Network Working Group
Internet-Draft
Intended status: Informational
Expires: April 26, 2007

H. Tschofenig
Siemens Networks GmbH & Co KG
H. Schulzrinne
Columbia U.
October 23, 2006

GEOPRIV Layer 7 Location Configuration Protocol; Problem Statement and Requirements

draft-tschofenig-geopriv-l7-lcp-ps-03.txt

Status of this Memo

Copyright Notice

Abstract

This document provides a problem statement, lists requirements and captures discussions for a GEOPRIV Layer 7 Location Configuration Protocol (LCP). This protocol aims to allow an end host to obtain location information, by value or by reference, from a Location Information Server (LIS) that is located in the access network. The obtained location information can then be used for a variety of different protocols and purposes. For example, it can be used as input to the Location-to-Service Translation Protocol (LoST) or to convey location within SIP to other entities.

Table of Contents

1. Introduction

This document provides a problem statement, lists requirements and captures discussions for a GEOPRIV Layer 7 Location Configuration Protocol (LCP). The protocol has two purposes:

- It is used to obtain location information from a special node, called the Location Information Server (LIS).

- It enables the end host to obtain a reference to location information. This reference can take the form of a subscription URI, such as a SIP presence URI, an HTTP/HTTPS URI, or any others.

The need for these two functions can be derived from the scenarios presented in Section 3.

This document splits the problem space into separate parts and discusses them in separate subsections. Section 4 discusses the challenge of discovering the Location Information Server in the access network. Section 5 compares different types of identifiers that can be used to retrieve location information. The concept of subscription URIs is described in Section 7. Digitally signing location information and the perceived benefits are covered in Section 8. A list of requirements for the GEOPRIV Layer 7 Location Configuration Protocol can be found in Section 9. This work is heavily influenced by security considerations. Hence, almost all sections address security concerns. A list of desired security properties can be found in Section 10 together with a discussion about possible threat models.

This document does not describe how the access network provider determines the location of the end host.

2. Terminology

In this document, the key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" are to be interpreted as described in RFC 2119 [1], with the qualification that unless otherwise stated these words apply to the design of the GEOPRIV Layer 7 Location Configuration Protocol.

We also use terminology from [2] and [3].

3. Scenarios

The following network types are within scope:

- DSL/Cable networks, WiMax-like fixed access
- Airport, City, Campus Wireless Networks, such as 802.11a/b/g, 802.16e/Wimax
- 3G networks
- Enterprise networks

We illustrate a few examples below.

3.1. Fixed Wired Environment

The following figure shows a DSL network scenario with the Access Network Provider and the customer premises. The Access Network Provider operates link and network layer devices (represented as Node) and the Location Information Server (LIS).

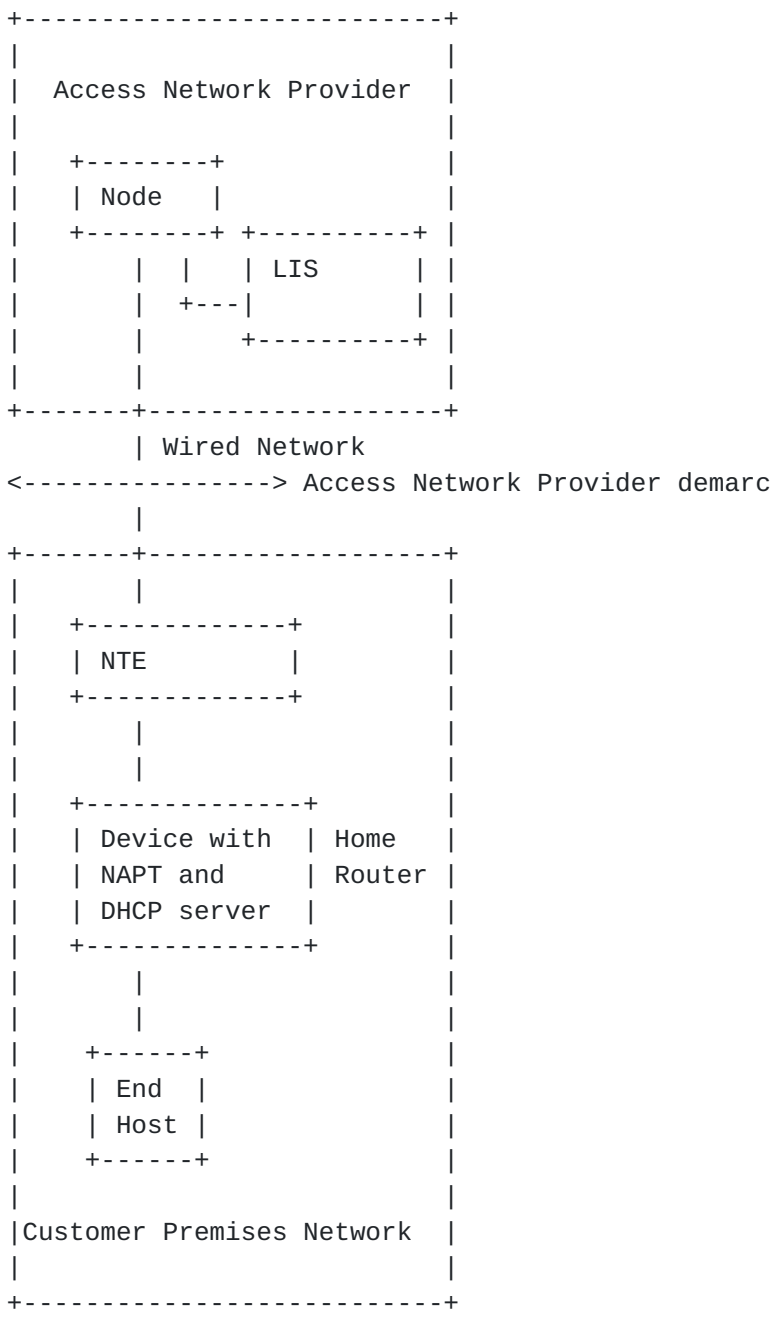

Figure 1: DSL Scenario

The customer premises consists of a router with a Network Address Translator with Port Address Translation (NAPT) and a DHCP server as used in most Customer Premises Networks (CPN) and the Network Termination Equipment (NTE) where Layer 1 and sometimes Layer 2 protocols are terminated. The router in the home network (e.g., broadband router, cable or DSL router) typically runs a NAPT and a DHCP server. The NTE is a legacy device and in many cases cannot be modified for the purpose of delivering location information to the

end host. The same is true of the device with the NAPT and DHCP server.

It is possible for the NTE and the home router to physically be in the same box, or for there to be no home router, or for the NTE and end host to be in the same physical box (with no home router). An example of this last case is where Ethernet service is delivered to customers' homes, and the Ethernet NIC in their PC serves as the NTE.

Current Customer Premises Network (CPN) deployments frequently show the following characteristics:

1. CPE = Single PC

    1. with Ethernet NIC [PPPoE or DHCP on PC]; there may be a bridged DSL or cable modem as NTE, or the Ethernet NIC might be the NTE

    2. with USB DSL or cable modem [PPPoA, PPPoE, or DHCP on PC]

    Note that the device with NAPT and DHCP of Figure 1 is not present in such a scenario.

2. One or more hosts with at least one router [DHCP Client or PPPoE, DHCP server in router; VoIP can be soft client on PC, stand-alone VoIP device, or Analog Terminal Adaptor (ATA) function embedded in router]

    1. combined router and NTE

    2. separate router with NTE in bridged mode

    3. separate router with NTE [NTE/router does PPPoE or DHCP to WAN, router provides DHCP server for hosts in LAN; double NAT

The majority of fixed access broadband customers use a router. The placement of the VoIP client is mentioned to describe what sorts of hosts may need to be able to request location information. Soft clients on PCs are frequently not launched until long after bootstrap is complete, and are not able to control any options that may be specified during bootstrap. They also cannot control whether a VPN client is operating on the PC.

3.2. Moving Network

An example of a moving network is a "WIMAX-like fixed wireless" scenario that is offered in several cities, like New Orleans, Biloxi, etc., where much of the communications infrastructure was destroyed

due to a natural disaster. The customer-side antenna for this service is rather small (about the size of a mass market paperback book) and can be run off battery power. The output of this little antenna is a RJ-45 Ethernet jack. A laptop can be plugged into this Ethernet jack. The user would then run a PPPoE client to connect to the network. Once the network connection is established, the user can run a SIP client on the laptop.

The network-side antenna is, for example, connected through ATM to the core network, and from there to the same BRASs that serve regular DSL customers. These Broadband Remote Access Servers (BRASs) terminate the PPPoE sessions, just like they do for regular DSL.

The laptop and SIP client are, in this case, unaware that they are "mobile". All they see is an Ethernet connection, and the IP address they get from PPPoE does not change over the coverage area. Only the user and the network are aware of the laptop's mobility.

Further examples of moving networks can be found in busses, trains, airplanes.

Figure 2 shows an example topology for a moving network.

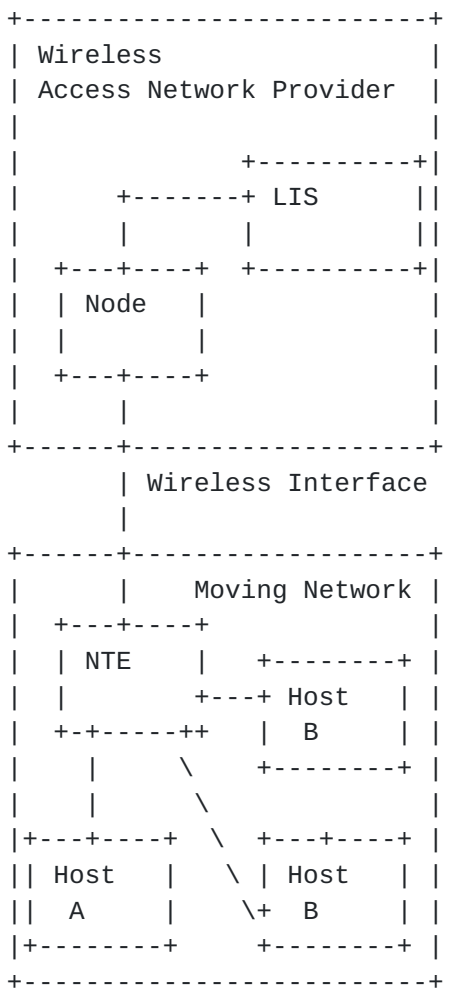

Figure 2: Moving Network

3.3. Wireless Access

Figure 3 shows a wireless access network where a moving end host obtains location information or references to location information from the LIS. The access equipment us, in many cases, link layer devices. This figure represents a hotspot network found in hotels, airports, coffee shops.

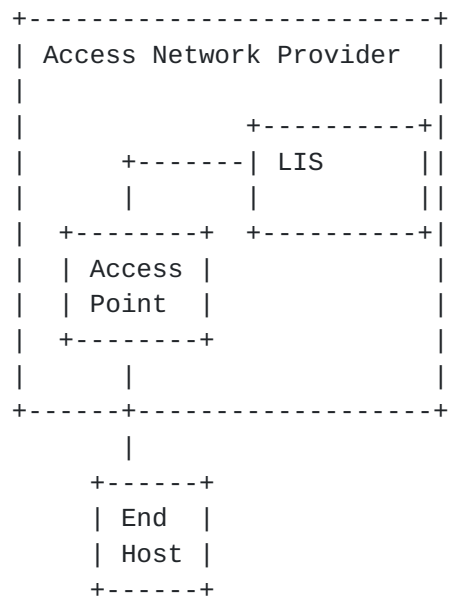

Figure 3: Wireless Access Scenario

4. Discovery of the Location Information Server

When an end host wants to retrieve location information from the LIS it first needs to discover it. Based on the problem statement of determining the location of the end host, which is known best by entities close to the end host itself, we assume that the LIS is located in the access network. Several procedures have been investigated that aim to discovery the LIS in such an access network.

DHCP-based Discovery:

In some environments the Dynamic Host Configuration Protocol might be a good choice for discovering the FQDN or the IP address of the LIS. In environments where DHCP can be used it is also possible to use the already defined location extensions. In environments with legacy devices, such as the one shown in Section 3.1, a DHCP based discovery solution is not possible.

DNS-based Discovery:

With this idea the end host obtains its public IP address (e.g., via STUN [4]) in order to obtain its domain name (via the usual reverse DNS lookup). Then, the SRV or NAPTR record for that domain is retrieved. This relies on the user's public IP address having a DNS entry.

Redirect Rule:

A redirect rule at a device in the access network, for example at the AAA client, will be used to redirect the Geopriv-L7 signalling messages (destined to a specific port) to the LIS. The end host could then discover the LIS by sending a packet to almost any address (as long it is not in the local network). The packet would be redirected to the respective LIS being configured. The same procedure is used by captive portals whereby any HTTP traffic is intercepted and redirected.

Multicast Query:

An end node could also discover a LIS by sending a multicast request to a well-known address. An example of such a mechanism is multicast DNS (see [5] and [6]).

The LIS discovery procedure raises deployment and security issues. When an end host discovers a LIS,

1. it does not talk to a man-in-the-middle adversary, and

2. it needs to ensure that the discovered entity is indeed an authorized LIS.

5. Identifier for Location Determination

The LIS returns location information to the end host when it receives a request. Some form of identifier is therefore needed to allow the LIS to determine the current location of the target or a good approximation of it.

The chosen identifier needs to have the following properties:

Ability for end host to learn or know the identifier:

The end host MUST know or MUST be able to learn the identifier (explicitly or implicitly) in order to send it to the LIS. Implicitly refers to the situation where a device along the path between the end host and the LIS modifies the identifier, as it is done by a NAT when an IP address based identifier is used.

Ability to use the identifier for location determination:

The LIS MUST be able to use the identifier (directly or indirectly) for location determination. Indirectly refers to the case where the LIS uses other identifiers locally within the access network, in addition to the one provided by the end host, for location determination.

Security properties of the identifier:

Misuse needs to be minimized whereby off-path adversary MUST NOT be able to obtain location information of other hosts. A on-path adversary in the same subnet SHOULD NOT be able to spoof the identifier of another host in the same subnet.

The problem is further complicated by the requirement that the end host should not be aware of the network topology and the LIS must be placed in such a way that it can determine location information with the available information. As shown in Figure 1 the host behind the NTE/NAPT-DHCP device is not visible to the access network and the LIS itself. In the DSL network environment some identifier used at the NTE is observable for by the LIS/access network.

The following list discusses frequently mentioned identifiers and their properties:

Host MAC address:

The host MAC address is known to the end system, but not carried over an IP hop.

ATM VCI/VPI:

The VPI/VCI is generally only seen by the DSL modem. Almost all routers in the US use 1 of 2 VPI/VCI value pairs: 0/35 and 8/35. This VC is terminated at the DSLAM, which uses a different VPI/VCI (per end customer) to connect to the ATM switch. Only the network provider is able to map VPI/VCI values through its network. With the arrival of VDSL, ATM will slowly be phased out in favor of Ethernet.

Switch/Port Number:

This identifier is available only in certain networks, such as enterprise networks, typically available via proprietary protocols like CDP or, in the future, 802.1ab.

Cell ID:

This identifier is available in cellular data networks and the cell ID might not be visible to the end host.

Authenticated User Identity:

In DSL networks the user credentials are, in many cases, only known by the router and not to the end host. To the network, the authenticated user identity is only available if a network access authentication procedure is executed. In case of roaming it still might not be available to the access network since security protocols might provide user identity confidentiality and thereby hide the real identity of the user allowing the access network to only see a pseudonym or a randomized string.

Host Identifier:

The Host Identifier introduced by the Host Identity Protocol [7] allows identification of a particular host. Unfortunately, the network can only use this identifier for location determination if the operator already stores an mapping of host identities to

location information. Furthermore, there is a deployment problem since the host identities are not used in todays networks.

Cryptographically Generated Address (CGA):

The concept of a Cryptographically Generated Address (CGA) was introduced by [8]. The basic idea is to put the truncated hash of a public key into the interface identifier part of an IPv6 address. In addition to the properties of an IP address it allows a proof of ownership. Hence, a return routability check can be omitted.

Network Access Identifiers:

A Network Access Identifier [9] is only used during the network access authentication procedure in RADIUS [10] or Diameter [11]. Furthermore, in a roaming scenario it does not help the access network to make meaningful decisions since the username part might be a pseudonym and there is no relationship to the end host's location.

Unique Client Identifier

The DSL Forum has defined that all devices that expect to be managed by the TR-069 interface be able to generate an identifier as described in DSL Forum TR-069v2 Section 3.4.4. It also has a requirement that routers that use DHCP to the WAN use RFC 4361 [12] to provide the DHCP server with a unique client identifier. This identifier is, however, not visible to the end host with the assumption of a legacy device like the NTE. If we assume that the LTE can be modified then a number of solutions come to mind including DHCP based location delivery.

IP Address:

The end host's IP address may be used for location determination. This IP address is not visible to the LIS if the end host is behind one or multipel NATs. This is, however, not a problem since the location of a host that is located behind a NAT cannot be determined by the access network. The LIS would in this case only see the public IP address of the NAT binding allocated by the NAT, which is the correct behavior. The property of the IP address for a return routability check is attractive as well to return location information only to a device that transmitted the

request. The LIS receives the request and provides location information back to the same IP address. If an adversary wants to learn location information from an IP address other than its own IP address then it would not see the response message (unless he is on the subnetwork or at a router along the path towards the LIS) since the LIS would return the message to the address where it came from.

On a shared medium an adversary could ask for location information of another host using its IP address. The adversary would be able to see the response message since he is sniffing on the shared medium unless security mechanisms (such as link layer encryption) is in place. With a network deployment as shown in Section 3.1 with multiple hosts in the Customer Premise being behind a NAT the LIS is unable to differentiate the individual end points. For WLAN deployments as found in hotels, as shown in as shown in Section 3.3, it is possible for an adversary to eavesdrop data traffic and subsequently to spoof the IP address in a query to the LIS to learn more detailed location information (e.g., specific room numbers). Such an attack might, for example, compromise the privacy of hotel guests. Note that DHCP would suffer from the same problem here unless each node uses link layer security mechanism.

Return routability checks are useful only if the adversary does not see the response message and if the goal is to delay state establishment. If the adversary is in a broadcast network then a return routability check alone is not sufficient to prevent the above attack since the adversary will observe the response.

6. Virtual Private Network (VPN) Considerations

To establish a VPN, a VPN client uses a particular VPN protocol to create a tunnel to a VPN server. VPNs can be established using a variety of protocols (e.g., IPsec, L2TP). The protocol used to establish the VPN does not impact LIS discovery or location acquisition.

VPN characteristics that can impact LIS discovery or acquiring a location from a LIS include the relationship of the VPN client to the communications application (e.g., VoIP phone), and whether the VPN server requires the device with the VPN client to send all outbound IP traffic across the VPN.

6.1. VPN Tunneled Internet Traffic

Any form of LIS discovery that would work without the VPN being established, will also be able to work after the VPN has been established. The DNS method of LIS discovery requires a device to discover the proper IP address for discovering and querying the LIS. Some devices may be expected to operate in a number of different networks, including corporate networks, hotspots, home networks, and protected networks by way of a VPN. The appropriate IP address to use for LIS discovery may vary depending on the network.

It may be useful for such devices to do a reverse DNS lookup, LIS discovery request, and LIS query for all IP addresses they can determine for themselves. When all LISs involved in these queries are properly configured, only one of these queries should be expected to succeed. LISs should not be configured to provide a location for an IP address that may be used by many geographically dispersed users, or when the LIS has no way to determine the geographic location of the device using the IP address.

This form of VPN will not interfere with queries to the LIS, once the LIS has been discovered. It will also not interfere with location dereferencing.

6.2. VPN Client and End Point Physically Co-Located

If LIS discovery and queries are done prior to establishing the VPN, then the VPN will not interfere. For this reason, it is highly desirable for devices that may support communications applications to do location acquisition shortly after initial bootstrap, and prior to establishing any VPN. As the communication application may not be running prior to establishing the VPN, it is best if the communication application is not responsible for location acquisition.

Once a VPN has been established, the device should not permit location acquisition to be attempted. Location acquisition done after a VPN is established will either fail, or provide the wrong location.

If the device does allow attempts at location acquisition after establishing the VPN, these attempts should fail. LIS discovery through DHCP, Redirect, and Multicast methods would fail due to lack of support by the VPN server (it is undesirable for a VPN server to support LIS discovery). For DNS discovery, the device might know a variety of IP addresses, such as the IP address obtained at bootstrap (which may be public or private, depending on whether the device is behind a NAT), the VPN IP address, and an IP address the VPN provider uses for Internet traffic through its firewall. RDNS of private LAN addresses will fail. Success for RDNS of the VPN address would depend on whether there are entries in the VPN provider's DNS server. If RDNS of the VPN IP address succeeds, and the VPN provider has a LIS in their network, LIS discovery of the VPN network's LIS should succeed. It is desirable for a LIS that may get queries from devices entering the network through a VPN, to provide an error response to location queries that use such IP addresses. The LIS should not be configured to return a location for these IP addresses.

RDNS of public IP addresses should generally succeed (assuming the VPN provider's DNS allows for these queries to succeed). For IP addresses used to connect the VPN network to the Internet, the returned domain of RDNS would be the owner of that IP address, which is either the VPN provider or its ISP. If the domain is that of the VPN provider, the VPN provider may or may not have a DNS LIS entry associated with that domain. If there is a LIS, that LIS should not be configured to return a location for its public IP addresses. If an ISP owns the domain of the VPN's public IP address, the device will discover the ISP's LIS, and that LIS will return the location where traffic from that IP address enters the access network. If the device knows its public IP address, and RDNS and LIS discovery succeeded, the LIS would not provide location information (assuming the LIS would not be able to authenticate the device through means other than return routability). The message that reached the LIS would not be using (in the IP Header) an IP address from its domain.

If the private network allows traffic to go to the Internet, dereferencing of a location reference will work.

6.3. VPN Client and End Point Physically Separated

In this case, it is possible for the device with the VPN client to participate in the location acquisition process, and to provide location to end devices. If the VPN client device does participate,

then it must acquire location information before setting up its VPN.

If the VPN client device that participates in location acquisition is also the DHCP server for the LAN, then it would be able to either provide its location by DHCP, or provide itself as the LIS by DHCP. If this device names itself as the DNS server for devices in the LAN, then it could support RDNS for LAN addresses and provide itself as the LIS. If it says it is the LIS, then it must be able to respond to LIS queries for location acquisition. This device would also be able to support Redirect or Multicast methods of LIS determination.

If the VPN client device does not participate in location acquisition, then location acquisition will either fail or provide the wrong location, with the same results as described in section X.2 for a device that attempts location acquisition after establishing a VPN.

If the private network allows traffic to go to the Internet, dereferencing of a location reference will work.

7. Location-by-Reference and Location Subscriptions

In mobile wireless networks it is not efficient for the end host to periodically query the LIS for up-to-date location information. Furthermore, the end host might want to delegate the task of retrieving and publishing location information to a third party, such as a presence server. Finally, in some deployments the network operator might not want to make location information available to the end hosts.

These usage scenarios motivated the introduction of the location-by-reference concept. Depending on the type of reference, such as HTTP/HTTPS or SIP presence URI, different operations can be performed. While an HTTP/HTTPS URI can be resolved to location information, a SIP presence URI provides further benefits from the SUBSCRIBE/NOTIFY concept that can additionally be combined with location filters [13].

Figure 4 shows the assumed communication model:

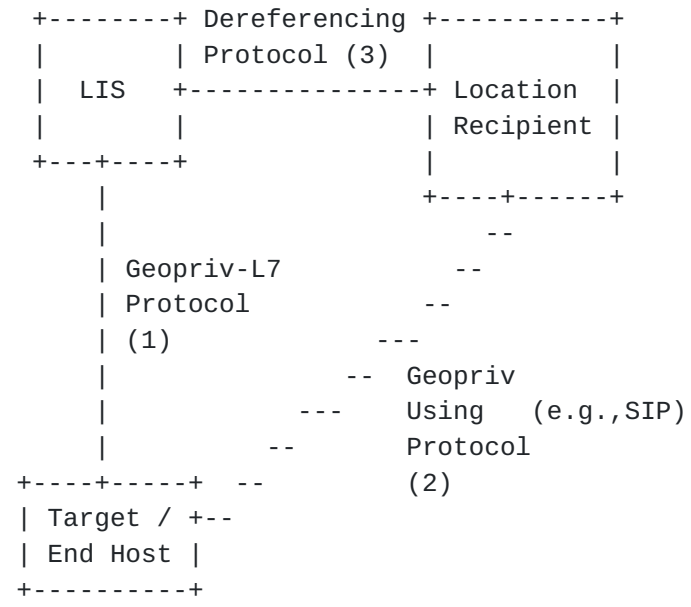

Figure 4: Communication Model

Note that there is no requirement for using the same protocol in (1) and (3).

The following list describes the location subscription idea:

1. The end host discovers the LIS.

2. The end host sends a request to the LIS asking for a location-by-reference, as shown in (1) of Figure 4.

3. The LIS responds to the request and includes a location object together with a subscription URI.

4. The Target puts the subscription URI into a SIP message as described in [14] forwards it to a Location Recipient, as shown in (2) of Figure 4. The Location Recipient subscribes to the obtained subscription URI (see (3) of Figure 4) and potentially uses a location filter (see [13]) to limit the notification rate.

5. If the Target moves outside a certain area, indicated by the location filter, then the Location Recipient will receive a notification.

Note that the Target may also act in the role of the Location Recipient whereby it would subscribe to its own location information. For example, the Target obtains a subscription URI from the Geopriv-L7 protocol. It subscribes to the URI in order to obtain its currently location information, which then serves as input to a LoST query (see [15]) in order to acquire the service boundary (e.g., PSAP boundary). The service boundary indicates the region where the device can move without the need to re-query since the returned answer remains unchanged. The Target uses this service boundary to location filters an updates the subscription. If the Target moves outside a certain area, indicated by the location filter, it will receive a notification and knows that re-querying LoST to obtain a new service boundary is necessary.

For location-by-reference, the LIS needs to maintain a list of randomized URIs for each host, timing out these URIs after the reference expires. References need to expire to prevent the recipient of such a URL from being able to (in some cases) permanently track a host. Furthermore, this mechanism also offers garbage collection capability for the LIS.

Location references must prevent adversaries from obtaining the Target's location. There are at least two approaches: The location reference contains a random component and allows any holder of the reference to obtain location information. Alternatively, the reference can be public and the LIS performs access control via a separate authentication mechanism, such as HTTP digest or TLS client side authentication, when resolving the reference to a location object.

8. Preventing Faked Location based DoS Attacks

This section describes a possible security threat in emergency related location conveyance and subsequently discusses countermeasures to overcome the threat.

8.1. Security Threat

Consider an end host that wants to act maliciously and creates its own location object with faked location information and uses this information in a subsequent SIP communication. In case of an emergency call the other communication partner, the Public Safety Answering Point (PSAP) operator, would use the information potentially without having a further possibility to verify the received location information. Emergency personnel would be sent to the indicated location noticing that there is no incident.

Hence, the PSAP operator, and the Location Recipient in general, would like to ensure that the provided location information is genuine, accurate and fresh to avoid taking wrong actions, such as dispatching emergency personnel to a wrong location.

There seems to be a need for preventing location forgery, replay and substitution attacks, which are all forms of sending a location which is deliberately not that of the end host. As shown below, various forms of countermeasures are possible to mitigate these attacks. Although some aspects are within the scope of the Geopriv-L7 Location Configuration Protocol (LCP), which is between a LIS and an Target, some aspects refer to other protocols, as shown in Figure 4. For example, in an emergency call, the PSAP (as a Location Recipient) wishes to verify that the location is indeed that of the calling party. Further, the Geopriv-L7 LCP is not the only protocol that could be used by an end host to acquire its location. Therefore, the topic of signatures on the location information was deemed out of scope. The subsequent discussion about countermeaures aims to capture the state of the discussions and illustrates the complexity in the overall design.

8.2. Discussion about Countermeasures

The goal of the above-described mechanism is to prevent prank calls and, in case of emergency services, unnecessary first-responder dispatch. As such, it is a mechanism to reduce the vulnerability of denial of service attacks. The benefit of a digital signature created by the LIS and covering the location information (plus some other fields) is to treat a missing or invalid signature as suspect during the call. The call would be treated differently in the sense that more questions might be asked (if an interaction with a human

person is possible). In case of emergency services, the call might get ranked differently if certain criteria are not fulfilled and if the PSAP operator is confronted with a massive amount of calls without the possiblity to respond to all of them.

8.2.1. Signed Location Information

One of the proposed countermeasures is to sign location information by the LIS before it is sent to the end host whereby the signed location information is verified by the final Location Recipient rather than the Target. This prevents the Target from tampering with the received location information since the digital signature would become invalid. The Location Recipient would be able to verify the source of the location information. Since the number of nodes that may play the role of a Location Recipient is large it is difficult to utilize a pre-shared secret key based infrastructure. Hence, a public key based infrastructure is required but authorization still remains challenging.

This solution approach is challenging when a PIDF-LO [16] has to be signed (instead of location information only) since the PIDF-LO contains more than just location information, such as "entity" attribute of the 'presence' element, and usage-rules (e.g., 'retransmission-allowed', 'retention-expires', 'ruleset-reference', 'note-well').

The value for the "entity" attribute of the 'presence' element is, in many cases, not known to the L2/L3 provider. If the LIS signs some layer-2/layer-3 (e.g., PPP/RADIUS/NAI) identity as entity URI, it will unlikely be the SIP URI.

To prevent adversaries from reusing an eavesdropped signed location object it is necessary to include additional information when generating the digital signature. For example, a timestamp and a validity field are useful to prevent certain replay attacks. Furthermore, the "entity" attribute may be included in the digital signature of a PIDF-LO with the following semantic: When using the signed location object (e.g., in SIP or another higher layer protocol), the Target needs to authenticate to the Location Recipient using the same identity carried in the "entity" attribute of the 'presence' element of the signed PIDF-LO. Using SIP, for example, a SIP proxy server could assert the entity URI corresponding to the Target using the SIP identity mechanism.

Including the layer 7 identity into the "entity" attribute of the 'presence' element poses a privacy problem since the access network provider can now see an identity that is in use. Hence, the LIS and possibly unauthorized listeners (if there's no privacy protection)

find out where the L7 entity is located, rather than just the location information.

With regard to the ability for an adversary to replay an eavesdropped a signed location object, the following two approaches need to be considered:

1. A signed PIDF-LO with the L7 identity included, conveyed without confidentiality protection from the Target to the Location Recipient, and

2. A signed PIDF-LO, without the L7 identity, conveyed with confidentiality protection from the Target to the Location Recipient.

Note that in both cases confidentiality protection for the communication between the LIS and the Target is provided. (2) has the same security properties as (1) in terms of the ability of somebody else to steal and re-use the PIDF-LO ("location theft") (assuming the Location Recipient and the Target are honest).

An adversary might, for example, want to perform a replay attack by eavesdropping the signed location object. If the LIS includes additional attributes, such as a timestamp and the validity time, the vulnerability can be reduced although not entirely prevented. The reason for an adversary to still be able to replay the location information is that there is no verifiable identifier is associated with the signed location information. For example, the LIS might include the IP address of the end host to the signed location object. Spoofing the IP address is, however, relatively easy. Moreover, the IP address that is used to associate the location information cannot be verified by the LR since the IP address can be modified legitimately (e.g., NAT reasons) or might not be seen due to tunneling techniques (e.g., VPN, Mobile IP).

Ideally, an "identifier" with the property of being non-spoofable by an adversary and verifiable by the LR when it receives a signed location object, which will ensure that the submitted location information is actually sent by the claimed end host and not replayed. One such verifiable identifier is a public key, the serial number of a certificate, a hash of a public key (in the sense of Purpose-Built-Keys or Cryptographically-Generated-Addresses) or the value of a hash chain. We call this identifier, key identifier or keyID for short.

In more details, the end host provides this identifier to the LIS and it is signed together with location information. The following steps are executed:

1. The end host interacts with the LIS to obtain its location information. The communication is secured using Transport Layer Security. This request carries the keyID. In this example, we use a keyID that represents the hash of a public key. The LIS ties the received keyID to the location object and signs it.

2. The LIS returns the signed location object that includes the keyID to the requesting end host.

3. Whenever the end host wants to distribute its location information to a LR, it attaches location information to a SIP message as described in [14]. The end host computes a digital signature over the SIP header fields and signed location object (as, for example, envisioned by SIP Identity [17]) with the private key that corresponds to the hashed public key found in the signed location object.

4. This message is sent to the LR.

5. The LR receives the message and it performs the following steps:

   * It retrieves the public key.

   * It computes the hash over the public key and compares it with the value in the key identifier included in the signed location object.

   * It verifies the digital signature and thereby ensures that the end host is indeed in possession of the private key corresponding to the obtained public key.

   * It verifies the digital signature protecting the location information and checks whether it was signed by a trusted access provider.

Even if an adversary eveasdrops the communication between the end host and the LR it cannot successfully replay a signed location object since it does not know the private key corresponding to the hashed public key found in the signed location information. The achieved security protection might even be stronger in context of CGAs.

8.2.2. Authenticated Calls

In many cases, authenticated calls, i.e., verifying the callers identity, are at least as useful as location signing since it establishes accountability for later prosecution.

If most of the legitimate calls are authenticated in some way, then it is possible, under attack conditions only, to give "dubious" calls lower priority or to have them go through some sort of turing test. As an example, PSAP operators do not want to reject emergency calls regardless of how they look like, but if the alternative is wasting 90% of the resources on bogus calls (and thus leaving many legitimate callers stranded) and not handling the unlucky unauthenticated, the expected outcome is better if you can separate. This is the standard "triage" model used in emergency medicine.

If somebody places a signed (known-third-party VSP-authenticated) call, there is at least the possibility of catching a malicious caller and the number of such calls is limited. Thus, there are only legitimate calls left

- that use end system location determination (e.g., GPS, manual configuration);
- that have no (known) VSP;
- that are not signed in some other way

In general, it is necessary to separate authentication from charging. There is no reason for tying authentication, authorization and charging together for this particular context. For example, certificates can be used, for example, for emergency service without being subscribed to either a VSP or ISP.

8.2.3. Location-by-Reference

The concept of location-by-reference was described in Section 7. The properties of location signing are very similar (if not equal) to the properties of the location-by-reference concept when the Location Recipient only authenticates the LIS (but not vice-versa). Both mechanisms allow the Location Recipient to authenticate the LIS (and potentially the access network provider).

There are also a few drawbacks with the location signing and the location-by-reference concept:

- Location signing has very limited utility if the number of signing parties is very large
- Location signing has very limited utility for commercial transactions. Commercial entities do not care whether a customer lies about their location, as long as they can make you pay for the service you asked for.

Authenticated calls also have their disadvantage since they require end-host or end-user certificates, which creates a deployment burden, unless mechanisms similar to SIP Identity [18] are used. Furthermore, authenticated calls do not prevent attacks where the location information was obtained unsecured from a LIS and an adversary in the access network was able to tamper with the in-flight location information.

9. Requirements

The following requirements and assumptions have been identified:

Requirement L7-1: Identifier Choice

The LIS MUST be presented with a unique identifier of its own addressing realm associated in some way with the physical location of the end host.

An identifier is only appropriate if it is from the same realm as the one for which the location information service maintains identifier to location mapping.

Requirement L7-2: Mobility Support

The GEOPRIV Layer 7 Location Configuration Protocol MUST support a broad range of mobility from devices that can only move between reboots, to devices that can change attachment points with the impact that their IP address is changed, to devices that do not change their IP address while roaming, to devices that continuously move by being attached to the same network attachment point.

Requirement L7-3: Layer 7 and Layer 2/3 Provider Relationship

The design of the GEOPRIV Layer 7 Location Configuration Protocol MUST NOT assume a business or trust relationship between the provider of application layer (e.g., SIP, XMPP, H.323) provider and the access network provider operating the LIS.

Requirement L7-4: Layer 2 and Layer 3 Provider Relationship

The design of the GEOPRIV Layer 7 Location Configuration Protocol MUST assume that there is a trust and business relationship between the L2 and the L3 provider. The L3 provider operates the LIS and needs to obtain location information from the L2 provider since this one is closest to the end host. If the L2 and L3 provider for the same host are different entities, they cooperate for the purposes needed to determine end system locations.

Requirement L7-5: Legacy Device Considerations

The design of the GEOPRIV Layer 7 Location Configuration Protocol MUST consider legacy residential NAT devices and NTEs in an DSL environment that cannot be upgraded to support additional protocols, for example to pass additional information through DHCP.

Requirement L7-6: VPN Awareness

The design of the GEOPRIV Layer 7 Location Configuration Protocol MUST assume that at least one end of a VPN is aware of the VPN functionality. In an enterprise scenario, the enterprise side will provide the LIS used by the client and can thereby detect whether the LIS request was initiated through a VPN tunnel.

Requirement L7-7: Network Access Authentication

The design of the GEOPRIV Layer 7 Location Configuration Protocol MUST NOT assume prior network access authentication.

Requirement L7-8: Network Topology Unawareness

The design of the GEOPRIV Layer 7 Location Configuration Protocol MUST NOT assume end systems being aware of the access network topology. End systems are, however, able to determine their public IP address(es) via mechanisms such as STUN [4] or NSIS NATFW NSLP [19] .

10. Security Considerations

10.1. Capabilities of the Adversary

As common elsewhere, several kinds of attackers can be distinguished. As always, Alice is the "good guy" and Trudy the attacker. Attackers can be:

- off-path, i.e., it cannot see packets between Alice and the LIS;
- on-path, i.e., can see such packets.

On-path attackers may be:

- passive, i.e., can only observe;
- semi-active, i.e., can inject packets with a bogus IP address, but cannot prevent the delivery of packets from the end system or modify these packets;
- active, i.e., can inject and modify packets at will.

10.2. Threats

When the reference to location information is communicated to the Location Recipient then on-path adversaries can eavesdrop the signaling communication together with the reference. Furthermore, the end-to-end communication might involve SIP proxies and they may not be trustworthy. Hence, they can eavesdrop the reference and misuse it (by resolving it).

Untrusted proxies that are involved in the communication lead to a requirement for the Target to selectively grant access to already known and trusted Location Recipients.

The following list presents threats specific to location information handling:

Place-Shifting (PS):

Trudy pretends to be at an arbitrary location.

Time-Shifting (TS):

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

Location-Theft (LT):

Trudy observes Alice's location and replays it as her own location object.

Location-Identity-Theft (LIT):

Trudy observes Alice's location and her identity (e.g., presence identity) and replays it.

Location-Swapping (LS):

Trudy' and Trudy'', located at different locations, can collude and swap location objects and pretend to be in each other's location.

Table 1 shows the different threats and the applicability of proposed countermeasures.

	Asserted Location	Timestamp	Encrypted Location	Authenticated Call	Location by Reference
PS	X	-	-	Track Offender	X
TS	-	X	-	Track Offender	Limits Impact
LT	-	-	X	Track Offender	-
LIT	-	-	X	-	-
LS	-	Limits Impact	-	Track Offender	-

Table 1

Legend:

-: Functionality not necessary to accomplish the desired functionality.

X: Functionality needed to prevent threat.

10.3. Requirements

The following requirements are placed on the location-by-value approach:

- No conclusion was reached whether a PIDF-LO or just location information has to be signed.
- No conclusion was reached whether location information should be signed.
- No conclusion was reached what could be signed.

The following requirements are placed on the location-by-reference approach:

- The reference MUST be valid for a limited amount of time.
- The reference MUST be hard to guess, i.e., it MUST contain a cryptographically random component.
- The reference MUST NOT contain any information that identifies the user, device or address of record
- The Location Recipient MUST be able to resolve the reference more than once (i.e., there is no implicit limit on the number of dereferencing actions).
- Possessing a reference to location information allows a Location Recipient to repeately obtain the latest information about the Target with the same granularity.
- The Target MUST be able to resolve the reference itself.

11. IANA Considerations

This document does not require actions by IANA.

12. Contributors

This contribution is a joint effort of the GEOPRIV Layer 7 Location Configuration Requirements Design Team of the Geopriv WG. The contributors include Henning Schulzrinne, Barbara Stark, Marc Linsner, Andrew Newton, James Winterbottom, Martin Thomson, Rohan Mahy, Brian Rosen, Jon Peterson and Hannes Tschofenig.

The design team members can be reached at:

Marc Linsner: mlinsner@cisco.com

Rohan Mahy: rohan@ekabal.com

Andrew Newton: andy@hxr.us

Jon Peterson: jon.peterson@neustar.biz

Brian Rosen: br@brianrosen.net

Henning Schulzrinne: hgs@cs.columbia.edu

Barbara Stark: Barbara.Stark@bellsouth.com

Martin Thomson: Martin.Thomson@andrew.com

Hannes Tschofenig: Hannes.Tschofenig@siemens.com

James Winterbottom: James.Winterbottom@andrew.com

The authors would like to thank Barbara Stark for her 'Virtual Private Network (VPN) Considerations' text proposal.

13. Acknowledgements

We would like to thanks the IETF GEOPRIV working group chairs, Andy Newton, Allison Mankin and Randall Gellens, for creating this design team. Furthermore, we would like thank Andy Newton for his support during the design team mailing list, the Jabber chat conference and the phone conference discussions. Finally, we would like to thank Murugaraj Shanmugam for his draft review.

14. References

14.1. Normative References

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

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

[3] Schulzrinne, H. and R. Marshall, "Requirements for Emergency Context Resolution with Internet Technologies", draft-ietf-ecrit-requirements-12 (work in progress), August 2006.

14.2. Informative References

[4] Rosenberg, J., Weinberger, J., Huitema, C., and R. Mahy, "STUN - Simple Traversal of User Datagram Protocol (UDP) Through Network Address Translators (NATs)", RFC 3489, March 2003.

[5] Aboba, B., "Link-local Multicast Name Resolution (LLMNR)", draft-ietf-dnsext-mdns-47 (work in progress), August 2006.

[6] Cheshire, S. and M. Krochmal, "Multicast DNS", draft-cheshire-dnsext-multicastdns-06 (work in progress), August 2006.

[7] Moskowitz, R., "Host Identity Protocol", draft-ietf-hip-base-06 (work in progress), June 2006.

[8] Aura, T., "Cryptographically Generated Addresses (CGA)", RFC 3972, March 2005.

[9] Aboba, B., Beadles, M., Arkko, J., and P. Eronen, "The Network Access Identifier", RFC 4282, December 2005.

[10] Rigney, C., Willens, S., Rubens, A., and W. Simpson, "Remote Authentication Dial In User Service (RADIUS)", RFC 2865, June 2000.

[11] Calhoun, P., Loughney, J., Guttman, E., Zorn, G., and J. Arkko, "Diameter Base Protocol", RFC 3588, September 2003.

[12] Lemon, T. and B. Sommerfeld, "Node-specific Client Identifiers for Dynamic Host Configuration Protocol Version Four (DHCPv4)", RFC 4361, February 2006.

[13] Mahy, R., "A Document Format for Filtering and Reporting Location Notications in the Presence Information Document Format Location Object (PIDF-LO)", draft-ietf-geopriv-loc-filters-00 (work in progress), March 2006.

[14] Polk, J. and B. Rosen, "Session Initiation Protocol Location Conveyance", draft-ietf-sip-location-conveyance-04 (work in progress), August 2006.

[15] Hardie, T., "LoST: A Location-to-Service Translation Protocol", draft-ietf-ecrit-lost-01 (work in progress), September 2006.

[16] Peterson, J., "A Presence-based GEOPRIV Location Object Format", RFC 4119, December 2005.

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

[18] Peterson, J. and C. Jennings, "Enhancements for Authenticated Identity Management in the Session Initiation Protocol (SIP)", draft-ietf-sip-identity-06 (work in progress), October 2005.

[19] Stiemerling, M., "NAT/Firewall NSIS Signaling Layer Protocol (NSLP)", draft-ietf-nsis-nslp-natfw-12 (work in progress), June 2006.

[20] Schulzrinne, H., "Common Policy: A Document Format for Expressing Privacy Preferences", draft-ietf-geopriv-common-policy-11 (work in progress), August 2006.

[21] Schulzrinne, H., "A Document Format for Expressing Privacy Preferences for Location Information", draft-ietf-geopriv-policy-08 (work in progress), February 2006.

Authors' Addresses

Hannes Tschofenig
Siemens Networks GmbH & Co KG
Otto-Hahn-Ring 6
Munich, Bavaria 81739
Germany

Phone: +49 89 636 40390
Email: Hannes.Tschofenig@siemens.com
URI: http://www.tschofenig.com

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

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

Full Copyright Statement

Intellectual Property

Acknowledgment

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).