TOC

ECRIT	H. Tschofenig
Internet-Draft	Nokia Siemens Networks
Intended status: Informational	H. Schulzrinne
Expires: September 7, 2010	Columbia University
	B. Aboba
	Microsoft Corporation
	March 06, 2010

Trustworthy Location Information
draft-tschofenig-ecrit-trustworthy-location-03.txt

Abstract

For some location-based applications, such as emergency calling or roadside assistance, it appears that the identity of the requestor is less important than accurate and trustworthy location information. To ensure adequate help location has to be left untouched by the end point or by entities in transit.

This document lists different threats, an adversary model, outlines three frequentlly discussed solutions and discusses operational considerations. Finally, the document concludes with a suggestion on how to move forward.

Status of this Memo

Copyright Notice

Table of Contents

1. Introduction

TOC

Much of the focus in trustable networks has been on ensuring the reliability of personal identity information or verifying privileges. However, in some cases, access to trustworthy location information is more important than identity since some services are meant to be widely available, regardless of the identity of the requestor. Emergency services, such as fire department, ambulance and police, but also commercial services such as food delivery and roadside assistance are among those. Customers, competitors or emergency callers lie about their location to harm the service provider or to deny services to others, by tying up the service capacity. In addition, if third parties can modify the information, they can deny services to the requestor.

Physical security is often based on location. As a trivial example, light switches in buildings are not typically protected by keycards or passwords, but are only accessible to those within the perimeter of the building. Merchants processing credit card payments already use location information to estimate the risk that a transaction is fraudulent, based on the HTTP client's IP address (that is then translated to location). In all these cases, trustworthy location information can be used to augment identity information or, in some cases, avoid the need for role-based authorization.

A number of standardization organizations have developed mechanisms to make civic and geodetic location available to the end host. Examples for these protocols are LLDP-MED [LLDP-MED] (, “Telecommunications: IP Telephony Infrastructure: Link Layer Discovery Protocol for Media Endpoint Devices, ANSI/TIA-1057-2006,” April 2006.), DHCP extensions (see [RFC4776] (Schulzrinne, H., “Dynamic Host Configuration Protocol (DHCPv4 and DHCPv6) Option for Civic Addresses Configuration Information,” November 2006.), [RFC3825] (Polk, J., Schnizlein, J., and M. Linsner, “Dynamic Host Configuration Protocol Option for Coordinate-based Location Configuration Information,” July 2004.)), HELD [I-D.ietf-geopriv-http-location-delivery] (Barnes, M., Winterbottom, J., Thomson, M., and B. Stark, “HTTP Enabled Location Delivery (HELD),” August 2009.), or the protocols developed within the IEEE as part of their link-layer specifications. The server offering this information is usually called a Location Information Server (LIS). More common with high-quality cellular devices is the ability for the end host itself to determine its own location using GPS. The location information is then provided, by reference or value, to the service-providing entities, i.e. location recipients, via application protocols, such as HTTP, SIP or XMPP.

This document investigates the security threats in Section 4 (Threats), and outlines three solutions that are frequently mentioned in Section 5 (Solution Proposals). We use emergency services an example to illustrate the security problems, as the problems have been typically discussed in that context since the stakes are high, but the issues apply also to other examples as cited earlier. We also take a look at the operational considerations in Section 6 (Operations Considerations) since there is a cost associated with the estbalishment of the necessary infrastructure. With the pros of the available technology being described and the cons of the operational complexity highlighted we offer a conclusion in Section 7 (Conclusion).

2. Terminology

TOC

This document re-uses a lot of the terminology defined in Section 3 of [RFC5012] (Schulzrinne, H. and R. Marshall, “Requirements for Emergency Context Resolution with Internet Technologies,” January 2008.).

3. Emergency Services

TOC

Users of the legacy telephone network can summon emergency services such as ambulance, fire and police using a well-known emergency service number (e.g., 9-1-1 in North America, 1-1-2 in Europe). Location information is used to route emergency calls to the appropriate regional Public Safety Answering Point (PSAP) that serves the caller to dispatch first-level responders to the emergency site.

Regulators have already started to demand emergency service support for voice over IP. However, enabling such critical public services using the Internet is challenging, as many of the assumptions of the public switched telephone network (PSTN) / public land mobile network (PLMN) no longer hold. In particular, while the local telephone company provides both the physical access and the phone service, VoIP allows and encourages to split these two roles between the Access Infrastructure Provider (AIP) and Application (Voice) Service Provider (VSP). The VSP may be located far away from the AIP and may either have no business relationship with that AIP or may be a competitor. It is also likely that the VSP will have no relationship with the PSAP and will therefore be unknown.

4. Threats

TOC

IP-based emergency calling faces many security threats, most of which are well-known from other realms, such as protecting the privacy of communications or against denial-of-service attacks using packet flooding. Here, we focus specifically on a higher-layer threat that is unique to services where semi-anonymous users can request expensive services.

Prank calls have been a problem for emergency services, dating back to the time of street corner call boxes. Individual prank calls waste emergency services and possibly endanger bystanders or emergency service personnel as they rush to the reported scene of a fire or accident. A more recent concern is that massive prank calls can be used to disrupt emergency services, e.g., during a mass-casualty event and thus be used as a means to amplify the effect of a terror attack, for example.

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 trustworthy 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.

Currently, emergency services rely on the fact that location spoofing is difficult for normal users. Additionally, the identity of most callers can be ascertained, so that the threat of severe punishments reduces prank calls. Mechanically placing a large number of emergency calls that appear to come from different locations is also difficult. Calls from payphones are subject to greater scrutiny by the call taker. 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'.

One of the main motivations of an adversary in the emergency services context is to prevent callers from utilizing emergency service support. This can be done by a variety of means, such as impersonating a PSAP or directory servers, attacking SIP signaling elements and location servers.

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 payphones 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 [I-D.ietf-ecrit-mapping-arch] (Schulzrinne, H., “Location-to-URL Mapping Architecture and Framework,” March 2009.) 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.
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.

We will focus only on the malicious end host adversary model since it follows today's most common adversary model on the Internet that includes bot nets.

4.1. Location Spoofing

TOC

An adversary can provide false location information in order to fool the emergency personnel. Such an attack is particularly easy if location information is attached to the emergency call by the end host and is either not verified or cannot be verified by anyone. Only entities that are close to the caller can verify the correctness of location information. Another form of this attack is to fool a VSP (and indirectly a LIS) in using a wrong identity (such as an IP address) for the location lookup. This type of attack can be accomplished in the PSTN today with the help of caller-id spoofing.
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.

4.2. Call Identity Spoofing

TOC

If an adversary can place emergency calls without disclosing its identity, then prank calls are more difficult to be traced. There are at least two different forms of authentication in this context: (a) network access authentication (e.g., using the Extensible Authentication Protocol (EAP) [RFC3748] (Aboba, B., Blunk, L., Vollbrecht, J., Carlson, J., and H. Levkowetz, “Extensible Authentication Protocol (EAP),” June 2004.) and (b) authentication of the emergency caller at the VoIP application layer. This differentiation is created by the split between the AIP and the VSP. Note that different identities are involved and that the are also managed by different parties and thus making the linkage between the two quite difficult.
Trying to find an adversary that did not authenticate itself to the VSP is difficult even though there is still a chance if network access authentication was executed. If there is no authentication (neither to the PSAP, to the VSP nor to the AIP) then it is very challenging to trace the call back in order to a make a particular entity accountable. This might, for example, be the case with an open IEEE 802.11 WLAN access point even if the owner of the access point can be determined. However, unlike for the existing telephone system, it is possible to imagine that VoIP emergency calls could require strong identity, as providing such identity information is not necessarily coupled to having a business relationship with the AIP, ISP or VSP. However, due to the time-critical nature of emergency calls, it is unlikely that multi-layers authentication can be used, so that in most cases, only the device placing the call will be able to be identified, making the system vulnerable to botnet attacks. Furthermore, deploying additional credentials for emergency service purposes, such as dedicated certificates, increases costs, introduces a significant administrative overhead and is only useful if widely used.

5. Solution Proposals

TOC

This section presents three solution approaches that have been discussed in order to mitigate the threats discussed.

5.1. Location Signing

TOC

One way to avoid location spoofing is to let a trusted location server sign the location information before it is sent to the end host, i.e., 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 (Location Signing) 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 [I-D.ietf-sip-location-conveyance] (Polk, J. and B. Rosen, “Location Conveyance for the Session Initiation Protocol,” March 2009.) 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

Additional information, such as timestamps or expiration times, has to be included together with the signed location to limit replay attacks. 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.
Bot nets are also unlikely to be deterred by location signing. However, accurate location information would limit the usable subset of the bot net, as only hosts within the PSAP serving area would be useful in placing calls.

To prevent location-swapping attacks it is necessary to include some some target specific identity information. The included information depends on the purpose, namely either real-time verification by the location recipient or for the purpose of a post-mortem analysis when the location recipient wants to determine the legal entity behind the target for prosecution (if this is possible). As argued in Section 6 (Operations Considerations) the operational considerations make a real-time verification difficult. A strawman proposal for location signing is provided by [I-D.thomson-geopriv-location-dependability] (Thomson, M. and J. Winterbottom, “Digital Signature Methods for Location Dependability,” January 2010.).

Still, for large-scale attacks launched by bot nets, this is unlikely to be helpful. Location signing is also difficult when the host provides its own location via GPS, which is likely to be a common occurrence for mobile devices. Trusted computing approaches, with tamper-proof GPS modules, may be needed in that case. After all, a device can always pretend to have a GPS device and the recipient has no way of verifying this or forcing disclosure of non-GPS-derived location information.

Location verification may be most useful if it is used in conjunction with other mechanisms. For example, a call taker can verify that the region that corresponds to the IP address of the media stream roughly corresponds to the location information reported by the caller. To make the use of bot nets more difficult, a CAPTCHA-style test may be applied to suspicious calls, although this idea is quite controversial for emergency services, at the danger of delaying or even rejecting valid calls.

5.2. Location by Reference

TOC

The location-by-reference concept was developed so that end hosts could 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.

Figure 2 (Location by Reference) 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 [I-D.ietf-sip-location-conveyance] (Polk, J. and B. Rosen, “Location Conveyance for the Session Initiation Protocol,” March 2009.) 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

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. HTTP-Enabled Location Delivery (HELD) [I-D.ietf-geopriv-http-location-delivery] (Barnes, M., Winterbottom, J., Thomson, M., and B. Stark, “HTTP Enabled Location Delivery (HELD),” August 2009.) is an example of a protocol that is able to return such references.

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.

However, location-by-reference does not offer significant security benefits if the end host uses GPS to determine its location. At best, a network provider can use cell tower or triangulation information to limit the inaccuracy of user-provided location information.

5.3. Proxy Adding Location

TOC

Instead of making location information available to the end host, it is possible to allow an entity in the AIP, or associated with the AIP, to retrieve the location information on behalf of the end point. This solution is possible when the application layer messages are routed through an entity with the ability to determine the location information of the end point, for example based on the end host's IP or MAC address.

When the untrustworthy end host does not have the ability to access location information, it cannot modify it either. Proxies can use various authentication security techniques, including SIP Identity [RFC4474] (Peterson, J. and C. Jennings, “Enhancements for Authenticated Identity Management in the Session Initiation Protocol (SIP),” August 2006.), to ensure that modifications to the location in transit can be detected by the location recipient (e.g., the PSAP). As noted above, this is unlikely to work for GPS-based location determination techniques.

The obvious disadvantage of this approach is that there is a 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, for example, based on SIP. This might impose a certain interoperability challenge. Additionally, the AIP more or less takes the 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] (Garcia-Martin, M., Belinchon, M., Pallares-Lopez, M., Canales-Valenzuela, C., and K. Tammi, “Diameter Session Initiation Protocol (SIP) Application,” November 2006.)). This interaction along the Authentication, Authorization and Accounting infrastructure (see) 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 the procedures described in [I-D.schulzrinne-ecrit-unauthenticated-access] (Schulzrinne, H., McCann, S., Bajko, G., Tschofenig, H., and D. Kroeselberg, “Extensions to the Emergency Services Architecture for dealing with Unauthenticated and Unauthorized Devices,” March 2010.) are applicable and 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.

6. Operations Considerations

TOC

6.1. Attribution to a Specific Trusted Source

TOC

[NENA-i2] (, “08-001 NENA Interim VoIP Architecture for Enhanced 9-1-1 Services (i2),” December 2005.) Section 3.7 describes some of the aspects of attribution as follows:

> The i2 solution proposes a Location Information Server (LIS) be the source for distributing location information within an access network. Furthermore the validity, integrity and authenticity of this information are directly attributed to the LIS operator.

Section 6.1.1 (Validity) describes the issues that arise in ensuring the validity of location information provided by the LIS operator. Section 6.1.2 (Location Signing) and Section 6.1.3 (Location by Reference) describe operational issues that arise in ensuring the integrity and authenticity of location information provided by the LIS operator.

6.1.1. Validity

TOC

In existing networks where location information is both determined by the access/voice service provider as well as communicated by the AIP/VSP, responsibility for location validity can be attributed entirely to a single party, namely the AIP/VSP.
However, on the Internet, not only may the AIP and VSP represent different parties, but location determination may depend on information contributed by parties trusted by neither the AIP nor VSP, or even the operator of the Location Information Server (LIS). In such circumstances, mechanisms for enhancing the integrity or authenticity of location data contribute little toward ensuring the validity of that data.
It should be understood that the means by which location is determined may not necessarily relate to the means by which the endpoint communicates with the LIS. Just because a Location Configuration Protocol (LCP) operates at a particular layer does not imply that the location data communicated by that protocol is derived solely based on information obtained at that layer. In some circumstances, LCP implementations may base their location determination on information gathered from a variety of sources which may merit varying levels of trust, such as information obtained from the calling endpoint, or

wiremap information that is time consuming to verify or may rapidly go out of date.

For example, consider the case of a Location Information Server (LIS) that utilizes LLDP-MED [LLDP-MED] (, “Telecommunications: IP Telephony Infrastructure: Link Layer Discovery Protocol for Media Endpoint Devices, ANSI/TIA-1057-2006,” April 2006.) endpoint move detection notifications in determining calling endpoint location. Regardless of whether the LIS implementation utilizes an LCP operating above the link layer (such as an application layer protocol such as HELD [I-D.ietf-geopriv-http-location-delivery] (Barnes, M., Winterbottom, J., Thomson, M., and B. Stark, “HTTP Enabled Location Delivery (HELD),” August 2009.)), the validity of the location information conveyed would be dependent on the security properties of LLDP-MED.

[LLDP-MED] (, “Telecommunications: IP Telephony Infrastructure: Link Layer Discovery Protocol for Media Endpoint Devices, ANSI/TIA-1057-2006,” April 2006.) Section 13.3 defines the endpoint move detection notification as follows:

```
        lldpXMedTopologyChangeDetected NOTIFICATION-TYPE
             OBJECTS { lldpRemChassisIdSubtype,
                       lldpRemChassisId,
                       lldpXMedRemDeviceClass
                     }
                   STATUS current
             DESCRIPTION
                       "A notification generated by the local device
                        sensing a change in the topology that
                        indicates a new remote device attached to a
                        local port, or a remote device disconnected
                        or moved from one port to another."
                    ::= { lldpXMedNotifications 1 }
```

Figure 3: Interworking Architecture

As noted in Section 7.4 of [LLDP-MED] (, “Telecommunications: IP Telephony Infrastructure: Link Layer Discovery Protocol for Media Endpoint Devices, ANSI/TIA-1057-2006,” April 2006.), the lldpRemChassisIdSubtype, lldpRemChassisId and lldpXMedRemDeviceClass variables are determined from the Chassis ID (1) and LLDP-MED Device Type Type-Length-Value (TLV) tuples provided within the LLDP advertisement of the calling device. As noted in [LLDP-MED] (, “Telecommunications: IP Telephony Infrastructure: Link Layer Discovery Protocol for Media Endpoint Devices, ANSI/TIA-1057-2006,” April 2006.) Section 9.2.3, all Endpoint Devices use the Network address ID subtype (5) by default. In order to provide topology change notifications in a timely way, it cannot necessarily be assumed that a Network Connectivity devices will validate the network address prior to transmission of the move detection notification. As a result, there is no guarantee that the network address reported by the endpoint will correspond to that utilized by the device.

The discrepancy need not be due to nefarious reasons. For example, an IPv6-capable endpoint may utilize multiple IPv6 addresses. Similarly, an IPv4-capable endpoint may initially utilize a Link- Local IPv4 address [RFC3927] (Cheshire, S., Aboba, B., and E. Guttman, “Dynamic Configuration of IPv4 Link-Local Addresses,” May 2005.) and then may subsequently acquire a DHCP-assigned routable address. All addresses utilized by the endpoint device may not be advertised in LLDP, or even if they are, endpoint move detection notification may not be triggered, either because no LinkUp/LinkDown notifications occur (e.g. the host adds or changes an address without rebooting) or because these notifications were not detectable by the Network Connectivity device (the endpoint device was connected to a hub rather than directly to a switch).

Similar issues may arise in situations where the LIS utilizes DHCP lease data to obtain location information. Where the endpoint address was not obtained via DHCP (such as via manual assignment, stateless autoconfiguration [RFC4862] (Thomson, S., Narten, T., and T. Jinmei, “IPv6 Stateless Address Autoconfiguration,” September 2007.) or Link-Local IPv4 self- assignment), no lease information will be available to enable determination of device location. This situation should be expected to become increasingly common as IPv6-capable endpoints are deployed, and Location Configuration Protocol (LCP) interactions occur over IPv6.

Even in scenarios in which the LIS relies on location data obtained from the IP MIB [RFC4293] (Routhier, S., “Management Information Base for the Internet Protocol (IP),” April 2006.) and the Bridge MIB [RFC4188] (Norseth, K. and E. Bell, “Definitions of Managed Objects for Bridges,” September 2005.), availability of location determination information is not assured. In an enterprise scale network, maintenance of current location information depends on the ability of the management station to retrieve data via polling of network devices. As the number of devices increases, constraints of network latency and packet loss may make it increasingly difficult to ensure that all devices are polled on a sufficiently frequent interval. In addition, in large networks, it is likely that tables will be large so that when UDP transport is used, query responses will fragment, resulting in increasing packet loss or even difficulties in firewall or NAT traversal.

Furthermore, even in situations where the location data can be presumed to exist and be valid, there may be issues with the integrity of the retrieval process. For example, where the LIS depends on location information obtained from a MIB notification or query, unless SNMPv3 [RFC3411] (Harrington, D., Presuhn, R., and B. Wijnen, “An Architecture for Describing Simple Network Management Protocol (SNMP) Management

Frameworks,” December 2002.) is used, data integrity and authenticity is not assured in transit between the network connectivity device and the LIS.

From these examples, it should be clear that the availability or validity of location data is a property of the LIS system design and implementation rather than an inherent property of the LCP. As a result, mechanisms utilized to protect the integrity and authenticity of location data do not necessarily provide assurances relating to the validity or provenance of that data.

6.1.2. Location Signing

TOC

[NENA-i2] (, “08-001 NENA Interim VoIP Architecture for Enhanced 9-1-1 Services (i2),” December 2005.) Section 3.7 includes 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. Antispoofing mechanisms should be applied to the Location Reporting method.

[Note: The term Valid Emergency Services Authority (VESA) refers to the root certificate authority.]

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.

TOC

6.1.3. Location by Reference

Where location by reference is provided, the recipient needs to deference the LbyR in order to obtain location. With the introduction of location by reference concept two authorization models were developed, see [I-D.winterbottom-geopriv-deref-protocol] (Winterbottom, J., Tschofenig, H., Schulzrinne, H., Thomson, M., and M. Dawson, “A Location Dereferencing Protocol Using HELD,” January 2010.), namely the "Authorization by Possession" and "Authorization via Access Control Lists" model. With the "Authorization by Possession" model everyone in possession of the reference is able to obtain the corresponding location information. This might, however, be incompatible with other requirements typically imposed by AIPs, such as location hiding (see [I-D.ietf-ecrit-location-hiding-req] (Schulzrinne, H., Liess, L., Tschofenig, H., Stark, B., and A. Kuett, “Location Hiding: Problem Statement and Requirements,” February 2010.)). As such, the "Authorization via Access Control Lists" model is likely to be the preferred model for many AIPs and subject for discussion in the subsequent paragraphs.
Just as with PIDF-LO signing, the operational considerations in managing credentials for use in LbyR dereferencing can be considerable without the introduction of some kind of hierarchy. It does not seem reasonable for a PSAP to manage client certificates or Digest credentials for all the LISes in its coverage area, so as to enable it to successfully dereference LbyRs. In some respects, this issue is even more formidable than the validation of signed PIDF- LOs. While PIDF-LO signing credentials are provided to the LIS operator, in the case of de-referencing, the PSAP needs to be obtain credentials compatible with the LIS configuration, a potentially more complex operational problem.
As with PIDF-LO signing, the operational issues of LbyR can be addressed to some extent by introduction of hierarchy. Rather than requiring the PSAP to obtain credentials for accessing each LIS, the local LIS could be required to upload location information to location aggregation points who would in turn manage the relationships with the PSAP. This would shift the management burden from the PSAPs to the location aggregation points.

TOC

6.2. Application to a Specific Point in Time

PIDF-LO objects contain a timestamp, which reflects the time at which the location was determined. 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 LISs. As a result, recipients need to validate the timestamp in order to determine whether it is credible.

6.3. Linkage to a Specific Endpoint

TOC

As noted in the "HTTP Enabled Location Delivery (HELD)" [I-D.ietf-geopriv-http-location-delivery] (Barnes, M., Winterbottom, J., Thomson, M., and B. Stark, “HTTP Enabled Location Delivery (HELD),” August 2009.) 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] (Cuellar, J., Morris, J., Mulligan, D., Peterson, J., and J. Polk, “Geopriv Requirements,” February 2004.). 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" element and the "deviceID" element [RFC4479] (Rosenberg, J., “A Data Model for Presence,” July 2006.) are not used.

Given the restrictions on inclusion of identification information within the PIDF-LO, it may not be possible for a recipient to verify that the entity on whose behalf location was determined represents the same entity conveying location to the recipient.
Where "Enhancements for Authenticated Identity Management in the Session Initiation Protocol (SIP)" [RFC4474] (Peterson, J. and C. Jennings, “Enhancements for Authenticated Identity Management in the Session Initiation Protocol (SIP),” August 2006.) is used, it is possible for the recipient to verify the identity assertion in the From: header. However, if PIDF parameters that contain identity are omitted or contain an unlinked pseudonym, then it may not be possible for the recipient to verify whether the conveyed location actually relates to the entity identified in the From: header.

This lack of binding between the entity obtaining the PIDF-LO and the entity conveying the PIDF-LO to the recipient enables cut and paste attacks which would enable an attacker to assert a bogus location, even where both the SIP message and PIDF-LO are signed. As a result, even implementation of both [RFC4474] (Peterson, J. and C. Jennings, “Enhancements for Authenticated Identity Management in the Session Initiation Protocol (SIP),” August 2006.) and location signing does not guarantee that location can be tied to a specific endpoint.

---

7. Conclusion

TOC

Emergency services raise a number of architectural questions, see [I-D.ietf-ecrit-framework] (Rosen, B., Schulzrinne, H., Polk, J., and A. Newton, “Framework for Emergency Calling using Internet Multimedia,” July 2009.), [I-D.schulzrinne-ecrit-unauthenticated-access] (Schulzrinne, H., McCann, S., Bajko, G., Tschofenig, H., and D. Kroeselberg, “Extensions to the Emergency Services Architecture for dealing with Unauthenticated and Unauthorized Devices,” March 2010.), [I-D.ietf-ecrit-location-hiding-req] (Schulzrinne, H., Liess, L., Tschofenig, H., Stark, B., and A. Kuett, “Location Hiding: Problem Statement and Requirements,” February 2010.), and [I-D.tschofenig-ecrit-architecture-overview] (Tschofenig, H. and H. Schulzrinne, “Emergency Services Architecture Overview: Sharing Responsibilities,” July 2007.). With the generalized emergency architecture considered within the ECRIT working group various security challenges need to be addressed, including the ability to report faked location and other attacks against the emergency services infrastructure. These types of attacks also show that the attack characteristics play an important role when dealing with the problems and lower-layer solutions, as they have been proposed as solutions for Denial of Service prevention (for example using cryptographic puzzles), have limited applicability.

Although it is important to ensure that location information cannot be faked there will be a larger number of GPS-enabled devices out there that make it difficult to utilize any of the security mechanisms described in Section 5 (Solution Proposals). It will be very unlikely that end users will upload their location information for "verification" to a nearby location server located in the access network.

Given the practical and operational limitations in the technology, it may be worthwhile to consider whether the goals of trustworthy location, as for example defined by NENA i2 [NENA-i2] (, “08-001 NENA Interim VoIP Architecture for Enhanced 9-1-1 Services (i2),” December 2005.), are attainable, or whether lesser goals (such as auditability) should be substituted instead.

The goal of auditability is to enable an investigator to determine the source of a rogue emergency call after the fact. Since such an investigation can rely on audit logs provided under court order, the information available to the investigator could be considerably greater than that present in messages conveyed in the emergency call. As a consequence the emergency caller becomes accountable for his actions. For example, in such a situation, 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. Auditability is likely to be of most benefits 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.
Where attacks are frequent and continuous, a reliance on non-automated mechanisms is unlikely to be satisfactory. As such, 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 might be valuable for the emergency services community.

8. IANA Considerations

TOC

This document does not require actions by IANA.

9. Acknowledgments

TOC

We would like to thank the members of the IETF ECRIT and the IETF GEOPRIV working group for their input to the discussions related to this topic. We would also like to thank Andrew Newton, Murugaraj Shanmugam, Richard Barnes and Matt Lepinski for their feedback to previous versions of this document. Martin Thomson provided valuable input to version -02 of this document.

10. References

TOC

10.1. Normative References

TOC

[RFC5012]

Schulzrinne, H. and R. Marshall, “Requirements for Emergency Context Resolution with Internet Technologies,” RFC 5012, January 2008 (TXT).

10.2. Informative references

TOC

[I-D.ietf-ecrit-framework] Rosen, B., Schulzrinne, H., Polk, J., and A. Newton, “Framework for Emergency Calling using Internet Multimedia,” draft-ietf-ecrit-framework-10 (work in progress), July 2009 (TXT).

[I-D.ietf-ecrit-location-hiding-req] Schulzrinne, H., Liess, L., Tschofenig, H., Stark, B., and A. Kuett, “Location Hiding: Problem Statement and Requirements,” draft-ietf-ecrit-location-hiding-req-04 (work in progress), February 2010 (TXT).

[I-D.ietf-ecrit-mapping-arch] Schulzrinne, H., “Location-to-URL Mapping Architecture and Framework,” draft-ietf-ecrit-mapping-arch-04 (work in progress), March 2009 (TXT).

[I-D.ietf-geopriv-http-location-delivery] Barnes, M., Winterbottom, J., Thomson, M., and B. Stark, “HTTP Enabled Location Delivery (HELD),” draft-ietf-geopriv-http-location-delivery-16 (work in progress), August 2009 (TXT).

[I-D.ietf-sip-location-conveyance] Polk, J. and B. Rosen, “Location Conveyance for the Session Initiation Protocol,” draft-ietf-sip-location-conveyance-13 (work in progress), March 2009 (TXT).

[I-D.schulzrinne-ecrit-unauthenticated-access] Schulzrinne, H., McCann, S., Bajko, G., Tschofenig, H., and D. Kroeselberg, “Extensions to the Emergency Services Architecture for dealing with Unauthenticated and Unauthorized Devices,” draft-schulzrinne-ecrit-unauthenticated-access-07 (work in progress), March 2010 (TXT).

[I-D.thomson-geopriv-location-dependability] Thomson, M. and J. Winterbottom, “Digital Signature Methods for Location Dependability,” draft-thomson-geopriv-location-dependability-05 (work in progress), January 2010 (TXT).

[I-D.tschofenig-ecrit-architecture-overview] Tschofenig, H. and H. Schulzrinne, “Emergency Services Architecture Overview: Sharing Responsibilities,” draft-tschofenig-ecrit-architecture-overview-00 (work in progress), July 2007 (TXT).

Winterbottom, J., Tschofenig, H., Schulzrinne, H., Thomson, M., and M. Dawson, “A Location

[I-D.winterbottom-geopriv-deref-protocol]	Dereferencing Protocol Using HELD," draft-winterbottom-geopriv-deref-protocol-05 (work in progress), January 2010 (TXT).
[LLDP-MED]	"Telecommunications: IP Telephony Infrastructure: Link Layer Discovery Protocol for Media Endpoint Devices, ANSI/TIA-1057-2006," April 2006.
[NENA-i2]	"08-001 NENA Interim VoIP Architecture for Enhanced 9-1-1 Services (i2)," December 2005.
[RFC3411]	Harrington, D., Presuhn, R., and B. Wijnen, "An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks," STD 62, RFC 3411, December 2002 (TXT).
[RFC3693]	Cuellar, J., Morris, J., Mulligan, D., Peterson, J., and J. Polk, "Geopriv Requirements," RFC 3693, February 2004 (TXT).
[RFC3748]	Aboba, B., Blunk, L., Vollbrecht, J., Carlson, J., and H. Levkowetz, "Extensible Authentication Protocol (EAP)," RFC 3748, June 2004 (TXT).
[RFC3825]	Polk, J., Schnizlein, J., and M. Linsner, "Dynamic Host Configuration Protocol Option for Coordinate-based Location Configuration Information," RFC 3825, July 2004 (TXT).
[RFC3927]	Cheshire, S., Aboba, B., and E. Guttman, "Dynamic Configuration of IPv4 Link-Local Addresses," RFC 3927, May 2005 (TXT).
[RFC4188]	Norseth, K. and E. Bell, "Definitions of Managed Objects for Bridges," RFC 4188, September 2005 (TXT).
[RFC4293]	Routhier, S., "Management Information Base for the Internet Protocol (IP)," RFC 4293, April 2006 (TXT).
[RFC4474]	Peterson, J. and C. Jennings, "Enhancements for Authenticated Identity Management in the Session Initiation Protocol (SIP)," RFC 4474, August 2006 (TXT).
[RFC4479]	Rosenberg, J., "A Data Model for Presence," RFC 4479, July 2006 (TXT).
[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 (TXT).
[RFC4776]	Schulzrinne, H., "Dynamic Host Configuration Protocol (DHCPv4 and DHCPv6) Option for Civic

	Addresses Configuration Information," RFC 4776, November 2006 (TXT).
[RFC4862]	Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration," RFC 4862, September 2007 (TXT).

Authors' Addresses

TOC

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: bernarda@microsoft.com