

Internet Engineering Task Force (IETF)
Request for Comments: 7378
Category: Informational
ISSN: 2070-1721

H. Tschofenig
Independent
H. Schulzrinne
Columbia University
B. Aboba, Ed.
Microsoft Corporation
December 2014

Trustworthy Location

Abstract

The trustworthiness of location information is critically important for some location-based applications, such as emergency calling or roadside assistance.

This document describes threats to conveying location, particularly for emergency calls, and describes techniques that improve the reliability and security of location information. It also provides guidelines for assessing the trustworthiness of location information.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc7378>.

Copyright Notice

Copyright (c) 2014 IETF Trust and the persons identified as the document authors. All rights reserved.

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

Table of Contents

1. Introduction	3
1.1. Terminology	3
1.2. Emergency Services Architecture	5
2. Threat Models	8
2.1. Existing Work	8
2.2. Adversary Model	9
2.3. Location Spoofing	10
2.4. Identity Spoofing	11
3. Mitigation Techniques	11
3.1. Signed Location-by-Value	12
3.2. Location-by-Reference	15
3.3. Proxy-Added Location	18
4. Location Trust Assessment	20
5. Security Considerations	23
6. Privacy Considerations	24
7. Informative References	26
Acknowledgments	30
Authors' Addresses	30

1. Introduction

Several public and commercial services need location information to operate. This includes emergency services (such as fire, ambulance, and police) as well as commercial services such as food delivery and roadside assistance.

For circuit-switched calls from landlines, as well as for Voice over IP (VoIP) services that only support emergency service calls from stationary Devices, location provided to the Public Safety Answering Point (PSAP) is determined from a lookup using the calling telephone number. As a result, for landlines or stationary VoIP, spoofing of caller identification can result in the PSAP incorrectly determining the caller's location. Problems relating to calling party number and Caller ID assurance have been analyzed by the Secure Telephone Identity Revisited [STIR] working group as described in "Secure Telephone Identity Problem Statement and Requirements" [RFC7340]. In addition to the work underway in STIR, other mechanisms exist for validating caller identification. For example, as noted in [EENA], one mechanism for validating caller identification information (as well as the existence of an emergency) is for the PSAP to call the user back, as described in [RFC7090].

Given the existing work on caller identification, this document focuses on the additional threats that are introduced by the support of IP-based emergency services in nomadic and mobile Devices, in which location may be conveyed to the PSAP within the emergency call. Ideally, a call taker at a PSAP should be able to assess, in real time, the level of trust that can be placed on the information provided within a call. This includes automated location conveyed along with the call and location information communicated by the caller, as well as identity information relating to the caller or the Device initiating the call. Where real-time assessment is not possible, it is important to be able to determine the source of the call in a post-incident investigation, so as to be able to enforce accountability.

This document defines terminology (including the meaning of "trustworthy location") in Section 1.1, reviews existing work in Section 1.2, describes threat models in Section 2, outlines potential mitigation techniques in Section 3, covers trust assessment in Section 4, and discusses security considerations in Section 5.

1.1. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

We use the definitions of "Internet Access Provider (IAP)", "Internet Service Provider (ISP)", and "Voice Service Provider (VSP)" found in "Requirements for Emergency Context Resolution with Internet Technologies" [RFC5012].

[EENA] defines a "hoax call" as follows: "A false or malicious call is when a person deliberately telephones the emergency services and tells them there is an emergency when there is not."

The definitions of "Device", "Target", and "Location Information Server" (LIS) are taken from "An Architecture for Location and Location Privacy in Internet Applications" [RFC6280], Section 7.

The term "Device" denotes the physical device, such as a mobile phone, PC, or embedded microcontroller, whose location is tracked as a proxy for the location of a Target.

The term "Target" denotes an individual or other entity whose location is sought in the Geopriv architecture [RFC6280]. In many cases, the Target will be the human user of a Device, or it may be an object such as a vehicle or shipping container to which a Device is attached. In some instances, the Target will be the Device itself. The Target is the entity whose privacy the architecture described in [RFC6280] seeks to protect.

The term "Location Information Server" denotes an entity responsible for providing Devices within an access network with information about their own locations. A Location Information Server uses knowledge of the access network and its physical topology to generate and distribute location information to Devices.

The term "location determination method" refers to the mechanism used to determine the location of a Target. This may be something employed by a LIS or by the Target itself. It specifically does not refer to the location configuration protocol (LCP) used to deliver location information to either the Target or the Recipient. This term is reused from "GEOPRIV Presence Information Data Format Location Object (PIDF-LO) Usage Clarification, Considerations, and Recommendations" [RFC5491].

The term "source" is used to refer to the LIS, node, or Device from which a Recipient (Target or third party) obtains location information.

Additionally, the terms "location-by-value" (LbyV), "location-by-reference" (LbyR), "Location Configuration Protocol", "Location Dereference Protocol", and "Location Uniform Resource Identifier" (URI) are reused from "Requirements for a Location-by-Reference Mechanism" [RFC5808].

"Trustworthy Location" is defined as location information that can be attributed to a trusted source, has been protected against modification in transmit, and has been assessed as trustworthy.

"Location Trust Assessment" refers to the process by which the reliability of location information can be assessed. This topic is discussed in Section 4.

"Identity Spoofing" occurs when the attacker forges or obscures their identity so as to prevent themselves from being identified as the source of the attack. One class of identity spoofing attack involves the forging of call origin identification.

The following additional terms apply to location spoofing (Section 2.3):

With "Place Shifting", attackers construct a Presence Information Data Format Location Object (PIDF-LO) for a location other than where they are currently located. In some cases, place shifting can be limited in range (e.g., within the coverage area of a particular cell tower).

"Time Shifting" occurs when the attacker uses or reuses location information that was valid in the past but is no longer valid because the attacker has moved.

"Location Theft" occurs when the attacker captures a Target's location information (possibly including a signature) and presents it as their own. Location theft can occur in a single instance or may be continuous (e.g., where the attacker has gained control over the victim's Device). Location theft may also be combined with time shifting to present someone else's location information after the original Target has moved.

1.2. Emergency Services Architecture

This section describes how location is utilized in the Internet Emergency Services Architecture, as well as the existing work on the problem of hoax calls.

1.2.1. Location

The Internet architecture for emergency calling is described in "Framework for Emergency Calling Using Internet Multimedia" [RFC6443]. Best practices for utilizing the architecture to make emergency calls are described in "Best Current Practice for Communications Services in Support of Emergency Calling" [RFC6881].

As noted in "An Architecture for Location and Location Privacy in Internet Applications" [RFC6280], Section 6.3:

there are three critical steps in the placement of an emergency call, each involving location information:

1. Determine the location of the caller.
2. Determine the proper Public Safety Answering Point (PSAP) for the caller's location.
3. Send a SIP INVITE message, including the caller's location, to the PSAP.

The conveyance of location information within the Session Initiation Protocol (SIP) is described in "Location Conveyance for the Session Initiation Protocol" [RFC6442]. Conveyance of location-by-value (LbyV) as well as conveyance of location-by-reference (LbyR) are supported. Section 7 of [RFC6442] ("Security Considerations") discusses privacy, authentication, and integrity concerns relating to conveyed location. This includes discussion of transmission-layer security for confidentiality and integrity protection of SIP, as well as (undeployed) end-to-end security mechanisms for protection of location information (e.g., S/MIME). Regardless of whether transmission-layer security is utilized, location information may be available for inspection by an intermediary that -- if it decides that the location value is unacceptable or insufficiently accurate -- may send an error indication or replace the location, as described in [RFC6442], Section 3.4.

Although the infrastructure for location-based routing described in [RFC6443] was developed for use in emergency services, [RFC6442] supports conveyance of location within non-emergency calls as well as emergency calls. Section 1 of "Implications of 'retransmission-allowed' for SIP Location Conveyance" [RFC5606] describes the overall architecture, as well as non-emergency usage scenarios (note: the [LOC-CONVEY] citation in the quote below refers to the document later published as [RFC6442]):

The Presence Information Data Format for Location Objects (PIDF-LO [RFC4119]) carries both location information (LI) and policy information set by the Rule Maker, as is stipulated in [RFC3693]. The policy carried along with LI allows the Rule Maker to restrict, among other things, the duration for which LI will be retained by recipients and the redistribution of LI by recipients.

The Session Initiation Protocol [RFC3261] is one proposed Using Protocol for PIDF-LO. The conveyance of PIDF-LO within SIP is specified in [LOC-CONVEY]. The common motivation for providing LI in SIP is to allow location to be considered in routing the SIP message. One example use case would be emergency services, in which the location will be used by dispatchers to direct the response. Another use case might be providing location to be used by services associated with the SIP session; a location associated with a call to a taxi service, for example, might be used to route to a local franchisee of a national service and also to route the taxi to pick up the caller.

1.2.2. Hoax Calls

Hoax calls have been a problem for emergency services dating back to the time of street corner call boxes. As the European Emergency Number Association (EENA) has noted [EENA]:

False emergency calls divert emergency services away from people who may be in life-threatening situations and who need urgent help. This can mean the difference between life and death for someone in trouble.

EENA [EENA] has attempted to define terminology and describe best current practices for dealing with false emergency calls. Reducing the number of hoax calls represents a challenge, since emergency services authorities in most countries are required to answer every call (whenever possible). Where the caller cannot be identified, the ability to prosecute is limited.

A particularly dangerous form of hoax call is "swatting" -- a hoax emergency call that draws a response from law enforcement prepared for a violent confrontation (e.g., a fake hostage situation that results in the dispatching of a "Special Weapons And Tactics" (SWAT) team). In 2008, the Federal Bureau of Investigation (FBI) issued a warning [Swatting] about an increase in the frequency and sophistication of these attacks.

Many documented cases of "swatting" (also sometimes referred to as "SWATing") involve not only the faking of an emergency but also falsification or obfuscation of identity [Swatting] [SWATing]. There are a number of techniques by which hoax callers attempt to avoid identification, and in general, the ability to identify the caller appears to influence the incidence of hoax calls.

Where a Voice Service Provider allows the caller to configure its outbound caller identification without checking it against the authenticated identity, forging caller identification is trivial. Similarly, where an attacker can gain entry to a Private Branch Exchange (PBX), they can then subsequently use that access to launch a denial-of-service attack against the PSAP or make fraudulent emergency calls. Where emergency calls have been allowed from handsets lacking a subscriber identification module (SIM) card, so-called non-service initialized (NSI) handsets, or where ownership of the SIM card cannot be determined, the frequency of hoax calls has often been unacceptably high [TASMANIA] [UK] [SA].

However, there are few documented cases of hoax calls that have arisen from conveyance of untrustworthy location information within an emergency call, which is the focus of this document.

2. Threat Models

This section reviews existing analyses of the security of emergency services, threats to geographic location privacy, threats relating to spoofing of caller identification, and threats related to modification of location information in transit. In addition, the threat model applying to this work is described.

2.1. Existing Work

"An Architecture for Location and Location Privacy in Internet Applications" [RFC6280] describes an architecture for privacy-preserving location-based services in the Internet, focusing on authorization, security, and privacy requirements for the data formats and protocols used by these services.

In Section 5 of [RFC6280] ("An Architecture for Location and Location Privacy in Internet Applications"), mechanisms for ensuring the security of the location distribution chain are discussed; these include mechanisms for hop-by-hop confidentiality and integrity protection as well as end-to-end assurance.

"Geopriv Requirements" [RFC3693] focuses on the authorization, security, and privacy requirements of location-dependent services, including emergency services. Section 8 of [RFC3693] includes discussion of emergency services authentication (Section 8.3), and issues relating to identity and anonymity (Section 8.4).

"Threat Analysis of the Geopriv Protocol" [RFC3694] describes threats against geographic location privacy, including protocol threats, threats resulting from the storage of geographic location data, and threats posed by the abuse of information.

"Security Threats and Requirements for Emergency Call Marking and Mapping" [RFC5069] reviews security threats associated with the marking of signaling messages and the process of mapping locations to Universal Resource Identifiers (URIs) that point to PSAPs. RFC 5069 describes attacks on the emergency services system, such as attempting to deny system services to all users in a given area, to gain fraudulent use of services and to divert emergency calls to non-emergency sites. In addition, it describes attacks against individuals, including attempts to prevent an individual from receiving aid, or to gain information about an emergency, as well as attacks on emergency services infrastructure elements, such as mapping discovery and mapping servers.

"Secure Telephone Identity Threat Model" [RFC7375] analyzes threats relating to impersonation and obscuring of calling party numbers, reviewing the capabilities available to attackers, and the scenarios in which attacks are launched.

2.2. Adversary Model

To provide a structured analysis, we distinguish between three adversary models:

External adversary model: The end host, e.g., an emergency caller whose location is going to be communicated, is honest, and the adversary may be located between the end host and the location server or between the end host and the PSAP. None of the emergency service infrastructure elements act maliciously.

Malicious infrastructure adversary model: The emergency call routing elements, such as the Location Information Server (LIS), the Location-to-Service Translation (LoST) infrastructure (which is used for mapping locations to PSAP addresses), or call routing elements, may act maliciously.

Malicious end host adversary model: The end host itself acts maliciously, whether the owner is aware of this or the end host is acting under the control of a third party.

Since previous work describes attacks against infrastructure elements (e.g., location servers, call route servers, mapping servers) or the emergency services IP network, as well as threats from attackers attempting to snoop location in transit, this document focuses on the threats arising from end hosts providing false location information within emergency calls (the malicious end host adversary model).

Since the focus is on malicious hosts, we do not cover threats that may arise from attacks on infrastructure that hosts depend on to obtain location. For example, end hosts may obtain location from civilian GPS, which is vulnerable to spoofing [GPSCounter], or from third-party Location Service Providers (LSPs) that may be vulnerable to attack or may not provide location accuracy suitable for emergency purposes.

Also, we do not cover threats arising from inadequate location infrastructure. For example, the LIS or end host could base its location determination on a stale wiremap or an inaccurate access point location database, leading to an inaccurate location estimate. Similarly, a Voice Service Provider (VSP) (and, indirectly, a LIS) could utilize the wrong identity (such as an IP address) for location lookup, thereby providing the end host with misleading location information.

2.3. Location Spoofing

Where location is attached to the emergency call by an end host, the end host can fabricate a PIDF-LO and convey it within an emergency call. The following represent examples of location spoofing:

Place shifting: Mallory, the adversary, pretends to be at an arbitrary location.

Time shifting: Mallory pretends to be at a location where she was a while ago.

Location theft: Mallory observes or obtains Alice's location and replays it as her own.

2.4. Identity Spoofing

While this document does not focus on the problems created by determination of location based on spoofed caller identification, the ability to ascertain identity is important, since the threat of punishment reduces hoax calls. As an example, calls from pay phones are subject to greater scrutiny by the call taker.

With calls originating on an IP network, at least two forms of identity are relevant, with the distinction created by the split between the IAP and the VSP:

- (a) network access identity such as might be determined via authentication (e.g., using the Extensible Authentication Protocol (EAP) [RFC3748]);
- (b) caller identity, such as might be determined from authentication of the emergency caller at the VoIP application layer.

If the adversary did not authenticate itself to the VSP, then accountability may depend on verification of the network access identity. However, the network access identity may also not have been authenticated, such as in the case where an open IEEE 802.11 Access Point is used to initiate a hoax emergency call. Although endpoint information such as the IP address or Media Access Control (MAC) address may have been logged, tying this back to the Device owner may be challenging.

Unlike the existing telephone system, VoIP emergency calls can provide an identity that need not necessarily be coupled to a business relationship with the IAP, ISP, or VSP. However, due to the time-critical nature of emergency calls, multi-layer authentication is undesirable. Thus, in most cases, only the Device placing the call will be able to be identified. Furthermore, deploying additional credentials for emergency service purposes (such as certificates) increases costs, introduces a significant administrative overhead, and is only useful if widely deployed.

3. Mitigation Techniques

The sections that follow present three mechanisms for mitigating the threats presented in Section 2:

1. Signed location-by-value (Section 3.1), which provides for authentication and integrity protection of the PIDF-LO. There is only an expired straw-man proposal for this mechanism [Loc-Dependability]; thus, as of the time of this writing this mechanism is not suitable for deployment.

- 2. Location-by-reference (Section 3.2), which enables location to be obtained by the PSAP directly from the location server, over a confidential and integrity-protected channel, avoiding modification by the end host or an intermediary. This mechanism is specified in [RFC6753].
- 3. Proxy-added location (Section 3.3), which protects against location forgery by the end host. This mechanism is specified in [RFC6442].

3.1. Signed Location-by-Value

With location signing, a location server signs the location information before it is sent to the Target. The signed location information is then sent to the Location Recipient, who verifies it.

Figure 1 shows the communication model with the Target requesting signed location in step (a); the location server returns it in step (b), and it is then conveyed to the Location Recipient, who verifies it (step (c)). For SIP, the procedures described in "Location Conveyance for the Session Initiation Protocol" [RFC6442] are applicable for location conveyance.

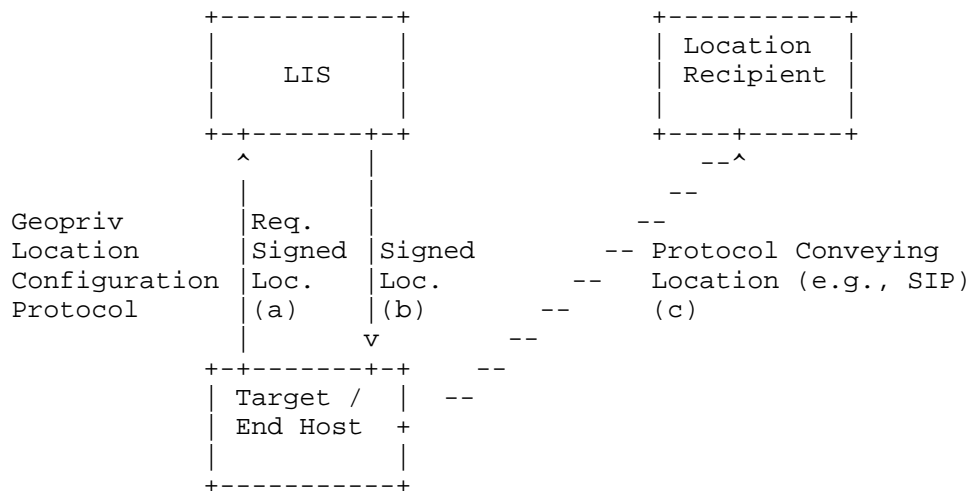


Figure 1: Location Signing

A straw-man proposal for location signing is provided in "Digital Signature Methods for Location Dependability" [Loc-Dependability]. Note that since [Loc-Dependability] is no longer under development, location signing cannot be considered deployable at the time of this writing.

In order to limit replay attacks, that proposal calls for the addition of a "validity" element to the PIDF-LO, including a "from" sub-element containing the time that location information was validated by the signer, as well as an "until" sub-element containing the last time that the signature can be considered valid.

One of the consequences of including an "until" element is that even a stationary Target would need to periodically obtain a fresh PIDF-LO, or incur the additional delay of querying during an emergency call.

Although privacy-preserving procedures may be disabled for emergency calls, by design, PIDF-LO objects limit the information available for real-time attribution. As noted in [RFC5985], Section 6.6:

The LIS MUST NOT include any means of identifying the Device in the PIDF-LO unless it is able to verify that the identifier is correct and inclusion of identity is expressly permitted by a Rule Maker. Therefore, PIDF parameters that contain identity are either omitted or contain unlinked pseudonyms [RFC3693]. A unique, unlinked presentity URI SHOULD be generated by the LIS for the mandatory presence "entity" attribute of the PIDF document.

Optional parameters such as the "contact" and "deviceID" elements [RFC4479] are not used.

Also, the Device referred to in the PIDF-LO may not necessarily be the same entity conveying the PIDF-LO to the PSAP. As noted in [RFC6442], Section 1:

In no way does this document assume that the SIP user agent client that sends a request containing a location object is necessarily the Target. The location of a Target conveyed within SIP typically corresponds to that of a Device controlled by the Target, for example, a mobile phone, but such Devices can be separated from their owners, and moreover, in some cases, the user agent may not know its own location.

Without the ability to tie the Target identity to the identity asserted in the SIP message, it is possible for an attacker to cut and paste a PIDF-LO obtained by a different Device or user into a SIP INVITE and send this to the PSAP. This cut-and-paste attack could succeed even when a PIDF-LO is signed or when [RFC4474] is implemented.

To address location-spoofing attacks, [Loc-Dependability] proposes the addition of an "identity" element that could include a SIP URI (enabling comparison against the identity asserted in the SIP headers) or an X.509v3 certificate. If the Target was authenticated by the LIS, an "authenticated" attribute is added. However, because the inclusion of an "identity" element could enable location tracking, a "hash" element is also proposed that could instead contain a hash of the content of the "identity" element. In practice, such a hash would not be much better for real-time validation than a pseudonym.

Location signing cannot deter attacks in which valid location information is provided. For example, an attacker in control of compromised hosts could launch a denial-of-service attack on the PSAP by initiating a large number of emergency calls, each containing valid signed location information. Since the work required to verify the location signature is considerable, this could overwhelm the PSAP infrastructure.

However, while DDoS attacks are unlikely to be deterred by location signing, accurate location information would limit the subset of compromised hosts that could be used for an attack, as only hosts within the PSAP serving area would be useful in placing emergency calls.

Location signing is also difficult when the host obtains location via mechanisms such as GPS, unless trusted computing approaches, with tamper-proof GPS modules, can be applied. Otherwise, an end host can pretend to have GPS, and the Recipient will need to rely on its ability to assess the level of trust that should be placed in the end host location claim.

Even though location-signing mechanisms have not been standardized, [NENA-i2], Section 4.7 includes operational recommendations relating to location signing:

Location configuration and conveyance requirements are described in NENA 08-752[27], but guidance is offered here on what should be considered when designing mechanisms to report location:

1. The location object should be digitally signed.
2. The certificate for the signer (LIS operator) should be rooted in VESA. For this purpose, VPC and ERDB operators should issue certificates to LIS operators.
3. The signature should include a timestamp.

4. Where possible, the Location Object should be refreshed periodically, with the signature (and thus the timestamp) being refreshed as a consequence.
5. Antispoofing mechanisms should be applied to the Location Reporting method.

(Note: The term "Valid Emergency Services Authority" (VESA) refers to the root certificate authority. "VPC" stands for VoIP Positioning Center, and "ERDB" stands for the Emergency Service Zone Routing Database.)

As noted above, signing of location objects implies the development of a trust hierarchy that would enable a certificate chain provided by the LIS operator to be verified by the PSAP. Rooting the trust hierarchy in the VESA can be accomplished either by having the VESA directly sign the LIS certificates or by the creation of intermediate Certificate Authorities (CAs) certified by the VESA, which will then issue certificates to the LIS. In terms of the workload imposed on the VESA, the latter approach is highly preferable. However, this raises the question of who would operate the intermediate CAs and what the expectations would be.

In particular, the question arises as to the requirements for LIS certificate issuance, and how they would compare to requirements for issuance of other certificates such as a Secure Socket Layer/Transport Layer Security (SSL/TLS) web certificate.

3.2. Location-by-Reference

Location-by-reference was developed so that end hosts can avoid having to periodically query the location server for up-to-date location information in a mobile environment. Additionally, if operators do not want to disclose location information to the end host without charging them, location-by-reference provides a reasonable alternative. Also, since location-by-reference enables the PSAP to directly contact the location server, it avoids potential attacks by intermediaries.

As noted in "A Location Dereference Protocol Using HTTP-Enabled Location Delivery (HELD)" [RFC6753], a location reference can be obtained via HELD [RFC5985]. In addition, "Location Configuration Extensions for Policy Management" [RFC7199] extends location configuration protocols such as HELD to provide hosts with a reference to the rules that apply to a location-by-reference so that the host can view or set these rules.

Figure 2 shows the communication model with the Target requesting a location reference in step (a); the location server returns the reference and, potentially, the policy in step (b), and it is then conveyed to the Location Recipient in step (c). The Location Recipient needs to resolve the reference with a request in step (d). Finally, location information is returned to the Location Recipient afterwards. For location conveyance in SIP, the procedures described in [RFC6442] are applicable.

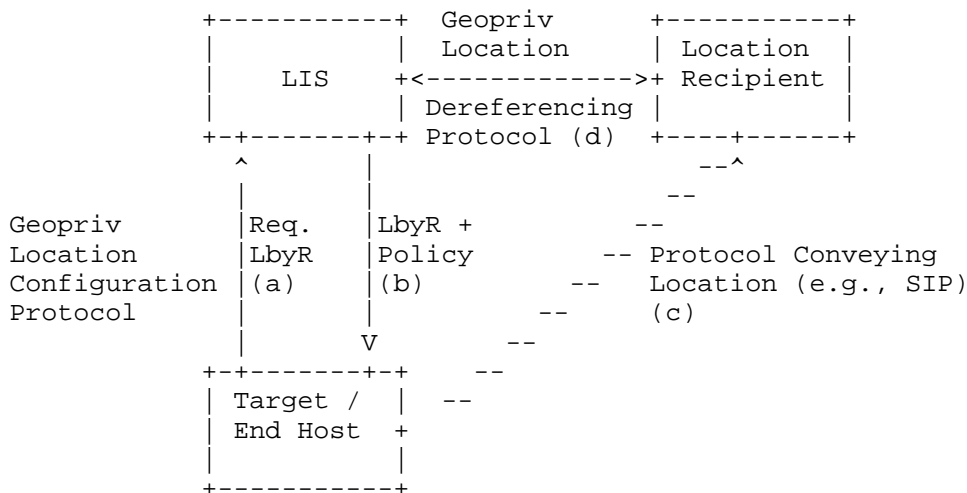


Figure 2: Location-by-Reference

Where location-by-reference is provided, the Recipient needs to dereference the LbyR in order to obtain location. The details for the dereferencing operations vary with the type of reference, such as an HTTP, HTTPS, SIP, secure SIP (SIPS), or SIP Presence URI.

For location-by-reference, the location server needs to maintain one or several URIs for each Target, timing out these URIs after a certain amount of time. References need to expire to prevent the Recipient of such a Uniform Resource Locator (URL) from being able to permanently track a host and to offer garbage collection functionality for the location server.

Off-path adversaries must be prevented from obtaining the Target's location. The reference contains a randomized component that prevents third parties from guessing it. When the Location Recipient fetches up-to-date location information from the location server, it can also be assured that the location information is fresh and not replayed. However, this does not address location theft.

With respect to the security of the dereference operation, [RFC6753], Section 6 states:

TLS MUST be used for dereferencing location URIs unless confidentiality and integrity are provided by some other mechanism, as discussed in Section 3. Location Recipients MUST authenticate the host identity using the domain name included in the location URI, using the procedure described in Section 3.1 of [RFC2818]. Local policy determines what a Location Recipient does if authentication fails or cannot be attempted.

The authorization by possession model (Section 4.1) further relies on TLS when transmitting the location URI to protect the secrecy of the URI. Possession of such a URI implies the same privacy considerations as possession of the PIDF-LO document that the URI references.

Location URIs MUST only be disclosed to authorized Location Recipients. The GEOPRIV architecture [RFC6280] designates the Rule Maker to authorize disclosure of the URI.

Protection of the location URI is necessary, since the policy attached to such a location URI permits anyone who has the URI to view the associated location information. This aspect of security is covered in more detail in the specification of location conveyance protocols, such as [RFC6442].

For authorizing access to location-by-reference, two authorization models were developed: "Authorization by Possession" and "Authorization via Access Control Lists". With respect to "Authorization by Possession", [RFC6753], Section 4.1 notes:

In this model, possession -- or knowledge -- of the location URI is used to control access to location information. A location URI might be constructed such that it is hard to guess (see C8 of [RFC5808]), and the set of entities that it is disclosed to can be limited. The only authentication this would require by the LS is evidence of possession of the URI. The LS could immediately authorize any request that indicates this URI.

Authorization by possession does not require direct interaction with a Rule Maker; it is assumed that the Rule Maker is able to exert control over the distribution of the location URI. Therefore, the LIS can operate with limited policy input from a Rule Maker.

Limited disclosure is an important aspect of this authorization model. The location URI is a secret; therefore, ensuring that adversaries are not able to acquire this information is paramount. Encryption, such as might be offered by TLS [RFC5246] or S/MIME [RFC5751], protects the information from eavesdroppers.

...

Using possession as a basis for authorization means that, once granted, authorization cannot be easily revoked. Cancellation of a location URI ensures that legitimate users are also affected; application of additional policy is theoretically possible but could be technically infeasible. Expiration of location URIs limits the usable time for a location URI, requiring that an attacker continue to learn new location URIs to retain access to current location information.

In situations where "Authorization by Possession" is not suitable (such as where location hiding [RFC6444] is required), the "Authorization via Access Control Lists" model may be preferred.

Without the introduction of a hierarchy, it would be necessary for the PSAP to obtain credentials, such as certificates or shared symmetric keys, for all the LISs in its coverage area, to enable it to successfully dereference LbyRs. In situations with more than a few LISs per PSAP, this would present operational challenges.

A certificate hierarchy providing PSAPs with client certificates chaining to the VESA could be used to enable the LIS to authenticate and authorize PSAPs for dereferencing. Note that unlike PIDF-LO signing (which mitigates modification of PIDF-LOs), this merely provides the PSAP with access to a (potentially unsigned) PIDF-LO, albeit over a protected TLS channel.

Another approach would be for the local LIS to upload location information to a location aggregation point who would in turn manage the relationships with the PSAP. This would shift the management burden from the PSAPs to the location aggregation points.

3.3. Proxy-Added Location

Instead of relying upon the end host to provide location, it is possible for a proxy that has the ability to determine the location of the end point (e.g., based on the end host IP or MAC address) to retrieve and add or override location information. This requires deployment of application-layer entities by ISPs, unlike the two other techniques. The proxies could be used for emergency or non-emergency communications, or both.

The use of proxy-added location is primarily applicable in scenarios where the end host does not provide location. As noted in [RFC6442], Section 4.1:

A SIP intermediary SHOULD NOT add location to a SIP request that already contains location. This will quite often lead to confusion within LRs. However, if a SIP intermediary adds location, even if location was not previously present in a SIP request, that SIP intermediary is fully responsible for addressing the concerns of any 424 (Bad Location Information) SIP response it receives about this location addition and MUST NOT pass on (upstream) the 424 response. A SIP intermediary that adds a locationValue MUST position the new locationValue as the last locationValue within the Geolocation header field of the SIP request.

...

A SIP intermediary MAY add a Geolocation header field if one is not present -- for example, when a user agent does not support the Geolocation mechanism but their outbound proxy does and knows the Target's location, or any of a number of other use cases (see Section 3).

As noted in [RFC6442], Section 3.3:

This document takes a "you break it, you bought it" approach to dealing with second locations placed into a SIP request by an intermediary entity. That entity becomes completely responsible for all location within that SIP request (more on this in Section 4).

While it is possible for the proxy to override location included by the end host, [RFC6442], Section 3.4 notes the operational limitations:

Overriding location information provided by the user requires a deployment where an intermediary necessarily knows better than an end user -- after all, it could be that Alice has an on-board GPS, and the SIP intermediary only knows her nearest cell tower. Which is more accurate location information? Currently, there is no way to tell which entity is more accurate or which is wrong, for that matter. This document will not specify how to indicate which location is more accurate than another.

The disadvantage of this approach is the need to deploy application-layer entities, such as SIP proxies, at IAPs or associated with IAPs. This requires that a standardized VoIP profile be deployed at every end Device and at every IAP. This might impose interoperability challenges.

Additionally, the IAP needs to take responsibility for emergency calls, even for customers with whom they have no direct or indirect relationship. To provide identity information about the emergency caller from the VSP, it would be necessary to let the IAP and the VSP interact for authentication (see, for example, "Diameter Session Initiation Protocol (SIP) Application" [RFC4740]). This interaction along the Authentication, Authorization, and Accounting infrastructure is often based on business relationships between the involved entities. An arbitrary IAP and VSP are unlikely to have a business relationship. If the interaction between the IAP and the VSP fails due to the lack of a business relationship, then typically a fall-back would be provided where no emergency caller identity information is made available to the PSAP and the emergency call still has to be completed.

4. Location Trust Assessment

The ability to assess the level of trustworthiness of conveyed location information is important, since this makes it possible to understand how much value should be placed on location information as part of the decision-making process. As an example, if automated location information is understood to be highly suspect or is absent, a call taker can put more effort into verifying the authenticity of the call and obtaining location information from the caller.

Location trust assessment has value, regardless of whether the location itself is authenticated (e.g., signed location) or is obtained directly from the location server (e.g., location-by-reference) over security transport, since these mechanisms do not provide assurance of the validity or provenance of location data.

To prevent location-theft attacks, the "entity" element of the PIDF-LO is of limited value if an unlinked pseudonym is provided in this field. However, if the LIS authenticates the Target, then the linkage between the pseudonym and the Target identity can be recovered in a post-incident investigation.

As noted in [Loc-Dependability], if the location object was signed, the Location Recipient has additional information on which to base their trust assessment, such as the validity of the signature, the identity of the Target, the identity of the LIS, whether the LIS authenticated the Target, and the identifier included in the "entity" field.

Caller accountability is also an important aspect of trust assessment. Can the individual purchasing the Device or activating service be identified, or did the call originate from a non-service initialized (NSI) Device whose owner cannot be determined? Prior to the call, was the caller authenticated at the network or application layer? In the event of a hoax call, can audit logs be made available to an investigator, or can information relating to the owner of an unlinked pseudonym be provided, enabling investigators to unravel the chain of events that led to the attack?

In practice, the source of the location data is important for location trust assessment. For example, location provided by a Location Information Server (LIS) whose administrator has an established history of meeting emergency location accuracy requirements (e.g., United States Phase II E-911 location accuracy) may be considered more reliable than location information provided by a third-party Location Service Provider (LSP) that disclaims use of location information for emergency purposes.

However, even where an LSP does not attempt to meet the accuracy requirements for emergency location, it still may be able to provide information useful in assessing how reliable location information is likely to be. For example, was location determined based on the nearest cell tower or 802.11 Access Point (AP), or was a triangulation method used? If based on cell tower or AP location data, was the information obtained from an authoritative source (e.g., the tower or AP owner), and when was the last time that the location of the tower or access point was verified?

For real-time validation, information in the signaling and media packets can be cross-checked against location information. For example, it may be possible to determine the city, state, country, or continent associated with the IP address included within SIP Via or Contact header fields, or the media source address, and compare this against the location information reported by the caller or conveyed in the PIDF-LO. However, in some situations, only entities close to the caller may be able to verify the correctness of location information.

Real-time validation of the timestamp contained within PIDF-LO objects (reflecting the time at which the location was determined) is also challenging. To address time-shifting attacks, the "timestamp" element of the PIDF-LO, defined in [RFC3863], can be examined and compared against timestamps included within the enclosing SIP message, to determine whether the location data is sufficiently fresh. However, the timestamp only represents an assertion by the LIS, which may or may not be trustworthy. For example, the Recipient of the signed PIDF-LO may not know whether the LIS supports time synchronization, or whether it is possible to reset the LIS clock manually without detection. Even if the timestamp was valid at the time location was determined, a time period may elapse between when the PIDF-LO was provided and when it is conveyed to the Recipient. Periodically refreshing location information to renew the timestamp even though the location information itself is unchanged puts additional load on LISs. As a result, Recipients need to validate the timestamp in order to determine whether it is credible.

While this document focuses on the discussion of real-time determination of suspicious emergency calls, the use of audit logs may help in enforcing accountability among emergency callers. For example, in the event of a hoax call, information relating to the owner of the unlinked pseudonym could be provided to investigators, enabling them to unravel the chain of events that led to the attack. However, while auditability is an important deterrent, it is likely to be of most benefit in situations where attacks on the emergency services system are likely to be relatively infrequent, since the resources required to pursue an investigation are likely to be considerable. However, although real-time validation based on PIDF-LO elements is challenging, where LIS audit logs are available (such as where a law enforcement agency can present a subpoena), linking of a pseudonym to the Device obtaining location can be accomplished during an investigation.

Where attacks are frequent and continuous, automated mechanisms are required. For example, it might be valuable to develop mechanisms to exchange audit trail information in a standardized format between ISPs and PSAPs / VSPs and PSAPs or heuristics to distinguish potentially fraudulent emergency calls from real emergencies. While a Completely Automated Public Turing test to tell Computers and Humans Apart (CAPTCHA) may be applied to suspicious calls to lower the risk from bot-nets, this is quite controversial for emergency services, due to the risk of delaying or rejecting valid calls.

5. Security Considerations

Although it is important to ensure that location information cannot be faked, the mitigation techniques presented in this document are not universally applicable. For example, there will be many GPS-enabled Devices that will find it difficult to utilize any of the solutions described in Section 3. It is also unlikely that users will be willing to upload their location information for "verification" to a nearby location server located in the access network.

This document focuses on threats that arise from conveyance of misleading location information, rather than caller identification or authentication and integrity protection of the messages in which location is conveyed. Nevertheless, these aspects are important. In some countries, regulators may not require the authenticated identity of the emergency caller (e.g., emergency calls placed from Public Switched Telephone Network (PSTN) pay phones or SIM-less cell phones). Furthermore, if identities can easily be crafted (as is the case with many VoIP offerings today), then the value of emergency caller authentication itself might be limited. As a result, attackers can forge emergency calls with a lower risk of being held accountable, which may encourage hoax calls.

In order to provide authentication and integrity protection for the Session Initiation Protocol (SIP) messages conveying location, several security approaches are available. It is possible to ensure that modification of the identity and location in transit can be detected by the Location Recipient (e.g., the PSAP), using cryptographic mechanisms, as described in "Enhancements for Authenticated Identity Management in the Session Initiation Protocol (SIP)" [RFC4474]. However, compatibility with Session Border Controllers (SBCs) that modify integrity-protected headers has proven to be an issue in practice, and as a result, a revision of [RFC4474] is in progress [SIP-Identity]. In the absence of an end-to-end solution, SIP over Transport Layer Security (TLS) can be used to provide message authentication and integrity protection hop by hop.

PSAPs remain vulnerable to distributed denial-of-service attacks, even where the mitigation techniques described in this document are utilized. Placing a large number of emergency calls that appear to come from different locations is an example of an attack that is difficult to carry out within the legacy system but is easier to imagine within IP-based emergency services. Also, in the current system, it would be very difficult for an attacker from one country to attack the emergency services infrastructure located in another country, but this attack is possible within IP-based emergency services.

While manually mounting the attacks described in Section 2 is non-trivial, the attacks described in this document can be automated. While manually carrying out a location theft would require that the attacker be in proximity to the location being spoofed, or to collude with another end host, an attacker able to run code on an end host can obtain its location and cause an emergency call to be made. While manually carrying out a time-shifting attack would require that the attacker visit the location and submit it before the location information is considered stale, while traveling rapidly away from that location to avoid apprehension, these limitations would not apply to an attacker able to run code on the end host. While obtaining a PIDF-LO from a spoofed IP address requires that the attacker be on the path between the HELD requester and the LIS, if the attacker is able to run code requesting the PIDF-LO, retrieve it from the LIS, and then make an emergency call using it, this attack becomes much easier. To mitigate the risk of automated attacks, service providers can limit the ability of untrusted code (such as WebRTC applications written in JavaScript) to make emergency calls.

Emergency services have three finite resources subject to denial-of-service attacks: the network and server infrastructure; call takers and dispatchers; and the first responders, such as firefighters and police officers. Protecting the network infrastructure is similar to protecting other high-value service providers, except that location information may be used to filter call setup requests, to weed out requests that are out of area. Even for large cities, PSAPs may only have a handful of call takers on duty. So, even if automated techniques are utilized to evaluate the trustworthiness of conveyed location and call takers can, by questioning the caller, eliminate many hoax calls, PSAPs can be overwhelmed even by a small-scale attack. Finally, first-responder resources are scarce, particularly during mass-casualty events.

6. Privacy Considerations

The emergency calling architecture described in [RFC6443] utilizes the PIDF-LO format defined in [RFC4119]. As described in the location privacy architecture [RFC6280], privacy rules that may include policy instructions are conveyed along with the location object.

The intent of the location privacy architecture was to provide strong privacy protections, as noted in [RFC6280], Section 1.1:

A central feature of the Geopriv architecture is that location information is always bound to privacy rules to ensure that entities that receive location information are informed of how they may use it. These rules can convey simple directives ("do not share my location with others"), or more robust preferences ("allow my spouse to know my exact location all of the time, but only allow my boss to know it during work hours")... The binding of privacy rules to location information can convey users' desire for and expectations of privacy, which in turn helps to bolster social and legal systems' protection of those expectations.

However, in practice this architecture has limitations that apply within emergency and non-emergency situations. As noted in Section 1.2.2, concerns about hoax calls have led to restrictions on anonymous emergency calls. Caller identification (potentially asserted in SIP via P-Asserted-Identity and SIP Identity) may be used during emergency calls. As a result, in many cases location information transmitted within SIP messages can be linked to caller identity. For example, in the case of a signed LbyV, there are privacy concerns arising from linking the location object to identifiers to prevent replay attacks, as described in Section 3.1.

The ability to observe location information during emergency calls may also represent a privacy risk. As a result, [RFC6443] requires transmission-layer security for SIP messages, as well as interactions with the location server. However, even where transmission-layer security is used, privacy rules associated with location information may not apply.

In many jurisdictions, an individual requesting emergency assistance is assumed to be granting permission to the PSAP, call taker, and first responders to obtain their location in order to accelerate dispatch. As a result, privacy policies associated with location are implicitly waived when an emergency call is initiated. In addition, when location information is included within SIP messages in either emergency or non-emergency uses, SIP entities receiving the SIP message are implicitly assumed to be authorized Location Recipients, as noted in [RFC5606], Section 3.2:

Consensus has emerged that any SIP entity that receives a SIP message containing LI through the operation of SIP's normal routing procedures or as a result of location-based routing should be considered an authorized recipient of that LI. Because of this presumption, one SIP element may pass the LI to another even if the LO it contains has <retransmission-allowed> set to "no"; this

sees the passing of the SIP message as part of the delivery to authorized recipients, rather than as retransmission. SIP entities are still enjoined from passing these messages outside the normal routing to external entities if <retransmission-allowed> is set to "no", as it is the passing to third parties that <retransmission-allowed> is meant to control.

Where LbyR is utilized rather than LbyV, it is possible to apply more restrictive authorization policies, limiting access to intermediaries and snoopers. However, this is not possible if the "authorization by possession" model is used.

7. Informative References

- [EENA] EENA, "False Emergency Calls", EENA Operations Document, Version 1.1, May 2011, <http://www.eena.org/ressource/static/files/2012_05_04-3.1.2.fc_v1.1.pdf>.
- [GPSCounter] Warner, J. and R. Johnston, "GPS Spoofing Countermeasures", Los Alamos research paper LAUR-03-6163, December 2003.
- [Loc-Dependability] Thomson, M. and J. Winterbottom, "Digital Signature Methods for Location Dependability", Work in Progress, draft-thomson-geopriv-location-dependability-07, March 2011.
- [NENA-i2] NENA 08-001, "NENA Interim VoIP Architecture for Enhanced 9-1-1 Services (i2)", Version 2, August 2010.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC2818] Rescorla, E., "HTTP Over TLS", RFC 2818, May 2000, <<http://www.rfc-editor.org/info/rfc2818>>.
- [RFC3261] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., and E. Schooler, "SIP: Session Initiation Protocol", RFC 3261, June 2002, <<http://www.rfc-editor.org/info/rfc3261>>.
- [RFC3693] Cuellar, J., Morris, J., Mulligan, D., Peterson, J., and J. Polk, "Geopriv Requirements", RFC 3693, February 2004, <<http://www.rfc-editor.org/info/rfc3693>>.

- [RFC3694] Danley, M., Mulligan, D., Morris, J., and J. Peterson, "Threat Analysis of the Geopriv Protocol", RFC 3694, February 2004, <<http://www.rfc-editor.org/info/rfc3694>>.
- [RFC3748] Aboba, B., Blunk, L., Vollbrecht, J., Carlson, J., and H. Levkowetz, Ed., "Extensible Authentication Protocol (EAP)", RFC 3748, June 2004, <<http://www.rfc-editor.org/info/rfc3748>>.
- [RFC3863] Sugano, H., Fujimoto, S., Klyne, G., Bateman, A., Carr, W., and J. Peterson, "Presence Information Data Format (PIDF)", RFC 3863, August 2004, <<http://www.rfc-editor.org/info/rfc3863>>.
- [RFC4119] Peterson, J., "A Presence-based GEOPRIV Location Object Format", RFC 4119, December 2005, <<http://www.rfc-editor.org/info/rfc4119>>.
- [RFC4474] Peterson, J. and C. Jennings, "Enhancements for Authenticated Identity Management in the Session Initiation Protocol (SIP)", RFC 4474, August 2006, <<http://www.rfc-editor.org/info/rfc4474>>.
- [RFC4479] Rosenberg, J., "A Data Model for Presence", RFC 4479, July 2006, <<http://www.rfc-editor.org/info/rfc4479>>.
- [RFC4740] Garcia-Martin, M., Ed., Belinchon, M., Pallares-Lopez, M., Canales-Valenzuela, C., and K. Tammi, "Diameter Session Initiation Protocol (SIP) Application", RFC 4740, November 2006, <<http://www.rfc-editor.org/info/rfc4740>>.
- [RFC5012] Schulzrinne, H. and R. Marshall, Ed., "Requirements for Emergency Context Resolution with Internet Technologies", RFC 5012, January 2008, <<http://www.rfc-editor.org/info/rfc5012>>.
- [RFC5069] Taylor, T., Ed., Tschofenig, H., Schulzrinne, H., and M. Shanmugam, "Security Threats and Requirements for Emergency Call Marking and Mapping", RFC 5069, January 2008, <<http://www.rfc-editor.org/info/rfc5069>>.
- [RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, August 2008, <<http://www.rfc-editor.org/info/rfc5246>>.

- [RFC5491] Winterbottom, J., Thomson, M., and H. Tschofenig, "GEOPRIV Presence Information Data Format Location Object (PIDF-LO) Usage Clarification, Considerations, and Recommendations", RFC 5491, March 2009, <<http://www.rfc-editor.org/info/rfc5491>>.
- [RFC5606] Peterson, J., Hardie, T., and J. Morris, "Implications of 'retransmission-allowed' for SIP Location Conveyance", RFC 5606, August 2009, <<http://www.rfc-editor.org/info/rfc5606>>.
- [RFC5751] Ramsdell, B. and S. Turner, "Secure/Multipurpose Internet Mail Extensions (S/MIME) Version 3.2 Message Specification", RFC 5751, January 2010, <<http://www.rfc-editor.org/info/rfc5751>>.
- [RFC5808] Marshall, R., Ed., "Requirements for a Location-by-Reference Mechanism", RFC 5808, May 2010, <<http://www.rfc-editor.org/info/rfc5808>>.
- [RFC5985] Barnes, M., Ed., "HTTP-Enabled Location Delivery (HELD)", RFC 5985, September 2010, <<http://www.rfc-editor.org/info/rfc5985>>.
- [RFC6280] Barnes, R., Lepinski, M., Cooper, A., Morris, J., Tschofenig, H., and H. Schulzrinne, "An Architecture for Location and Location Privacy in Internet Applications", BCP 160, RFC 6280, July 2011, <<http://www.rfc-editor.org/info/rfc6280>>.
- [RFC6442] Polk, J., Rosen, B., and J. Peterson, "Location Conveyance for the Session Initiation Protocol", RFC 6442, December 2011, <<http://www.rfc-editor.org/info/rfc6442>>.
- [RFC6443] Rosen, B., Schulzrinne, H., Polk, J., and A. Newton, "Framework for Emergency Calling Using Internet Multimedia", RFC 6443, December 2011, <<http://www.rfc-editor.org/info/rfc6443>>.
- [RFC6444] Schulzrinne, H., Liess, L., Tschofenig, H., Stark, B., and A. Kuett, "Location Hiding: Problem Statement and Requirements", RFC 6444, January 2012, <<http://www.rfc-editor.org/info/rfc6444>>.
- [RFC6753] Winterbottom, J., Tschofenig, H., Schulzrinne, H., and M. Thomson, "A Location Dereference Protocol Using HTTP-Enabled Location Delivery (HELD)", RFC 6753, October 2012, <<http://www.rfc-editor.org/info/rfc6753>>.

- [RFC6881] Rosen, B. and J. Polk, "Best Current Practice for Communications Services in Support of Emergency Calling", BCP 181, RFC 6881, March 2013, <<http://www.rfc-editor.org/info/rfc6881>>.
- [RFC7090] Schulzrinne, H., Tschofenig, H., Holmberg, C., and M. Patel, "Public Safety Answering Point (PSAP) Callback", RFC 7090, April 2014, <<http://www.rfc-editor.org/info/rfc7090>>.
- [RFC7199] Barnes, R., Thomson, M., Winterbottom, J., and H. Tschofenig, "Location Configuration Extensions for Policy Management", RFC 7199, April 2014, <<http://www.rfc-editor.org/info/rfc7199>>.
- [RFC7340] Peterson, J., Schulzrinne, H., and H. Tschofenig, "Secure Telephone Identity Problem Statement and Requirements", RFC 7340, September 2014, <<http://www.rfc-editor.org/info/rfc7340>>.
- [RFC7375] Peterson, J., "Secure Telephone Identity Threat Model", RFC 7375, October 2014, <<http://www.rfc-editor.org/info/rfc7375>>.
- [SA] "Saudi Arabia - Illegal sale of SIMs blamed for surge in hoax calls", Arab News, April 5, 2010, <<http://www.arabnews.com/node/341463>>.
- [SIP-Identity] Peterson, J., Jennings, C. and E. Rescorla, "Authenticated Identity Management in the Session Initiation Protocol (SIP)", Work in Progress, draft-ietf-stir-rfc4474bis-02, October 2014.
- [STIR] IETF, "Secure Telephone Identity Revisited (stir) Working Group", October 2013, <<http://datatracker.ietf.org/wg/stir/charter/>>.
- [SWATing] "SWATing 911 Calls", Dispatch Magazine On-Line, April 6, 2013, <<http://www.911dispatch.com/swating-911-calls/>>.
- [Swatting] "Don't Make the Call: The New Phenomenon of 'Swatting'", Federal Bureau of Investigation, February 4, 2008, <<http://www.fbi.gov/news/stories/2008/february/swatting020408>>.

- [TASMANIA] "Emergency services seek SIM-less calls block", ABC News Online, August 18, 2006, <<http://www.abc.net.au/elections/tas/2006/news/stories/1717956.htm?elections/tas/2006/>>.
- [UK] "Rapper makes thousands of prank 999 emergency calls to UK police", Digital Journal, June 24, 2010, <<http://www.digitaljournal.com/article/293796?tp=1>>.

Acknowledgments

We would like to thank the members of the IETF ECRIT working group, including Marc Linsner and Brian Rosen, for their input at IETF 85 that helped get this document pointed in the right direction. We would also like to thank members of the IETF GEOPRIV working group, including Richard Barnes, Matt Lepinski, Andrew Newton, Murugaraj Shanmugam, and Martin Thomson for their feedback on previous versions of this document. Alissa Cooper, Adrian Farrel, Pete Resnick, Meral Shirazipour, and Bert Wijnen provided helpful review comments during the IETF last call.

Authors' Addresses

Hannes Tschofenig
Austria

E-Mail: Hannes.tschofenig@gmx.net
URI: <http://www.tschofenig.priv.at>

Henning Schulzrinne
Columbia University
Department of Computer Science
450 Computer Science Building
New York, NY 10027
United States

Phone: +1 212 939 7004
E-Mail: hgs@cs.columbia.edu
URI: <http://www.cs.columbia.edu>

Bernard Aboba (editor)
Microsoft Corporation
One Microsoft Way
Redmond, WA 98052
United States

E-Mail: bernard_aboba@hotmail.com