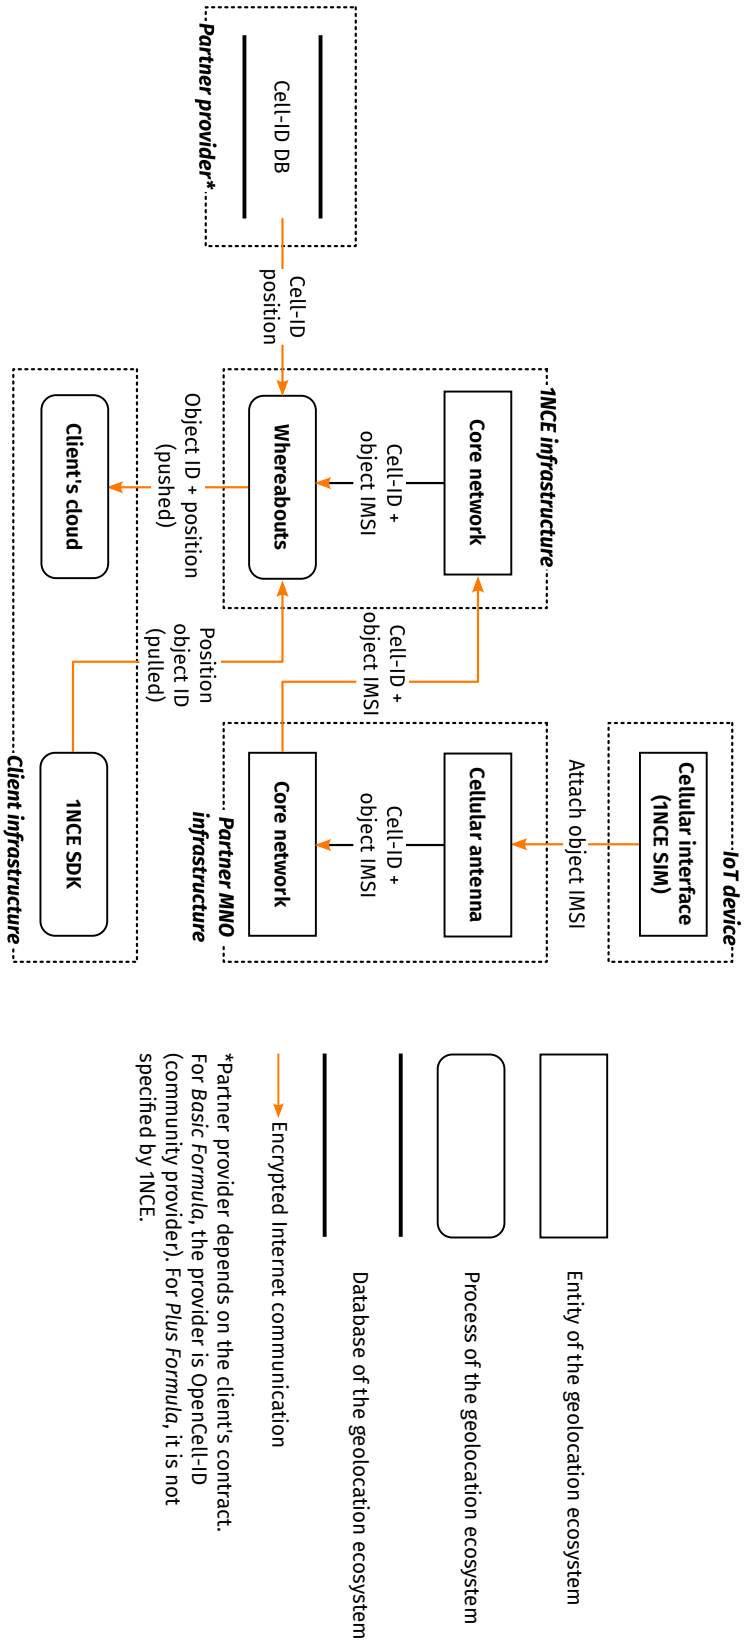

Figure 12: Data Flow Diagrams for 1NCE Ecosystem