ECRIT Working Group
INTERNET-DRAFT
Category: Informational
Expires: September 12, 2013

H. Tschofenig
Nokia Siemens Networks
H. Schulzrinne
Columbia University
B. Aboba (ed.)
Microsoft Corporation
13 March 2013

Trustworthy Location
draft-ietf-ecrit-trustworthy-location-05.txt

Abstract

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

This document describes how to convey location in a manner that is inherently secure and reliable. It also provides guidelines for assessing the trustworthiness of location information.

Copyright Notice

Table of Contents

1. Introduction

Several public and commercial services depend upon location information in their operations. This includes emergency services (such as fire, ambulance and police) as well as commercial services such as food delivery and roadside assistance.

Services that depend on location commonly experience security issues today. While prank calls have been a problem for emergency services dating back to the time of street corner call boxes, with the move to IP-based emergency services, the ability to launch automated attacks has increased. 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, which in certain European countries can constitute as much as 70% of all emergency calls. Reducing the number of prank 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.

Since prank emergency calls can endanger bystanders or emergency services personnel, or divert resources away from legitimate emergencies, they can be life threatening. A particularly dangerous form of prank call is "swatting" - an prank emergency call that draws a response from law enforcement (e.g. a fake hostage situation that results in dispatching of a "Special Weapons And Tactics" (SWAT) team). In 2008 the FBI issued a warning [Swatting] about an increase in the frequency and sophistication of these attacks.

Many documented cases of "swatting" involve not only the faking of an emergency, but also the absence of accurate caller identification and the delivery of misleading location data. Today these attacks are often carried out by providing false caller identification, since for circuit-switched calls from landlines, location provided to the PSAP is determined from a lookup using the calling telephone number. With IP-based emergency services, in addition to the potential for false caller identification, it is also possible to attach misleading location information to the emergency call.

Ideally, a call taker at a PSAP should be put in the position 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 about the caller. Where real-time assessment is not possible, it is important to be able to determine the source of the call in a post-mortem, so as to be able to enforce accountability.

This document defines terminology (including the meaning of "trustworthy location") in Section 1.1, investigates security threats in Section 2, outlines potential solutions 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].

The definition for "Target" is taken from "Geopriv Requirements" [RFC3693].

The term "location determination method" refers to the mechanism used to determine the location of a Target. This may be something employed by a location information server (LIS), or by the Target itself. It specifically does not refer to the location configuration protocol (LCP) used to deliver location information either to the Target or the Recipient. This term is re-used from "GEOPRIV 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 URI are re-used 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.

2. Threats

While previous IETF documents have analyzed aspects of the security of emergency services or threats to geographic location privacy, those documents do not cover the threats arising from unreliable location information.

A threat analysis of the emergency services system is provided in "Security Threats and Requirements for Emergency Call Marking and Mapping" [RFC5069]. 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. [RFC5069] also describes attacks against individuals, including attempts to prevent an individual from receiving aid, or to gain information about an emergency. "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.

This document focuses on threats from attackers providing false location information within emergency calls. Since we do not focus on attackers gaining control of infrastructure elements (e.g., location servers, call route servers) or the emergency services IP network, the threats are derived from the introduction of untrustworthy location information by end hosts. In addition to threats arising from the intentional forging of location information, end hosts may be induced to provide untrustworthy location information. For example, end hosts may obtain location from civilian GPS, which is vulnerable to spoofing [GPSCounter] or from third party Location Service Providers (LSPs) which may be vulnerable to attack or may not warrant the use of their services for emergency purposes.

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 LIS, the LoST infrastructure, used for mapping locations to PSAP address, 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 whether it is acting as a bot.

In this document, we focus only on the malicious end host adversary model.

2.1. Location Spoofing

An adversary can provide false location information in an emergency call in order to misdirect emergency resources. For calls originating within the PSTN, this attack can be carried out via caller-id spoofing. Where location is attached to the emergency call by an end host, several avenues are available to provide false location information:

1. The end host could fabricate a PIDF-LO and convey it within an emergency call;

2. The VSP (and indirectly a LIS) could be fooled into using the wrong identity (such as an IP address) for location lookup, thereby providing the end host with misleading location information;

3. Inaccurate or out-of-date information (such spoofed GPS signals, a stale wiremap or an inaccurate access point location database) could be utilized by the LIS or the end host in its location determination, thereby leading to an inaccurate determination of location.

By analysis of the SIP headers, it may be possible to flag situations where the conveyed location is suspect (e.g. potentially wrong city, state, country or continent). However, in other situations only entities close to the caller may be able to verify the correctness of location information.

The following list presents threats specific to location information handling:

Place shifting: Trudy, the adversary, pretends to be at an arbitrary location. In some cases, place shifting can be limited in range, e.g., to the coverage area of a particular cell tower.

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

Location theft: Trudy observes Alice's location and replays it as her own.

Location swapping: Trudy and Malory, located in different locations, can collude and swap location information and pretend to be in each other's location.

2.2. Identity Spoofing

With calls originating on an IP network, at least two forms of identity are relevant, with the distinction created by the split between the AIP 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, this also may not have been authenticated, such as in the case where an open IEEE 802.11 Access Point is used to initiate a nuisance emergency call. Although endpoint information such as the IP or MAC address may have been logged, tying this back to the device owner may be challenging.

Unlike the existing telephone system, VoIP emergency calls could require strong identity, which need not necessarily be coupled to a business relationship with the AIP, ISP or VSP. However, due to the time-critical nature of emergency calls, multi-layer authentication is undesirable, so that in most cases, only the device placing the call will be able to be identified, making the system vulnerable to bot-net attacks. 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. Solutions

This section presents three mechanisms which can be used to convey location: signed location by value (Section 3.1), location by reference (Section 3.2) and proxy added location (Section 3.3).

In order for to provide authentication and integrity protection for the SIP messages conveying location, several security approaches are available. While it is possible for proxies to use security mechanisms such as SIP Identity [RFC4474] to ensure that modifications to the location in transit can be detected by the location recipient (e.g., the PSAP), compatibility with Session

Border Controllers (SBCs) that modify integrity-protected headers has proven to be an issue in practice. As a result, the use of SIP over TLS is at present a more likely mechanism to provide per-message authentication and integrity protection.

3.1. Signed Location by Value

With location signing, a location server signs the location information before it is sent to the end host, (the entity subject to the location determination process).

The signed location information is then verified by the location recipient and not by the target. 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 in step (c) who verifies it. For SIP, the procedures described in "Location Conveyance for the Session Initiation Protocol" [RFC6442] are applicable for location conveyance.

```
              +-----------+              +-----------+
              |           |              | Location  |
              |    LIS    |              | Recipient |
              |           |              |           |
              +-+-------+-+              +----+------+
                ^       |                   --^
                |       |                 --
 Geopriv        |Req.   |               --
 Location       |Signed |Signed       -- Geopriv
 Configuration  |Loc.   |Loc.       --   Using Protocol
 Protocol       |(a)    |(b)      --     (e.g., SIP)
                |       v       --       (c)
              +-+-------+-+   --
              | Target /  | --
              | End Host  +
              |           |
              +-----------+
```

Figure 1: Location Signing

In order to limit replay attacks, additional information, such as timestamps or expiration times, has to be included together with the signed location. If the location is retrieved from a location server, even a stationary end host has to periodically obtain a fresh signed location, or incur the additional delay of querying during the emergency call.

While bot-nets are unlikely to be deterred by location signing,

accurate location information would limit the subset of the bot-net that could be used for an attack, as only hosts within the PSAP serving area would be useful in placing emergency calls.

To prevent location-swapping attacks it is necessary to include some some target-specific identity information. The required information depends on whether the goal is real-time verification by the location recipient or post-mortem analysis (where the goal is determination of the legal entity responsible for the attack). As argued in Section 4, real-time verification is not always possible.

Location signing is unlikely to deter attacks launched by bot-nets, since the work required to verify the location signature is considerable. 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 a GPS device, 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.

A straw-man proposal for location signing is provided in [I-D.thomson-geopriv-location-dependability], and [NENA-i2] Section 3.7 includes operational recommendations relating to location signing:

> Location determination is out of scope for NENA, but we can offer guidance 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 certs 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. Anti-spoofing mechanisms should be applied to the Location Reporting method.
>
> [Note: The term Valid Emergency Services Authority (VESA) refers to the root certificate authority.]

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 VESA can be accomplished either by having the VESA directly sign the LIS certificates, or by the creation of intermediate 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 whether they are significantly different from say, requirements for issuance of an 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. As noted in "A Location Dereference Protocol Using HTTP-Enabled Location Delivery (HELD)" [RFC6753], a location reference can be obtained via HTTP-Enabled Location Delivery (HELD) [RFC5985] or the Dynamic Host Configuration Protocol (DHCP) location URI option [DHCP-URI-OPT].

Figure 2 shows the communication model with the target requesting a location reference in step (a), the location server returns the reference 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.

```
                +-----------+  Geopriv      +-----------+
                |           |  Location     | Location  |
                |    LIS    +<------------->+ Recipient |
                |           | Dereferencing |           |
                +-+-------+-+ Protocol (d)  +----+------+
                  ^       |                    --^
                  |       |                  --
    Geopriv       |Req.   |                --
    Location      |LbyR   |LbyR          -- Geopriv
    Configuration |(a)    |(b)         --   Using Protocol
    Protocol      |       |          --     (e.g., SIP)
                  |       V        --       (c)
                +-+-------+-+    --
                | Target /  |  --
                | End Host  +
                |           |
                +-----------+
```

Figure 2: Location by Reference

Where location by reference is provided, the recipient needs to deference the LbyR in order to obtain location. The details for the dereferencing operations vary with the type of reference, such as a HTTP, HTTPS, SIP, SIPS URI or a 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 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 swapping.

With respect to the security of the de-reference 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 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 o 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 hierarchy, it would be necessary for the PSAP to obtain client certificates or Digest credentials for all the LISes in its coverage area, to enable it to successfully dereference LbyRs. In situations with more than a few LISes 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 against 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 Adding Location

Instead of relying upon the end host to provide location, 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.

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 AIPs or associated with AIPs. This requires a standardized VoIP profile to be deployed at every end device and at every AIP. This might impose interoperability challenges.

Additionally, the AIP needs to take responsibility for emergency calls, even for customers they have no direct or indirect relationship with. To provide identity information about the emergency caller from the VSP it would be necessary to let the AIP and the VSP to interact for authentication (see, for example, [RFC4740]). This interaction along the Authentication, Authorization and Accounting infrastructure is often based on business relationships between the involved entities. The AIP and the VSP are very likely to have no such business relationship, particularly when talking about an arbitrary VSP somewhere on the Internet. In case that the interaction between the AIP 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, a call taker can put more effort into obtaining location information from the caller.

Caller accountability is another 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 prank 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 lead to the attack? In practice, the ability to identify a caller may decrease the likelihood of caller misbehavior. For example, where emergency calls have been allowed from handsets lacking a SIM card, or where ownership of the SIM card cannot be determined, the frequency of nuisance calls has often been unacceptably high [TASMANIA][UK][SA].

Note that location trust assessment has value regardless of whether the location has been conveyed securely (via signed location, location-by-reference or proxy-added location) or not (via location-by-value without location signing), since secure conveyance does not provide assurance relating to the validity or provenance of location data.

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. Phase II) 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 about 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 region associated with the IP address included within SIP Via: or Contact: headers, or the media source address, and compare this against the location information reported by the caller or conveyed in the PIDF-LO. While a CAPTCHA-style test 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.

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.

Due to these design choices, 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. While PIDF-LO signing would prevent modification of a PIDF-LO or invention of one out of whole cloth, it would not prevent this cut and paste attack. Neither would implementation of "Enhancements for Authenticated Identity Management in the Session Initiation Protocol (SIP)" [RFC4474], allowing the recipient to verify the identity assertion in the From: header. However, while it might not be possible to detect the cut and paste

in real-time, examination of the audit logs might provide enough information to enable events to be reconstructed.

Real-time validation of the timestamp contained within PIDF-LO objects (reflecting the time at which the location was determined) is also challenging. Even if the PIDF-LO is signed 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 LISes. 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 prank call, information relating to the owner of the unlinked pseudonym could be provided to investigators, enabling them to unravel the chain of events that lead 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 in a post-mortem.

Where attacks are frequent and continuous, automated mechanisms are required. For example, it might be valuable to develop mechanisms to exchange audit trails information in a standardized format between ISPs and PSAPs / VSPs and PSAPs or heuristics to distinguish potentially fraudulent emergency calls from real emergencies.

5. Security Considerations

IP-based emergency services face a number of security threats that do not exist within the legacy system. In order to limit prank calls, legacy emergency services rely on the ability to identify callers, as well as on the difficulty of location spoofing for normal users. The ability to ascertain identity is important, since the threat of punishment reduces prank calls; as an example, calls from pay phones

are subject to greater scrutiny by the call taker.

Mechanically placing a large number of emergency calls that appear to come from different locations is difficult in a legacy environment. Also, in the current system, it would be very difficult for an attacker from country 'Foo' to attack the emergency services infrastructure located in country 'Bar'.

However, within an IP-based emergency services a number of these attacks become much easier to mount. 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 fire fighters 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. PSAPs even for large cities may only have a handful of PSAP call takers on duty, so even if they can, by questioning the caller, eliminate a lot of prank calls, they are quickly overwhelmed by even a small-scale attack. Finally, first responder resources are scarce, particularly during mass-casualty events.

Attackers may want to modify, prevent or delay emergency calls. In some cases, this will lead the PSAP to dispatch emergency personnel to an emergency that does not exist and, hence, the personnel might not be available to other callers. It might also be possible for an attacker to impede the users from reaching an appropriate PSAP by modifying the location of an end host or the information returned from the mapping protocol. In some countries, regulators may not require the authenticated identity of the emergency caller, as is true for PSTN-based emergency calls placed from pay phones or SIM-less cell phones today. Furthermore, if identities can easily be crafted (as it is the case with many VoIP offerings today), then the value of emergency caller authentication itself might be limited. As a consequence, an attacker can forge emergency call information without the chance of being held accountable for its own actions.

The above-mentioned attacks are mostly targeting individual emergency callers or a very small fraction of them. If attacks are, however, launched against the mapping architecture (see [RFC5582] or against the emergency services IP network (including PSAPs), a larger region and a large number of potential emergency callers are affected. The call takers themselves are a particularly scarce resource and if human interaction by these call takers is required then this can very quickly have severe consequences.

Although it is important to ensure that location information cannot be faked 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.

6. IANA Considerations

This document does not require actions by IANA.

7. References

7.1. Informative References

[DHCP-URI-OPT]
Polk, J., "Dynamic Host Configuration Protocol (DHCP) IPv4 and IPv6 Option for a Location Uniform Resource Identifier (URI)", Internet draft (work in progress), draft-ietf-geopriv-dhcp-lbyr-uri-option-19, February 2013.

[EENA] EENA, "False Emergency Calls", EENA Operations Document, Version 1.0, March 2011, http://www.eena.org/ressource/static/files/2011_03_15_3.1.2.fc_v1.0.pdf

[GPSCounter]
Warner, J. S. and R. G. Johnston, "GPS Spoofing Countermeasures", Los Alamos research paper LAUR-03-6163, December 2003.

[NENA-i2] "08-001 NENA Interim VoIP Architecture for Enhanced 9-1-1 Services (i2)", December 2005.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2818] Rescorla, E., "HTTP over TLS", RFC 2818, May 2000.

[RFC3693] Cuellar, J., Morris, J., Mulligan, D., Peterson, J., and J. Polk, "Geopriv Requirements", RFC 3693, February 2004.

[RFC3694] Danley, M., Mulligan, D., Morris, J. and J. Peterson, "Threat Analysis of the Geopriv Protocol", RFC 3694, February 2004.

[RFC3748] Aboba, B., Blunk, L., Vollbrecht, J., Carlson, J., and H. Levkowetz, "Extensible Authentication Protocol (EAP)", RFC 3748, June 2004.

[RFC4474] Peterson, J. and C. Jennings, "Enhancements for Authenticated Identity Management in the Session Initiation Protocol (SIP)", RFC 4474, August 2006.

[RFC4479] Rosenberg, J., "A Data Model for Presence", RFC 4479, July 2006.

[RFC4740] Garcia-Martin, M., Belinchon, M., Pallares-Lopez, M., Canales-Valenzuela, C., and K. Tammi, "Diameter Session Initiation Protocol (SIP) Application", RFC 4740, November 2006.

[RFC5069] Taylor, T., Tschofenig, H., Schulzrinne, H. and M. Shanmugam, "Security Threats and Requirements for Emergency Call Marking and Mapping", RFC 5069, January 2008.

[RFC5246] Dierks, T. and E. Rescorla, "The Transport Level Security (TLS) Protocol Version 1.2", RFC 5246, August 2008.

[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.

[RFC5582] Schulzrinne, H., "Location-to-URL Mapping Architecture and Framework", RFC 5582, September 2009.

[RFC5751] Ramsdell, B. and S. Turner, "Secure/Multipurpose Internet Mail Extensions (S/MIME) Version 3.2 Message Specification", RFC 5751, January 2010.

[RFC5808] Marshall, R., "Requirements for a Location-by-Reference Mechanism", RFC 5808, May 2010.

[RFC5985] Barnes, M., "HTTP Enabled Location Delivery (HELD)", RFC 5985, September 2010.

[RFC6280] Barnes, R., et. al, "An Architecture for Location and Location Privacy in Internet Applications", RFC 6280, July 2011.

[RFC6442] Polk, J., Rosen, B. and J. Peterson, "Location Conveyance for the Session Initiation Protocol", RFC 6442, December 2011.

[RFC6444] Schulzrinne, H., Liess, L., Tschofenig, H., Stark, B., and A. Kuett, "Location Hiding: Problem Statement and Requirements", RFC 6444, January 2012.

[RFC6753] Winterbottom, J., Tschofenig. H., Schulzrinne, H. and M. Thomson, "A Location Dereference Protocol Using HTTP-Enabled

Location Delivery (HELD)", RFC 6753, October 2012.

[SA] "Saudi Arabia - Illegal sale of SIMs blamed for surge in prank calls", Arab News, May 4, 2010, http://www.menafn.com/qn_news_story_s.asp?StoryId=1093319384

[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/news/newsitems/200608/s1717956.htm

[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, Henning Schulzrinne and Brian Rosen, for their input at IETF 85 that helped get this documented pointed in the right direction. We would also like to thank members of the IETF GEOPRIV WG, including Andrew Newton, Murugaraj Shanmugam, Martin Thomson, Richard Barnes and Matt Lepinski for their feedback to previous versions of this document.

Authors' Addresses

Hannes Tschofenig
Nokia Siemens Networks
Linnoitustie 6
Espoo 02600
Finland

Phone: +358 (50) 4871445
Email: 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
US

Phone: +1 212 939 7004
Email: hgs@cs.columbia.edu
URI: http://www.cs.columbia.edu

Bernard Aboba
Microsoft Corporation
One Microsoft Way
Redmond, WA 98052
US

Email: bernard_aboba@hotmail.com