TOC

GEOPRIV	R. Barnes
Internet-Draft	M. Lepinski
Updates: 3693, 3694	BBN Technologies
(if approved)	A. Cooper
Intended status: BCP	J. Morris
Expires: September 10, 2009	Center for Democracy &
	Technology
	H. Tschofenig
	Nokia Siemens Networks
	H. Schulzrinne
	Columbia University
	March 09, 2009

An Architecture for Location and Location Privacy in Internet Applications
draft-barnes-geopriv-lo-sec-05

Status of this Memo

Copyright Notice

Abstract

Location-based services (such as navigation applications, emergency services, management of equipment in the field) need geographic location information about Internet hosts, their users, and other related entities. These applications need to securely gather and transfer location information for location services, and at the same time protect the privacy of the individuals involved. This document describes an architecture for privacy-preserving location-based services in the Internet, focusing on authorization, security, and privacy requirements for the data formats and protocols used by these services.

Table of Contents

1. Introduction

TOC

Location-based services (applications that require information about the geographic location of an individual or device) are becoming increasingly common on the Internet. Navigation and direction services, emergency services, friend finders, management of equipment in the field and many other applications require geographic location information about Internet hosts, their users, and other related entities. As the accuracy of location information improves and the expense of calculating and obtaining it declines, the distribution and use of location information in Internet-based services will likely become increasingly pervasive. Ensuring that location information is transmitted and accessed in a secure and privacy-protective way is essential to the future success of these services, as well as the minimization of the privacy harms that could flow from their wide deployment and use.

Standards for communicating location information over the Internet have an important role to play in providing a technical basis for privacy and security protection. This document describes a standardized privacy- and security-focused architecture for location-based services in the Internet: the Geopriv architecture. The central component of the Geopriv architecture is the location object, which is used to convey both location information about an individual or device and user-specified privacy rules governing that location information. As location information moves through its life cycle -- positioning, distribution, and finally receipt and use by its ultimate recipient(s) -- Geopriv provides mechanisms to guarantee the integrity and confidentiality of location objects and to ensure that location information is only transmitted in compliance with the user's privacy rules.

The goals of this document are two-fold: First, the architecture described revises and expands on the basic Geopriv Requirements [2] (Cuellar, J., Morris, J., Mulligan, D., Peterson, J., and J. Polk, “Geopriv Requirements,” February 2004.)[3] (Danley, M., Mulligan, D., Morris, J., and J. Peterson, “Threat Analysis of the Geopriv Protocol,” February 2004.), in order to clarify how these privacy concerns and the Geopriv architecture apply to use cases that have arisen since the publication of those documents. Second, this document should provide a general introduction to Geopriv and Internet location-based services, and be useful as a good first document for readers new to Geopriv.

1.1. Binding Rules to Data

TOC

A central feature of the Geopriv architecture is that location information is always bound to privacy rules, in order to ensure that entities that receive location are informed of how to they may use it. By creating a structure to convey the user's preferences along with location information, the likelihood that those preferences will be honored necessarily increases. In particular, no recipient of the location information can disavow knowledge of users' preferences for how their location may be used. The binding of privacy rules to location information can convey users' desire for and expectations of privacy, which in turn helps to bolster social and legal systems' protection of those expectations.

Binding of usage rules to sensitive information is a common way of protecting information. Several emerging schemes for expressing copyright information provide for rules to be transmitted together with copyrighted works. The Creative Commons model is the most prominent example, allowing an owner of a work to set four types of rules ("Attribution," "Noncommercial," "No Derivative Works" and "ShareAlike") governing the subsequent use of the work. After the author sets these rules, the rules are conveyed together with the work itself, so that every recipient is aware of the copyright terms of the work.

Classification systems for controlling sensitive documents within an organization are another example. In these systems, when a document is created, it is marked with a classification such as "SECRET" or "PROPRIETARY." Each recipient of the document knows from this marking that the document should only be shared with other people who are authorized to access documents with that marking. Classification markings can also convey other sorts of rules, such as a specification for how long the marking is valid (a declassification date). For example, the United States Department of Defense guidelines for classification [4] (U.S. Department of Defense, “National Industrial Security Program Operating Manual,” January 1995.) allow the creator of a document to mark it with a classification level that restricts access (e.g., "SECRET") and an indication of when the document should be declassified or downgraded to a lower classification (e.g., "DECLASSIFY ON December 31, 2011" or "DOWNGRADE TO CONFIDENTIAL ON December 31, 2011").

1.2. Location-Specific Privacy Risks

TOC

While location-based services raise some privacy concerns that are common to all forms of personal information, many of them are heightened and others are uniquely applicable in the context of location information.

Location information is frequently generated on or by mobile devices. Because individuals often carry their mobile devices with them, location information may be used to form a comprehensive record of an individual's movements and activities. While other kinds of data could arguably be considered more sensitive than location information in certain contexts -- an individual's medical records or bank statements, for instance -- these kinds of data provide mere snapshots of an individual's activities at discrete moments in time, or within discrete aspects of their lives. Location information, on the other hand, may be collected everywhere and at any time, often without explicit user interaction, and it may potentially describe both what a person is doing and where he or she is doing it. The fact that an individual's mobile device location is obtained when he is at the bank can reveal that he was at the bank, when he was there, and which branch he uses. Location-based services may allow for amassing such data points about an individual's every movement, potentially spurring the creation of richly detailed profiles of individual behavior.

The availability of location information may also allow an individual's whereabouts to unwittingly become more public than desired, with potentially serious consequences. Location information may reveal the fact that an individual was in a particular medical clinic or government building, for example, implying potentially very sensitive information about the individual that was not meant to be shared. The ubiquity of location information may also increase the risks of stalking and domestic violence if perpetrators are able to use (or abuse) location-based services to gain access to location information about their victims. Location information additionally raises significant child safety concerns as more and more children access location-aware devices.

Finally, location information is and will continue to be of particular interest to governments and law enforcers around the world. The existence of detailed records of individuals' movements should not automatically facilitate the ability for governments to track their citizens, but in some jurisdictions, laws dictating what government agents must do to obtain location data are either non-existent or out-of-date.

1.3. Privacy Paradigms

TOC

Traditionally, the extent to which data about individuals enjoys privacy protections on the Internet has largely been decided by the recipients of the data. Internet users may or may not be aware of the privacy practices of the entities with whom they share data. Even if they are aware, they have generally been limited to making a binary choice between sharing data with a particular entity or not sharing it. Internet users have not historically been granted the opportunity to express their own privacy preferences to the recipients of their data and to have those preferences honored.

This paradigm is problematic because the interests of data recipients are often not aligned with the interests of data subjects. While both parties may agree that data should be collected, used, disclosed and retained as necessary to deliver a particular service to the data subject, they may not agree about how the data should otherwised be used. For example, an Internet user may gladly provide his email address on a Web site to receive a newsletter, but he may not want the Web site to share his email address with marketers, whereas the Web site may profit from such sharing. Neither providing the address for both purposes nor deciding not to provide it is an optimal option from the Internet user's perspective.

The Geopriv model departs from this paradigm for privacy protection. As explained above, location information can be uniquely sensitive. And as siloed location-based services emerge and proliferate, they increasingly require standardized protocols for communicating location information between services and entities. Recognizing both of these dynamics, Geopriv gives data subjects the ability to express their choices with respect to their own location information, rather than allowing the recipients of the information to define how it will be used. The combination of heightened privacy risk and the need for standardization compelled the Geopriv designers to shift away from the prevailing Internet privacy model, instead empowering users to express their privacy preferences about the use of their location information.

Geopriv does not, by itself, provide technical means through which it can be guaranteed that users' location privacy rules will be honored by recipients. The privacy protections in the Geopriv architecture are largely provided by virtue of the fact that recipients of location (Location Servers and Location Recipients in the below discussion) are informed of relevant privacy rules, and must only use location in accordance with those rules. The distributed nature of the architecture inherently limits the degree to which compliance privacy controls -- the fact that an entity has not used location in an unauthorized way -- can be guaranteed and verified by technical means. (Some security mechanisms can address this problem to a limited extent; see Section 4 (Security Considerations).)

By binding privacy rules to location information, however, Geopriv provides valuable information about users' privacy preferences, so that non-technical forces such as legal contracts, governmental consumer protection authorities, and marketplace feedback can better enforce those privacy preferences. If a commercial recipient of location information, for example, violates the location rules bound to the information, the recipient can in a growing number of countries be charged with violating consumer or data protection laws. In the absence of a binding of rules with location information, consumer protection authorities would be less able to protect consumers whose location information has been abused.

2. Overview of the Architecture

TOC

This section provides an overview of the Geopriv architecture for the secure and private distribution of location information on the Internet. We describe the three phases of the "location life cycle" -- positioning, distribution and receipt -- and discuss how the components of the architecture fit within each phase. The next section provides additional detail about how each phase can be achieved in a private and secure manner.

The risks discussed in the previous section all arise from unauthorized disclosure or usage of location information. Thus, the Geopriv architecture has two fundamental privacy goals:

1. Ensure that location information is distributed only to authorized entities, and

2. Provide information to those entities about how they are authorized to use the location information.

If these two goals are met, all parties that receive location information will also receive directives about how they can use that information. Privacy-preserving entities will only engage in authorized uses, and entities that violate privacy will do so knowingly, since they have been informed of what is authorized (and thus, implicitly, of what is not).

Privacy rules and their distribution are thus the central technical components of the privacy system, since they inform location recipients about how they are authorized to use that information. The two goals in the preceding paragraph are enabled by two classes of rules:

1. Access control rules: Rules that describe which entities may receive location information and in what form

2. Usage rules: Rules that describe what uses of location information are authorized

Within this framework for privacy, security mechanisms provide support for the application of privacy rules. For example, authentication mechanisms validate the identities of entities requesting location (so that authorization and access-control policies can be applied), and confidentiality mechanisms protect location information en route between privacy-preserving entities. Security mechanisms can also provide assurances that are outside the purview of privacy by, for example, assuring location recipients that location information has been faithfully transmitted to them by its creator.

2.1. Basic Geopriv Scenario

TOC

As location information is transmitted among Internet hosts, it goes through a "location life-cycle:" first, the location is computed based on some external information (positioning), then it is transmitted from one host to another (distribution) until finally it is used by a recipient (receipt).

For example, suppose Alice learns of her location from a wireless location service and wishes to share it privately with her friends by way of a presence service. Alice clearly needs to provide the presence server with her location and a list of friends to whom the server can grant access to the location. To enable Alice's friends to preserve her privacy, they need to be provided with privacy rules. Alice may tell some of her friends the rules directly, or she can have the presence server provide the rules to her friends when it provides them with her location. In this way, every friend who receives Alice's location is authorized by Alice to receive it, and every friend who receives it knows the rules. Good friends will obey the rules. If a bad friend breaks them and Alice finds out, the bad friend cannot claim that he was unaware of the rules.

Some of Alice's friends will be interested in using Alice's location only for their own purposes (to meet up with her or plot her location over time, for example). The usage rules that they receive direct them as to what they can or cannot do (for example, Alice might not want them keeping her location for more than, say, two weeks).

Consider one friend, Bob, who wants to send Alice's location to some of his friends. Bob needs not only usage rules for himself, but also access control rules that describe who he can send information to and rules to give to the recipients. If the rules he received from the presence server authorize him to give Alice's location to others, he may do so; otherwise, he will require additional rules from Alice before he is authorized to distribute her location. If recipients who receive Alice's location from Bob want to distribute the location on further, they must go through the same process as Bob.

The whole example is illustrated in the following figure:

```
     +----------+
     | Wireless |
     | Location |
     | Service  |
     +----------+
        |
        |
       Location
        |
        |
        |    +------------More-Rules---------------------->+-----+
        |    |                                       +---->| Bob |--> ...
        |    |                                       |     +-----+
        v    |                                       |
     +-------+            +----------+               |
     |       |--Location->| Presence |--Location-+   |     +----------+
     | Alice |            | Service  |           |---+---->| Friend-1 |
     |       |---Rules--->|          |---Rules---+   |     +----------+
     +-------+            +----------+               |
                                                                |
                                                     |     +----------+
                                                     +---->| Friend-2 |
                                                           +----------+
```

Figure 1: Basic Geopriv Scenario

2.2. Roles and Data Formats

TOC

The above example illustrates the five basic roles in the Geopriv architecture:

Target: An individual or other entity whose location is conveyed in the Geopriv architecture. The Target is the entity whose privacy Geopriv seeks to protect. Alice is the Target in the figure above.

Rule Maker (RM): An individual or entity that creates rules governing access to location information for a Target. In some cases the Rule Maker and the Target will be the same individual or entity (as is the case with Alice), and in other cases they will be separate. For example, a parent may serve as the Rule Maker when the Target is his child, or a corporate security officer may be the Rule Maker for devices owned by the corporation but used by employees. The Rule Maker, however, is not necessarily the owner of a Target device. For example, a corporation may provide a device to an employee but permit the employee to serve as the Rule Maker and set her own privacy rules.

Location Generator (LG): The entity that initially determines or gathers the location of the Target. Location Generators may be any sort of software or hardware used to obtain the Target's position (examples include GPS chips and cellular networks). A Target may even be its own Location Generator; devices capable of unassisted satellite-based positioning and devices that accept manually entered location information are two examples. The wireless location service is the Location Generator in the figure above.

Location Server (LS): An entity that receives both location information and rules, and applies the rules to the location information to determine what other entities, if any, can receive location information. The first LS in the Geopriv process receives location information from Location Generators and rules from Rule Makers, and then applies the rules to the location information. Location Servers may not necessarily be "servers" in the colloquial sense of hosts in remote data centers servicing requests. Rather, a Location Server can be any software or hardware component that receives and distributes location information. Examples include a server in an access network, a presence server, or a Web browser or other software running on a Target's device. The above example includes four Location Servers: the wireless location service, Alice, the presence service and Bob.

Location Recipient (LR): The ultimate end point entity to which location information is distributed. A Location Recipient may ask for location explicitly (by sending a query to a Location Server), or it may receive location asynchronously. Location

> Recipients do not distribute location information to any other Geopriv entities. Friend-1 an Friend-2 are Location Recipients in the figure above.

In general, these entities may or may not be physically separate from each other.
Within this architecture, entities acting in Geopriv roles communicate using three types of protocols, which carry Location Objects and Privacy Rules in well-defined data formats:

> **Privacy Rule:** A directive that regulates an entity's activities with respect to location information, including the collection, use, disclosure, and retention of the location information. Privacy Rules describe how location information may be used by an entity, the level of detail with which location information may be described to an entity, and the conditions under which location information may be disclosed to an entity.
>
> **Location Object (LO):** An object used to convey location information together with Privacy Rules. Geopriv supports both geodetic location data (latitude/longitude/altitude/etc.) and civic location data (street/city/state/etc.). Either or both types of location information may be present in a single LO. In the positioning phase, a LO may contain location information without Privacy Rules (which are passed from one entity to another during the distribution phase).
>
> **Positioning Protocol:** A protocol used by a Location Generator and a source external to the LG (the Target, for example) to exchange information necessary to determine the Target's location. Many Positioning Protocols also carry a Location Object representing the location derived from this information.
>
> **Conveyance Protocol:** A protocol used by a Location Server to send a Location Object to a Location Recipient or another Location Server.
>
> **Rules Protocol:** A protocol used by a Rule Maker to provide Privacy Rules to a Location Server.

The whole example, using Geopriv roles, objects and protocols, is illustrated in the following figure:

```
     +----+
     | LG |
     +----+
       |
       |
     Positioning (LO w/o Rules)
     Protocol
       |
       |    +------------Rules Protocol------------------->+----+
       |    |            (Privacy Rules)           +---->| LS |--> ...
       |    |                                      |     +----+
       v    |                                      |
     +-------+                                     |
     |Target | Conv. Proto.   +----+ Conv. Proto.  |     +----+
     |  RM   |--------------->| LS |---------------+---->| LR |
     |  LS   | (LO w/ Rules)  +----+ (LO w/ Rules) |     +----+
     +-------+                                     |
                                                   |     +----+
                                                   +---->| LR |
                                                         +----+
```

Figure 2: Basic Geopriv Scenario

2.3. Relationships Between Geopriv Roles

TOC

Although in the above example there is only a single Location Generator and a single Rule Maker, in some cases a Location Server may receive Location Objects from multiple Location Generators or Rules from multiple Rule Makers. Likewise, a single Location Generator may publish location information to multiple Location Servers, and a single Location Recipient may receive Location Objects from multiple Location Servers.

The term "Target" may refer not only to an individual whose location is described by a LO, but also to that individual's device, since the device engages in protocol interactions, not the individual. For the remainder of this document, the term "Target" refers to the device. Geopriv can also be used to convey location information about a device that is not directly linked to a single individual, such as a package or product containing a location-capable sensor, or a device linked to multiple individuals.

Although a single individual may use multiple devices, the Geopriv protocols address one device per individual at a time; it is outside the scope of Geopriv to interpolate one individual's location information across multiple devices or to arbitrate between privacy preferences that differ across the individual's devices.

There is also typically a close binding between a Location Generator and the first Location Server in the distribution path. These two roles may be played by the same entity, for example, a positioning server that can also be queried for location. While these two roles can be played by different entities, the relationship between them needs to be closely prescribed in order to preserve privacy, since an LS is a privacy-aware entity and an LG may not be. The specific constraints on the relationship between an LG and an LS are described in Section 3.1.2 (Privacy Considerations).

3. The Location Life-Cycle

TOC

The previous section gave an example of how an individual's location can be distributed through the Internet. In general, the location life-cycle breaks down into three phases:

1. Positioning: A Location Generator determines the Target's location
2. Distribution: Location Servers send location from one Location Server to another (possibly several times)
3. Receipt: A Location Recipient receives the location and uses it.

Each of these phases involves a different set of Geopriv roles, and each has a different set of privacy implications. The Geopriv roles are mapped onto the location life-cycle in the figure below.

```
+----------+                +----------+
|          |                |  Rule    |+
|  Target  |                |  Maker(s)||
|          |                |          ||
+----------+                +----------+|
     ^|                      +----------+
     || Positioning              | Rules
     || Protocol                 | Protocol
     ||                          |
     |V                          V        LO (w/ Rules)
+----------+                +----------+   Conveyance   +----------+
|Location  | LO (w/o Rules) | Location |+   Protocol    |Location  |
|Generator |--------------->| Server(s)||-------+------>|Recipient |
|          |                |          ||               |          |
+----------+                +----------+|               +----------+
                            +----------+
<----------><------------------------------------------><------------>
 Positioning                Distribution                 Receipt
```

Figure 3: Location Life-Cycle

3.1. Positioning

TOC

Positioning is the process by which the physical location of the Target is computed, based on some observations about the Target's situation in the physical world. (This process goes by several other names, including Location Determination or Sighting.) The input to the positioning process is some information about the Target, and the outcome is that the Location Generator knows the location of the Target. Said differently, positioning is the process by which a Location Generator generates a Location Object from other information about a Target.

Given that they are situated at the beginning of the life-cycle of location information, positioning mechanisms and the protocols that support them play a central role in determining who has access to a Target's location information. At the end of the positioning process, the Location Generator (which may be the Target itself) knows the location of the Target. The LG, and possibly the Target, thus have the capability to distribute the Target's position, and the responsibility to do so in a privacy-preserving manner.

In this section, we give a brief taxonomy of current positioning systems, their requirements for protocol support, and the privacy and security requirements for these protocols.

3.1.1. Determination Mechanisms and Protocols

TOC

While the specific positioning mechanisms that can be applied for a given Target are strongly dependent on the physical situation and capabilities of the Target, these mechanisms generally fall into the three categories described in detail below:

1. Target-based
2. Network-based
3. Network-assisted

As suggested by the above names, a positioning scheme can rely on the Target, an Internet-accessible resource (not necessarily a network operator), or a combination of the two. For a given scheme, the nature of this reliance will dictate the protocol mechanisms needed to support it.

With Target-based positioning mechanisms, the Target is capable of determining its location by itself. This is the case for manually-entered location or for (unassisted) satellite-based positioning (using a Global Navigation Satellite System, or GNSS). In these cases, the Target itself is a Location Generator, and there are no protocols required to support positioning (since no information needs to be communicated).

In network-based positioning schemes, a third-party Location Generator (i.e., an Internet host other than the Target) has access to sufficient information about the Target, through out-of-band channels, to establish the position of the Target. In these cases, the Location Generator is the entity that has access to the out-of-band information used for positioning. The most common examples of this type of LG are entities that have a physical relationship to the Target (such as ISPs). In wired networks, wiremap-based location is a network-based technique; in wireless networks, timing and signal-strength based techniques that use measurements from base stations are considered to be network-based. Large-scale IP-to-geo databases (for example, those based on WHOIS data or latency measurements) are also considered to be network-based positioning mechanisms.

For network-based positioning as for Target-based, no protocols are strictly necessary to support positioning, since positioning information is collected outside of the location distribution system (at lower layers of the network stack, for example). This does not rule out the use of other Internet protocols (like SNMP) to collect inputs to the positioning process; rather, since these inputs can only be used by certain Location Generators to determine location, they are not controlled as private information. (However, when used in this way, they are Positioning Protocols.) Network-based positioning often provides location to protocols by which the network informs a Target device of its position (Location Configuration Protocols [5] (Tschofenig, H. and H. Schulzrinne, “GEOPRIV Layer 7 Location Configuration Protocol; Problem Statement and Requirements,” July 2009.)).

Network-assisted systems account for the greatest number and diversity of positioning schemes. In these systems, the work of positioning is divided between the Target and an external Location Generator via some communication (possibly over the Internet), typically in one of two ways:

1. The Target provides measurements to the LG

2. The LG provides assistance data to the Target

In this case, "measurements" are understood to be observations about the Target's environment, ranging from wireless signal strengths to the MAC address of a first-hop router. "Assistance" is the complement to measurement, namely the information that enables the computation of location-based on measurements. A set of wireless base station locations (or wireless calibration information) would be an assistance datum, as would be a table mapping routers to buildings in a corporate campus.

For example, wireless and wired networks can serve as the basis for network-assisted positioning. In several current 802.11 positioning systems, the Target sends measurements (e.g., MAC addresses and signal strengths) to a Location Generator, and the Location Generator returns a location to the client. In fixed networks, the Target can send its MAC address to the Location Generator, which can query the MAC-layer infrastructure to determine the switch and port to which that MAC address is connected, then query a wire map to determine the location at which the wire connected to that port terminates.

As an aside, the common phrase "assisted GPS" ("assisted GNSS" more broadly) actually encompasses techniques that transmit both measurements and assistance data. Systems in which the Target provides the assistance server with data such as pseudo-ranges are measurement-based, while those in which the assistance server provide ephemeris or alamanac data are assistance-based in the above terminology. (Those familiar with GNSS positioning will note that there are of course cases in which both of these interactions occur within a single location determination protocol, so the categories are not mutually exclusive.)

Naturally, the exchange of measurement or positioning data between the Target and the LG requires a protocol over which the information is carried. The structure of this protocol will depend on which of the two patterns a network-assisted scheme follows. Conversely, the structure of the protocol will determine which of the two parties (the Target, the LG, or both) is aware of the Target's location at the end of the protocol.

In summary, the positioning process can involve three Geopriv roles, and three protocols:

> **Location Generator:** An LG supports the positioning process, and if it has applicable Rules may act as a Location Server for Location Objects generated through the positioning process.
>
> **Target:** The Target can act as either as an LG (if it receives location as a result of positioning) or simply as a source of inputs to the positioning process.
>
> **Location Server:** At the end of the positioning process, either the Target or an external Location Generator is enabled to act as an LS, making it subject to privacy requirements, or the LG transmits the information to the first LS (which in this case must obtain Rules from a Rule Maker).
>
> **Positioning Protocol:** The protocol used by the Target and a Location Generator to exchange measurement and assistance data.
>
> **Rules Protocol:** The protocol used by the a Rule Maker to provide privacy rules to a Location Server or Location Generator.
>
> **Conveyance Protocol:** The protocol used by a Location Generator or Location Server to send a Location Object to a Location Recipient or another Location Server.

3.1.2. Privacy Considerations

TOC

At the conclusion of the positioning process, either the Target or an external LG (or both) has access to the location of the Target, and those entities (if they have applicable Rules from the Rule Maker) can act as Location Servers, or they will transmit the location to the first LS in the Geopriv process.

If either entity chooses to act as an LS by distributing the Target's location, then it must only do so in authorized ways. This requirement means that an LS must be provided with Privacy Rules for the Target's location, which dictate where the location may be sent, and what Privacy Rules should be sent with it. If no Rules are available to an LS, then it must obey a set of privacy-preserving default rules (namely, the LCP policy described below).

The simplest case is when the Target acts as the initiating LS, which happens with a Target-based positioning scheme or a network-assisted scheme that results in the Target knowing its location. In this case, Rules are always available, since the Target (or a user of a Target device) can act as a Rule Maker. This requirement implies that software that implements LS functions on a host (i.e., software that distributes location) MUST provide interfaces for the user to specify both (1) to whom the software may send location information and (2) what Rules should be transmitted along with that location.

When an entity other than the Target acts as an initiating LS, then it must follow appropriate Rules to protect the privacy of location that it distributes. If Rules have been provided to it by an authorized Rule Maker -- either as part of the Positioning Protocol or through another channel, e.g., a Rules Protocol -- then the LG (acting as an initial LS) MUST transmit location only as allowed by these Rules. (Note that as for other Location Servers, the decision as to which Rule Makers are authorized is a matter of local policy.) Where possible, Positioning Protocols SHOULD enable the Target to convey Rules to the LG. If an LG supports Positioning Protocols that do not convey Rules and the LG is to serve as an LS, it MUST be able to receive Rules via a Rules Protocol.

If an LS (especially an initiating LS) is not provisioned with Privacy Rules that authorize transmission of location information for a given Target, then it MUST transmit location only to the Target, and consider all other recipients unauthorized. This default rule is known as the "LCP Policy", since it underlies the privacy aspects of Location Configuration Protocols (LCPs) [5] (Tschofenig, H. and H. Schulzrinne, “GEOPRIV Layer 7 Location Configuration Protocol; Problem Statement and Requirements,” July 2009.).

As an aside: There are several Location Configuration Protocols that have been developed within the IETF, both using DHCP [6] (Polk, J., Schnizlein, J., and M. Linsner, “Dynamic Host Configuration Protocol Option for Coordinate-based Location Configuration Information,” July 2004.)[7] (Schulzrinne, H., “Dynamic Host Configuration Protocol (DHCPv4 and DHCPv6) Option for Civic Addresses Configuration Information,” November 2006.)[8] (Polk, J., “Dynamic Host Configuration Protocol (DHCP) IPv4 and IPv6 Option for a Location Uniform Resource Identifier (URI),” March 2010.) and using HTTP [9] (Barnes, M., Winterbottom, J., Thomson, M., and B. Stark, “HTTP Enabled Location Delivery (HELD),” August 2009.). Within the architecture described in this document, these protocols are Conveyance Protocols: the Location Information Server that provides location through an LCP is a Location Server that provides location to the Target, following the LCP policy.

In some deployment scenarios, positioning functions and distribution functions may need to be provided by separate entities. That is, the LG and LS roles may need to be separated, with the LG acting as a "dumb," non-privacy-aware positioning resource, and the LS providing the privacy logic necessary to support distribution (possibly with multiple LSs using the same LG). In order to allow the privacy-unaware LG to distribute location to these LSs while maintaining privacy, the relationship between the LG and set of LSs MUST be tightly constrained, effectively "hard-wired." That is, the LG MUST provide location only to a small fixed set of LSs, and each of these LSs MUST comply with the requirements of this section and those in Section 3.2 (Location Distribution).

3.1.3. Security Considerations

TOC

Manipulation of Positioning Protocols can expose location through two mechanisms: If a third party can guess measurements that a given Target would provide to the LG, and then use them to get the location of that Target, or if a third party can obtain assistance data that indicate the rough position of the client. To mitigate this risk, a Positioning Protocol SHOULD allow the LG to authenticate Positioning Protocol clients (i.e., the Target or other information sources), in the sense of verifying that measurements presented by a client are likely to be the actual physical values measured by that client (and likewise, that the requested assistance data are consistent with the client's actual rough position). These authentication mechanisms will necessarily rely on the nature of the positioning being done, and may not be technically feasible in all cases.

In any case, Positioning Protocols MUST provide confidentiality and integrity protections in order to prevent observation and modification of transmitted positioning data and Location Objects while en route between the positioning client and the LG.

If a Location Generator or a Target choose to act as an initiating Location Server, they inherit the security requirements for an LS, described in Section 3.2.4 (Security Considerations).

3.2. Location Distribution

TOC

When an entity receives location (from an LG or another LS) and redistributes it to other entities, it acts as a Location Server. Location Distribution is the process by which one or more Location Servers move location information from its source (a Location Generator) to its destination (a Location Recipient), in a privacy-preserving manner.

The role of a Location Server is thus two-fold: First, it must collect location information and Rules that control access to that information. Rules can be communicated within a Location Object, within a Conveyance Protocol that carries LOs, or through a separate Rules Protocol. Second, the Location Server must process requests for location and apply the Rules to these requests in order to determine whether it is authorized to fulfill them by returning location information.

A Location Server thus has at least two types of interactions with other hosts, namely receiving and sending Location Objects through a Conveyance Protocol. An LS may optionally implement a third interaction, allowing Rule Makers to provision it with Rules via a Rules Protocol. The distinction between these two cases is important in practice, because it determines whether the LS has a direct relationship with a Rule Maker: An LS that accepts Rules via a Rules Protocol (or via a direct interaction with the Rule Maker) is known as an "Authorized LS" (LSa), while an LS that acquires all its Rules through a Conveyance Protocol is known as an "Independent LS" (LSi). The location distribution process thus involves three roles and three protocols:

> **Location Servers:** Entities that perform the actual transmissions of location, in accordance with available Rules
>
> **Rule Makers:** Entities that set Rules that constrain how location information is disseminated
>
> **Location Recipients:** The ultimate destinations for location information
>
> **Conveyance Protocols:** Protocols used by an LS to transmit Location Objects
>
> **Rules Protocols:** Protocols used by an RM to supply Rules to an LS
>
> **LO Formats:** The structure of a LO, including the available location and Rules semantics.

3.2.1. Privacy Rules

TOC

Privacy Rules are the central mechanism in Geopriv for maintaining a Target's privacy, because they provide a recipient of a LO (an LS or LR) with information on how the LO may be used.
Throughout the Geopriv architecture, Privacy Rules are communicated in a rules language with a defined syntax and semantics (a Rules Format). For example, the Common Policy rules language has been defined [10] (Schulzrinne, H., Tschofenig, H., Morris, J., Cuellar, J., Polk, J., and J. Rosenberg, “Common Policy: A Document Format for Expressing Privacy Preferences,” February 2007.) to provide a framework for broad-based rule specifications. Geopriv Policy [11] (Schulzrinne, H., Tschofenig, H., Morris, J., Cuellar, J., and J. Polk, “Geolocation Policy: A Document Format for Expressing Privacy Preferences for Location Information,” January 2010.) defines a language for creating location-specific rules. XCAP [12] (Rosenberg, J., “The Extensible

Markup Language (XML) Configuration Access Protocol (XCAP),” May 2007.) can be used as a Rules Protocol to install rules in both of these formats.

Privacy Rules follow a default-deny pattern: an empty set of Rules implies that all requests for location should be denied (other than requests made by the Target itself), with each Rule added to the set granting a specific permission. Adding a Rule to a set can never reduce existing permissions; it can only augment them.

The following are examples of Privacy Rules governing location distribution:

- *Retransmit location when requested from example.com
- *Retransmit location when requested by a specific group of requesters
- *Retransmit only geodetic location
- *Retransmit location accurate within 100 meters
- *Retransmit location only to the first three recipients who request it
- *Retransmit location only before midnight on December 31, 2009

Location Servers enforce Privacy Rules in two ways: by denying requests for location, or by transforming the location information before retransmitting it. Some Rules will only be enforceable through denial. For example, if the entire Rule set for a particular Target consisted of the first Rule listed above, then a Location Server would be required to deny all requests for the Target's location from any recipients other than example.com. On the other hand, the second rule above could be enforced either by rejecting requests for civic location or by stripping out all civic location from a LO received from an LG before retransmitting the LO to any LR that requests it.

Location Servers may also receive Rules governing location retention, such as:

- *Retain location only for 24 hours
- *Retain location only until December 31, 2009

These Rules are simply directives about how long the Target's location information can be retained.

Privacy Rules can govern the behavior of both Location Servers and Location Recipients. Rules that direct Location Servers about how to treat a Target's location information are known as Local Rules. Local Rules are used internally by the Location Server to handle requests from Location Recipients. They are not distributed to Location Recipients.

Rules that travel inside LOs are known as Forwarded Rules. Forwarded Rules direct Location Servers and Location Recipients about how to handle the location information they receive. Because the Rules themselves may reveal potentially sensitive information about the Target, only the minimal subset of Forwarded Rules necessary to handle the LO is distributed.

An example can illustrate the interaction between Local Rules and Forwarded Rules. Suppose Alice provides the following Local Rules to a Location Server:

- *The LS may retain location only for 24 hours
- *The LS may retransmit Alice's precise location to Bob, who in turn is permitted to retain the location information for one month
- *The LS may retransmit Alice's city, state, and country to Steve, who in turn is permitted to retain the location information for one hour
- *The LS may retransmit Alice's country to a photo-sharing website, which in turn is permitted to retain the location information indefinitely and retransmit it to any requesters

When Steve asks for Alice's location, the Location Server (assuming the request is within the authorized 24-hour window for the LS to retain Alice's location) can transmit to Steve the limited location information (city, state, and country) along with Forwarded Rules instructing Steve to (a) not further retransmit Alice's location information, and (b) only retain the location information for one hour. By only sending these specifically applicable Forwarded Rules to Steve (as opposed to the full set of Local Rules), the LS is protecting Alice's privacy by not disclosing to Steve that (for example) Alice allows Bob to obtain more precise location information than Alice allows Steve to receive.

Geopriv is designed to be usable even by devices with constrained processing capabilities. To ensure that Forwarded Rules can be processed on constrained devices, LOs are required to carry only a limited set of Forwarded Rules, with an option to reference a more robust set of external Rules. The limited Rule set covers two privacy aspects: how long the Target's location may be retained ("Retention"), and whether or not the Target's location may be retransmitted ("Retransmission"). (The latter rule will never grant an LR the permission to retransmit, since by definition an LR is a final end point for a Target's location, but an LS that receives a LO may be granted this permission.) A LO may contain a pointer to more robust Rules, such as those shown in the set of six Rules at the beginning of this section.

TOC

3.2.2. Location References

The location distribution process occurs through a series of transmissions of Location Objects: transmissions of location "by value." Location "by value" can be expressed in terms of geodetic location data (latitude/longitude/altitude/etc.) and civic location data (street/city/state/etc.).

Location can also be distributed "by reference." A Location Object is represented by reference when it is represented by a URI that can be dereferenced to obtain the LO. This document summarizes the concerns about location by reference that are discussed at length in [13] (Marshall, R., “Requirements for a Location-by-Reference Mechanism,” November 2009.).

Distribution of location by reference (distribution of location URIs) offer several benefits. From a practical perspective, it can make location more compact, more recent, and more easily discoverable. Location URIs are a more compact way of transmitting location, since URIs are usually smaller than LOs. A recipient of location can make multiple requests to a URI over time to receive updated location (if the URI is configured to provide fresh location rather than a single "snapshot"). Location URIs can serve as an "LS discovery" mechanism, in that an entity can provide a location URI to an LS or LR in order to inform the recipient about which LS it should query to obtain a Location Object.

From a positioning perspective (i.e., for an LG), location by reference can offer the additional benefit of "just in time" positioning. If location is distributed by reference until it is needed for consumption, the LG (here acting as the referenced LS) only needs to perform positioning operations when a recipient makes a request for location.

From a privacy perspective, distributing location as a URI instead of a Location Object can help protect privacy by forcing each recipient of the location to request location from the referenced LS, which can then apply access controls individually to each recipient. Note, however, that the benefit provided here is contingent on the LS applying access controls. If the LS does not apply an access control policy to requests for a location URI (in other words, if enforces the "possession model" defined in [13] (Marshall, R., “Requirements for a Location-by-Reference Mechanism,” November 2009.)), then transmitting a location URI presents the same privacy risks as transmitting the Object itself. Moreover, the use of location URIs without access controls can introduce additional privacy risks: If URIs are more predictable than the location (e.g., if they are issued in sequence), then an attacker to whom the URI has not been sent may be able to guess the URI and use it to obtain the referenced LO. To mitigate this, location URIs without access controls MUST be constructed so that they are unpredictable.

3.2.3. Privacy Considerations

TOC

Two types of information are distributed in the location distribution process: location information and Privacy Rules. Rules can be communicated either independently (through a Rules Protocol) between an RM and an LS, or as part of a Location Object carrying location information. Location information, however, MUST always be accompanied by Rules -- otherwise, a recipient (whether an LS or an LR) will not know what uses are authorized, and will not be able to use the LO. Consequently, LO formats MUST be able to express Rules that convey appropriate authorizations.

An LS MUST only accept Rules from authorized Rule Makers. For an LSi, this requirement is met by applying the Rules provided in a LO to the distribution of that LO. For an LSa, this requirement means that the LS MUST be configurable with an RM authorization policy. An LS SHOULD define a prescribed set of RMs that may define Rules for a given Target or LO. For example, an LS may only allow the Target to set Rules for itself, or it might allow an RM to set Rules for several Targets (e.g., a parent for children, or a corporate security officer for employees).

No matter how Rules are provided to an LS, for each LO it receives, it MUST combine all Rules that apply to the LO into a rule set that defines which transmissions are authorized, and it MUST transmit location only in ways that are authorized by these Rules.

For an LSi, all Rules are provided in the LO. When an LSi receives a LO, it MUST examine the Rules that accompany that LO in order to determine how it may use the LO (if any Rules are included by reference, the LSi SHOULD attempt to download them). If the LO includes no Rules that allow the LSi to transmit the LO to another entity, then the LSi MUST NOT transmit the LO. It may, however use the LO for other purposes, e.g., logging, if these other actions are authorized. If the LO contains no Rules at all (e.g., if it is in a format with no Rules syntax), then the LSi MUST delete it.

When an LSa receives a LO, it MUST combine the Rules in the LO with Rules it has received from RMs. The strategy the LSa uses to combine these sets of Rules is a matter for local policy, depending on the relative priority that the LS grants to each source of Rules. Some example policies:

Union: A transmission of location is authorized if it is authorized by either a rule in the LO or an RM-provided rule.

Intersection: A transmission of location is authorized if it is authorized by both a rule in the LO and an RM-provided rule.

RM Override: A transmission of location is authorized if it is authorized by an RM-provided rule (regardless of the LO Rules)

LO Override: A transmission of location is authorized if it is authorized by a LO-provided rule (regardless of the RM Rules)

In general, it is RECOMMENDED that an LSi follow either the "Intersection" policy, since it grants equal weight to all RMs (including the LO creator). In cases where an external RM is more trusted than the source of the LO, the "Override" policy may be more suitable (e.g., if the external RM is the Target, and the LO is provided by a third party). Conversely, the "LO Override" policy is best suited to cases where the LO provider is more trused than the RM (e.g., if the RM is the user of a mobile device LS and the LO is provided with Rules from the RM's parents or corporate security office).

3.2.4. Security Considerations

TOC

An LS's decisions about how to transmit location are based on the identities of entities requesting information and other aspects of requests for location. In order to ensure that these decisions are made properly, the LS needs assurance of the reliability of information on the identities of the entities with which the LS interacts (including LRs, LSs, and RMs) and other information in the request.

Conveyance Protocols and Rules Protocols MUST provide information on the identity of the recipient of location (an LR or LS) and the identity of the RM, respectively. In order to ensure the validity of this information, these protocols MUST allow for mutual authentication of both parties, and MUST provide integrity protection for protocol messages. These security features ensure that the LG has sufficient information (and sufficiently reliable information) to make privacy decisions.

As they travel through the Internet within a Conveyance Protocol, Location Objects necessarily pass through a sequence of intermediaries, ranging from layer-2 switches to IP routers to application-layer proxies and gateways. The ability of an LS to protect privacy by making access-control decisions is reduced if these intermediaries have access to a Location Object as it travels between privacy-preserving entities.

A Conveyance Protocol MUST provide end-to-end confidentiality between an LS that transmits location and the LS or LR that receives it. When the protocol itself is protected end-to-end between the LS and the recipient, carrying an unprotected Location Object within this encrypted channel is sufficient. When the protocol has a mode in which messages are either unprotected or protected on a hop-by-hop basis (e.g., between intermediaries in a store-and-forward protocol), the protocol SHOULD allow the use of encrypted LOs, or for the transmission of a reference to location in place of a LO [13] (Marshall, R., “Requirements for a Location-by-Reference Mechanism,” November 2009.).
It is RECOMMENDED that Rule Makers, Location Servers, and Location Recipients use the security features of Rules Protocols and Coveyance Protocols to ensure that Rules are installed and applied properly, and that location is protected en route.

3.3. Receipt of Location Information

TOC

After location information has been distributed via a series of Location Servers, it finally comes to rest with a Location Recipient. Location Recipients are consumers of location; they do not forward location information to other entities. (Any recipient of location information that forwards it to other entities is acting as a Location Server in the distribution chain. The privacy requirements for an LS are described in Section 3.2 (Location Distribution).)
The primary privacy requirement of an LR is to constrain its usage of location to the set of uses authorized by the Rules in an LO. If an LR only uses a LO in ways that do not have a privacy impact -- specifically, if it does not transmit the LO to any other entity, and does not retain the LO for longer than is required to execute the Conveyance Protocol -- then no further action is necessary for the LR to comply with the requirements of this document.
As an example of this simplest case, if a Location Recipient (a) receives a location, (b) immediately provides to the Target information or a service based on the location, (c) does not retain the information, and (d) does not retransmit the location to any other entity, then the LR will comply with any set of Rules that are permissible under Geopriv. Thus, a service that, for example, only provides directions to the closest bookstore in response to an input of location, and promptly then discards the input location, will be in compliance with any Geopriv rule set.
LRs that make other uses of a LO (e.g., those that store LOs, or send them to other service providers to obtain location-based services) MUST meet the requirements below to assure that these uses are authorized.
The Location Receipt process thus involves one Geopriv role and two protocols:

> **Location Recipients:** Entities that accept LOs and use them, subject to the Privacy Rules they contain
>
> **Conveyance Protocols:** Protocols by which LOs are delivered to Location Recipients
>
> **LO Formats:** The structure of a LO, including the available location and Rules semantics.

3.3.1. Privacy Considerations

TOC

The principle privacy requirement for Location Recipients is to follow usage rules. When an LR receives a LO, it is REQUIRED to examine the Rules included with that LO. Any usage the LR makes of the LO MUST be explicitly authorized by these Rules. Since Rules are positive grants of permission, any action not explicitly authorized is denied by default.

In particular, given a LO in a particular format, an LR MUST NOT take any action that could be authorized by a rule within that format, unless such an rule is present in the LO to authorize the action. If such an action were authorized, then the RM would have included a rule to express this authorization. For instance, the PIDF-LO format [14] (Peterson, J., “A Presence-based GEOPRIV Location Object Format,” December 2005.) defines a rule that allows an LR to retain the LO for a specified amount of time; if an LR receives a LO that does not have such a rule, then it MUST NOT retain the location.

3.3.2. Security Considerations

TOC

Since a Location Recipient does not transmit location, there are no protocol security considerations required to support privacy (only the LR's compliance with Rules, as described above).

Aside from privacy, Location Recipients often require some assurance that a LO is reliable (assurance of the integrity, authenticity, and validity of an LO), since LRs use LOs in order to deliver location-based services. Threats against this reliability and corresponding mitigations are discussed in the Security Considerations below.

TOC

4. Security Considerations

Security considerations related to the privacy of Location Objects are discussed throughout this document. In this section we summarize those concerns and consider security risks not related to privacy.

The life-cycle of a Location Object often consists of a series of location transmissions (see Figure 4 (Location Life-Cycle)). For example, location might initially be published to a location configuration server which then transmits the location to the Target. The Target may then act as a Location Server and convey this location to a service provider (who acts as Location Recipient in this transmission) to facilitate some location-based service.

(Note that although Figure 4 (Location Life-Cycle) depicts a single "path", a single location server may transmit location to multiple location recipients over time; groups of these paths together form a logical distribution tree, with the location generator as the root node.)

```
      +-----+    +-----+    +-----+    +-----+    +-----+
      | LG |--->| LS |--->| LS |--->| LS |--->| LR |
      +-----+    +-----+    +-----+    +-----+    +-----+
                    |          |          |
                 +-----+    +-----+    +-----+
                 | RM |     | RM |     | RM |
                 +-----+    +-----+    +-----+
```

Figure 4: Location Life-Cycle

The location life-cycle gives rise to additional security concerns. For example, in a scenario where some intermediate location servers are untrusted, a location recipient may desire additional assurances that the LO was generated by a trusted LG, and not modified by these untrusted entities. In this section, we first consider threats and possible attacks against a Location Object throughout its entire life cycle. We then describe the assurances that various parties require to mitigate these threats. Finally, we discuss possible mechanisms that protocols or location object formats should make available to provide such assurances.

TOC

4.1. Threats to Location Objects

The major threats to the end-to-end security of Location Objects can be grouped into two categories: First, threats against the integrity and authenticity of Location Objects can expose entities that rely on Location Objects to many types of fraud. Second, threats against the confidentiality of Location Objects can reduce the ability of location servers to control access to location.

4.1.1. Threats to Location Integrity and Authenticity

TOC

A Location Object contains four essential types of information: Identifiers for the described Target, location information, time-stamps, and Rules. By grouping values of these various types together within a single structure, a Location Object encodes a set of bindings among them. That is, the Location Object asserts that the identified Target was present at the given location at the given time; and that the given Rules express the Target's desired policy, at the given time, for the distribution of his location. Below, we provide a set of attacks that a malicious party (e.g. an intermediate LS, an eavesdropper on the path between LS and LR, or the Target himself) might conduct to falsify one or more of the bindings asserted by the Location Object.

Note that in all cases the Target identity provided in a Location Object should be based on an authentication between the Target and the location generator (e.g. an explicit authentication based on a shared secret, or an implicit authentication based on the ability to receive a message). Therefore, the identity binding in a received Location Object is only as strong as the authentication between the Target and the location generator (that is, the Location Object can only attest to the fact that someone at the given location is capable of authenticating as the given identity). It is vital to the authenticity of location information that this authentication be as strong as is feasible in any deployment scenario. However, mechanisms within a Geopriv Location Object or protocol can provide no protection from attacks against this authentication mechanism and thus we do not explicitly consider such attacks.

Place Shifting: Falsifying the location in an otherwise valid Location Object. For example, Alice pretends to that she is currently in a location that she has never previously visited.

Time Shifting: Falsifying the time-stamp in an otherwise valid Location Object. For example, Alice pretends that she is currently in a location that she has not visited since last year.

Location Theft: Falsifying the identity in an otherwise valid Location Object. For example, a malicious intermediary sees a valid Location Object for Alice and produces a Location Object asserting that Bob is the given location at the given time.

Location-Identity Theft: An attacker replays a stale Location Object as though it were current. For example, a malicious intermediary sees a valid Location Object for Alice and replays it later to make it seem that Alice has not moved.

Location Swapping: Two malicious Targets conspire to produce two Location Objects asserting that each Target is at the other's location. For example, Alice pretends that she is at Bob's location and Bob pretends that he is at Alice's location. (Note that this attack cannot be prevented if the two attackers are willing to exchange authentication credentials. Because the identity assertions in a Location Object are only as strong as the Target authentication, the goal of Geopriv protocols is to ensure that this attack is not possible unless both Alice and Bob can successfully authenticate as the other.)

4.1.2. Threats to Location Privacy

TOC

In the Geopriv model, the privacy of location information is protected by the application of Privacy Rules specified by authorized rule makers, and by confidentiality protection en route. (For more information on privacy rule enforcement, see Section 3.2.3 (Privacy Considerations).) Below, we provide a set of attacks that a malicious party might conduct to allow distribution of a Location Object to unauthorized parties.

Eavesdropping: An unauthorized party observes the Location Object in transit. For example, a device on the path between a trusted

> LS and an authorized LR observes a Location Object sent in the clear.
>
> **Rule Tampering:** A malicious party modifies a Target's Privacy Rules and thus causes a trusted LS to unknowingly distribute the Location Object to unauthorized parties. For example, a device on the path between an LG and a trusted LS deletes the Privacy Rules contained in a Location Object and replaces them with a new set of Rules authorizing all parties to receive the Location Object.
>
> **Server Impersonation:** A malicious party impersonates a trusted location server and then knowingly disregards the Privacy Rules. For example, a man-in-the-middle between the LG and the trusted LS pretends to be the trusted LS, and then proceeds to distribute the Location Object to unauthorized entities.

4.2. Required Assurances

TOC

We now describe the assurances required by each party involved in location distribution in order to mitigate the attacks described in the previous two sections:

> **Rule Maker:** The rule maker is responsible for distributing the Target's Privacy Rules to the location servers. The primary assurance required by the Rule Maker is thus that the binding between the Target's Privacy Rules and the Target's identity is correctly conveyed to each location server that handles the Location Object. Ensuring the integrity of the Privacy Rules distributed to the location servers prevents rule-tampering attacks. (Note that in many circumstances, the privacy policy of the Target may itself be sensitive information, in these cases the Rule Maker also requires the assurance that the binding between the Target's identity and the Target's Privacy Rules are not deducible by anyone other than an authorized location server).
>
> **Location Server:** The Location Server is responsible for enforcing the Target's privacy policy. The first assurance required by the location server is that the binding between the Target's Privacy Rules and the Target's identity is authentic. Authenticating the rule-maker who created the Privacy Rules prevents rule-tampering attacks. The second assurance required by the location server is that the binding between the Target's identity and the Target's location are not deducible by any entity except as allowed the Target's privacy policy. Ensuring the confidentiality of these bindings prevents eavesdropping attacks. (Note that ensuring the

confidentiality of the Location Object also helps to mitigate location-theft and location-identity-theft attacks, since it makes it more difficult for an attacker to obtain a valid Location Object to replay.)

Location Recipient: The Location Recipient is the end consumer of the Location Object. The location recipient thus requires assurances about the authenticity of the bindings between the Target's location, the Target's identity and the time. Ensuring the authenticity of these bindings prevents place-shifting, time-shifting, location-theft, and location-identity-theft attacks; and mitigates location-swapping attacks to the greatest possible extent.

Location Generator: The Location Generator shares responsibility for ensuring that the Target's privacy policy is enforced. The primary assurance required by the Location Generator is that the Location Server to which the Location Object is initially published is one that is trusted to enforce the Target's privacy policy. Authenticating the trusted Location Server mitigates the risk of server impersonation attacks. (Additionally, in some scenarios, there may be no Location Server which can be trusted to sufficiently safe-guard the Target's location information, in which case the Location Generator may require assurance that intermediate location servers are unable to deduce the binding between the Target's identity and the Target's location.)

4.3. Protocol mechanisms

TOC

Protocols that carry location can provide strong assurances, but only for a single segment of the Location Object's life cycle. In particular, a protocol can provide integrity protection and confidentiality for the data exchanged, and mutual authentication of the parties involved in the protocol, by using a secure transport such as IPsec or TLS.

Additionally, note that if (1) the protocol provides mutual authentication for every segment; and (2) every entity in the location distribution exchanges information only with entities with whom it has a trust relationship, then entities can transitively obtain assurances regarding the origin and ultimate destination of the Location Object. Of course, direct assurances are always preferred over assurances requiring transitive trust, since they require fewer assumptions.

Using protocol mechanisms alone, the entities can receive assurances only about a single hop in the distribution chain. For example, suppose that an LR retrieves location from an LS over an integrity- and confidentiality-protected channel. The LR knows that the transmitted LO has not been modified or observed en route. However, the assurances provided by the protocol do not guarantee that the transmitted LO was not corrupted before it was sent (e.g., by a previous LS). Likewise, the LR can verify that the LO was transmitted by the LS, but cannot verify the origin of the LO if it is different from the LS.

Security mechanisms in protocols are thus unable to provide direct assurances over multiple transmissions of an LO. However, it should be noted that the transmission of location "by reference" can be used to effectively turn multi-hop paths into single-hop paths. If the multiple transmissions of a LO are replaced by multiple transmissions of an identifier (a multi-hop dissemination channel), then the LO need only traverse a single hop, namely the dereference transaction between the LR and the dereference server.

4.4. Mechanisms within the Location Object

TOC

Assurances as to the integrity and confidentiality of a Location Object can be provided directly through the Location Object format. Additionally, the Location Object format can be used to authenticate the originator of a Location Object. In particular, integrity and origin authentication can be assured by signing a Location Object (e.g., using S/MIME or XMLSIG), and confidentiality can be assured by encrypting the Location Object using a public encryption key belonging to the intended recipient (e.g. using S/MIME). Recipients of Location Objects secured in this fashion can obtain assurance as to the integrity and authenticity of the Location Object even after it has been handled by untrusted intermediaries. Similarly, a Location Server (or Location Generator) that guarantees confidentiality in this fashion can be assured that the Location Object is protected from unauthorized viewing even in the presence of untrusted intermediaries.

Although such direct, end-to-end assurances are desirable, and these mechanisms should be used whenever possible, there are many deployment scenarios where directly securing a Location Object is impractical. In particular, in some deployment scenarios a direct trust relationship may not exist between the creator of the Location Object and the ultimate recipient. Additionally, in a scenario where many recipients are authorized to receive a given Location Object, the creator of the Location Object cannot guarantee end-to-end confidentiality without knowing precisely which recipient will receive the Location Object.

An additional challenge in providing end-to-end authenticity guarantees by signing the Location Object is that in many deployments different entities may assert different bindings within the same Location Object. Consider, for example, a scenario where a Location Generator produces a Location Object that asserts a binding between a time, a location, and a pseudonym for the Target. Additionally, a Rule Maker creates a binding between a set of Privacy Rules and a public Target identity. A presence server receives the Rules binding from Rule Maker and the Location Object from the Location Generator. The presence server then generates a new Location Object binding together the time, the location, the public Target identity and the Privacy Rules. In such a scenario there is no single entity who can directly assert the validity of the entire Location Object. In such a case, a mechanism is needed within the Location Object format that allows multiple originators to jointly assert various components of the Location Object bindings.

5. Example Scenarios

TOC

This section contains a set of example of how the Geopriv architecture can be deployed in practice. These examples are meant to illustrate key points of the architecture, rather than to form an exhaustive set of use cases.

For convenience and clarity in these examples, we assume that the Privacy Rules that a LO carries are equivalent to those in a PIDF-LO Location Object (namely, that the principal Rules that can be set are limits on the retransmission and retention of the LO). It should be noted that while these two Rules are the most well-known and important examples, the specific types of Rules an LS or LR must consider will in general depend on the types of LO it processes. It is possible that in some cases, Geopriv entities will have to consider additional Rules and in others, retention and retransmission will be unconstrained. For the rest of this section, however, we assume for simplicity that limiting retention and controlling access to location are the two primary responsibilities incumbent on a recipient of location (an LS or LR).

5.1. Minimal Scenario

TOC

One of the simplest scenarios in the Geopriv architecture is when a Target determines its own location and uses that LO to request a service (e.g., by including the LO in an HTTP POST request or SIP INVITE message), and the server delivers that service immediately (e.g., in a 200 OK response in HTTP or SIP), without retaining or retransmitting the Target's location. The Target acts as an LG by using a Target-based positioning algorithm (e.g., manual entry), as a Rule Maker by specifying that the location should be sent to the server, and as an initial Location Server by interpreting the rule and transmitting the LO. The server acts as a Location Recipient by receiving and using the LO.

In this case, the privacy of location information is maintained in two steps: The first step is that location is only transmitted as directed by the single Rule Maker, namely the Target. The second step is simply the fact that the server (i.e., the LR) did not do anything that created a privacy risk -- it did not retain or retransmit location. Because the server limits its behavior in this way, it does not need to read the Rules in the LO (even though they were provided) -- no rule would prevent it from using location in this safe manner.

The following outline summarizes this scenario:

- Positioning: Target-based, Target=LG=initial LS
- Distribution hop 1: HTTP UA --> Ephemeral web service, privacy via user indication
- Receipt: Ephemeral web service delivers response without retaining or retransmitting location
- Key points:
    - LRs that do not behave in ways that risk privacy are Geopriv-compliant by default. No further action is necessary.

5.2. Location-based Web Services

TOC

Many location-based services are delivered over the Web, using Javascript code to orchestrate a series of HTTP requests for location specific information. To support these applications, browser extensions have been developed that support Target-based positioning (manual entry and GPS) and network-assisted positioning (via AGPS, and multilateration with 802.11 and cellular signals), exposing position to web pages through Javascript APIs.

In this scenario, we consider a Target that uses a browser with a network-assisted positioning extension. When the Target uses this browser to request location-based services from a web page, the browser prompts the user to grant the page permission to access the user's location. If the user grants permission, the browser extension sends 802.11 signal strength measurements to a positioning server, which then returns the position of the host. The extension constructs a Location Object with this location and Rules set by the user, then passes the LO to the page through its Javascript API. The page then obtains location-relevant information using an XMLHttpRequest [15] (World Wide Web Consortium, “The XMLHttpRequest Object,” April 2008.) to a server in the same domain as the page and renders this information to the user.

At first blush, this scenario seems much more complicated than the minimal scenario above. However, most of the privacy considerations are actually the same.

The positioning phase in this scenario begins when the browser extension contacts the positioning server. The positioning server acts as a Location Generator, and the protocol that supports this interaction is a Positioning Protocol. The positioning server supports the privacy of the location information it provides by following the LCP Policy, i.e., by providing location information only to the entity being located.

The distribution phase actually occurs entirely within the Target host: The single hop in distribution occurs when the browser extension (an entity under the control of the Target) passes a LO to the web page (an entity under the control of its author). In this phase, the browser extension acts as the initiating LS, with the user/Target as the sole Rule Maker; the user interface for rule-making is effectively a Rules Protocol, and the extension's API effectively defines a Conveyance Protocol and LO Format. The web site acts as Location Recipient when the web page accepts the LO.

The receipt phase encompasses the web site's use of the LO. In this context, the phrase "web site" encompasses not only the web page, but also the dedicated supporting logic behind it. Considering the entire web site as a recipient, rather than a single page, it becomes clear that by sending the LO in an XMLHttpRequest to a back-end server is more like passing it to a separate component of the LR (as opposed to retransmitting it to another entity). Thus, even in this case, where location-relevant information is obtained from a back-end server, the LR does not retain or retransmit location, so its behavior is "privacy-safe" -- it doesn't need to interpret the Rules in the LO.

However, consider a variation on this scenario where the web page requests additional information (e.g., a map) from a third-party site. In this case, since location is being transmitted to a third party, the web site (either in the web page or in a back-end server) would need to verify that this transmission is allowed by the LO's Privacy Rules. Similarly, if the site wanted to log the user's location information, then it would need to examine the LO to determine how long this information can be retained. In such a case, if the LR needs to do something that is not allowed by the Rules, it may have to deny service to the user (hopefully providing a message with the reason). Nonetheless, if the Rules permit retention or retransmission (even if this retransmission is limited by access control rules), then the LR may do so to the extent the Rules allow.

The following outline summarizes this scenario:

*Positioning: Network-assisted, positioning server=LG, privacy via LCP Policy

*Rule installation: RM (=Target/user) gives permission to sites and sets LO Rules

*Distribution hop 1: Browser=LS --> Web site=LR, privacy via user confirmation

*Receipt: Back-end server delivers location-relevant information without further retransmission, then deletes location; privacy via safe behavior

*Key points:

- Privacy in this scenario is provided by a combination of explicit user direction and Rules in an LO

- Distribution can occur within a host, between mutually untrusting components

- Some transmissions of location are actually internal to an LR

- LRs that do things that might be constrained by Rules need to verify that these actions are allowed for a particular LO

5.3. Emergency Calling

TOC

Support for emergency calls by Voice-over-IP devices is a critical use case for location information about Internet hosts. The details of the Internet architecture for emergency calling are described in [16] (Rosen, B., Schulzrinne, H., Polk, J., and A. Newton, “Framework for Emergency Calling using Internet Multimedia,” July 2009.)[17] (Rosen, B. and J. Polk, “Best Current Practice for Communications Services in support of Emergency Calling,” January 2010.). In this architecture, there are three critical steps in the placement of an emergency call, each involving location information:

1. Determine the location of the caller

2. Determine the proper Public Safety Answering Point (PSAP) for the caller's location

3. Send a SIP INVITE message (including the caller's location) to the PSAP

The first step in an emergency call is to determine the location of the caller. This step is the positioning phase of the location life-cycle. Location is determined by whatever means are available to the caller's device, or to the network, if this step is being done by a proxy. Whichever entity does the positioning (either the caller or a proxy) acts as the initiating Location Server, preserving the privacy of location information by only including it in emergency calls.

The second step in an emergency call encompasses location distribution and receipt. The entity that is routing the emergency call sends location though the LoST protocol [18] (Hardie, T., Newton, A., Schulzrinne, H., and H. Tschofenig, “LoST: A Location-to-Service Translation Protocol,” August 2008.) to a mapping server. In this role, the routing entity acts as a Location Server, LoST acts as a Location Conveyance protocol, and the LoST server acts as a Location Recipient. The LO format within LoST does not allow Rules to be sent along with location, but because LoST is an application-specific protocol, the sending of location within a LoST message authorizes the LoST server to use the location to complete the protocol, namely to route the message as necessary through the LoST mapping architecture [19] (Schulzrinne, H., “Location-to-URL Mapping Architecture and Framework,” March 2009.). That is, the LoST server is authorized to complete the LoST protocol, but to do nothing else.

The third step in an emergency call is again a combination of distribution and receipt. The caller (or another entity that inserts the caller's location) acts as an LS, SIP acts as a Conveyance Protocol [20] (Polk, J. and B. Rosen, “Location Conveyance for the Session Initiation Protocol,” March 2009.), and the PSAP acts as a Location Recipient. In this specific example, the caller's location is transmitted either as a PIDF-LO object or as a reference that returns a PIDF-LO (or both); in the latter case, the reference should be appropriately protected so that only the PSAP has access. In any case, the receipt of a LO implies that the PSAP should obey the Rules in those LOs in order to preserve privacy. Depending on the regulatory environment, the PSAP may have the option to ignore those constraints in order to respond to an emergency, or it may be bound to respect these Rules (in spite of the emergency situation).

The following outline summarizes this scenario:

- Positioning: Any, Target=initial LS
- Distribution/receipt hop 1: Target=LS --> LoST infrastructure (no Rules), privacy via authorization implicit in protocol
- Distribution/receipt hop 2: Target=LS --> PSAP, privacy via Rules in LO
- Receipt: PSAP uses location to deliver emergency services
- Key points:
  - Privacy in this scenario is provided by a combination of explicit user direction, implicit authorization particular to a protocol, and Rules in an LO
  - LRs may be constrained to respect or ignore Privacy Rules by local regulation

5.4. Combination of Services

TOC

In modern Internet applications, users frequently receive information via one channel and broadcast it via another. In this sense, both users and channels (e.g., web services) become location servers. Here we consider a more complex example that illustrates this pattern across multiple logical hops.

Suppose Alice (the Target) subscribes to a wireless ISP that determines her location using a network-based positioning technique (e.g., via the location of the base station serving the Target), and provides that information directly to a location-enhanced presence provider (which might use SIP, XMPP, or another protocol). The location-enhanced presence provider allows Alice to specify Rules for how this location is distributed: which friends should receive Alice's location and what Rules they should get with it. Alice uses a few other location-enhanced services as well, so she sends Rules that allows her location to be shared with those services, and allows those services to retain and retransmit her location.

Bob is one of Alice's friends, and he receives her location via this location-enhanced presence service. Noting that she's at their favorite coffee shop, Bob wants to upload a photo of the two of them at the coffee shop to a photo-sharing site, along with a LO that marks the location. Bob checks the Rules in Alice's LO and verifies that the photo sharing site is one of the services that Alice authorized. Seeing that Alice has authorized him to give the LO to the photo-sharing site, he attaches it to the photo and uploads it.

Once the geo-tagged photo is uploaded, the photo sharing site reads the Rules in the LO and verifies that the site is authorized to store the photo and to share it with others. Since Alice has allowed the site the retransmit and retain without any constraints, the site fulfills Bob's request to make the geo-tagged photo publicly accessible.

Eve, another user of the photo sharing site, downloads the photo of Alice and Bob at the coffee shop and receives Alice's LO along with it. Eve posts the photo and location to her public page on a social networking site without checking the Rules, even though the LO doesn't allow Eve to send the location anywhere else. The social networking site, however, observes that no retransmission or retention are allowed (both of which it needs for a public posting), and rejects the upload.

In terms of the location life-cycle, this scenario consists of a positioning step, followed by four distribution hops and receipt.

Positioning is the simplest step: An LG in Alice's ISP monitors her location and transmits it to the presence service, maintaining privacy by only transmitting location to a single entity (to which privacy responsibilities are delegated).

The first distribution hop occurs when the presence server sends location to Bob. In this transaction, the presence server acts as an LS, Alice acts as an RM, and Bob acts as another LS (on the receiving side). The privacy of this transaction is assured by the fact that

Alice has installed Rules on the presence server that dictate who it may allow to access her location. The second distribution hop is when Bob uploads the LO to the photo-sharing site. Here Bob again acts an LS (on the sending side), preserving the privacy of location information by verifying that the Rules in the LO allow him to upload it. The third distribution hop is when the photo-sharing site sends the LO to Eve, likewise following the Rules -- but a different set of Rules than Bob, since a LO can specify different rulesets for different Location Servers.

Eve is the fourth LS in the chain, and fails to comply with Geopriv by not checking the rule in the LO prior to uploading it to the social networking site. The site, however, is a responsible LR -- it checks the Rules in the LO, sees that they don't allow it to use the location as it needs to, and discards the LO.

The following outline summarizes this scenario:

*Positioning: Network-based, LG in network, privacy via exclusive relationship with presence service

*Distribution hop 1: Presence server --> Bob, privacy via Alice's access control rules (installed via Rules Protocol)

*Distribution hop 2: Bob --> photo sharing site, privacy via Rules for Bob in LO

*Distribution hop 3: Photo sharing site --> Eve, privacy via Rules for site in LO

*Distribution hop 4: Eve --> Social networking site, violates privacy by retransmitting

*Recipient: Social networking site, privacy via checking Rules and discarding

*Key points:

-Privacy can be preserved through multiple hops

-A LO can specify different Rules for different entities

-An LS can still disobey the Rules, but even then, the architecture still works in some cases

TOC

6. Glossary

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [1] (Bradner, S., “Key words for use in RFCs to Indicate Requirement Levels,” March 1997.).

$ **Access Control Rule** A rule that describe which entities may receive location information and in what form.

$ **Authorized Location Server (LSa)** A Location Server that receives Rules from a Rule Maker (in addition to Rules provided in LOs). An Authorized Location Server may receive Location Objects containing Privacy Rules from Location Generators and other Location Servers, and it may also receive Privacy Rules directly from Rule Makers.

$ **civic location** The geographic position of an entity in terms of a postal address or civic landmark. Examples of such data are room number, street number, street name, city, ZIP code, county, state and country.

$ **geodetic location** The geographic position of an entity in a particular coordinate system (for example, a latitude-longitude pair).

$ **Independent Location Server (LSi)** A Location Server that has no relationship with a Rule Maker. An Independent Location Server may receive Location Objects containing Privacy Rules from Location Generators and other Location Servers, but it does not receive Privacy Rules directly from Rule Makers.

$ **Local Rule** A Privacy Rules that directs a Location Server about how to treat a Target's location information. Local Rules are used internally by a Location Server to handle requests from Location Recipients. They are not distributed to Location Recipients.

$ **Location Generator (LG)** An entity that initially determines or gathers the location of a Target. Location Generators may be any sort of software or hardware used to obtain a Target's position (examples include GPS chips and cellular networks).

$ **Location Information Server (LIS)** An entity responsible for providing devices within an access network with information about their own locations. A Location Information Server uses knowledge of the access network and its physical topology to generate and distribute location information to devices.

$ **Location Object (LO)** A data unit that conveys location information together with Privacy Rules within the Geopriv architecture. A Location Object may convey geodetic location data (latitiude/longitude/altitude), civic location data (street/city/state/etc.), or both.

$ **Location Recipient (LR)** An ultimate end point entity to which a Location Object is distributed. Location Recipients request location information about a particular Target from a Location Server. If allowed by the appropriate Privacy Rules, a Location Recipient will receive Location Objects describing the Target's location from the Location Server.

$ **Location Server (LS)** An entity that receives Location Objects from Location Generators, Privacy Rules from Rule Makers, and location requests from Location Recipients. A Location Server applies the appropriate Privacy Rules to a Location Object received from a Location Generator and may disclose the Location Object, in compliance with the Rules, to Location Recipients.

Location Servers may not necessarily be "servers" in the colloquial sense of hosts in remote data centers servicing requests. Rather, a Location Server can be any software or hardware component that receives and distributes location information. Examples include a positioning server (with a location interface) in an access network, a presence server, or a Web browser or other software running on a Target's device.

$ **Privacy Rule** A directive that regulates an entity's activities with respect to a Target's location information, including the collection, use, disclosure, and retention of the location information. Privacy Rules describe how location information may be used by an entity, the level of detail with which location information may be described to an entity, and the conditions under which location information may be disclosed to an entity. Privacy Rules are communicated from Rule Makers to Location Servers and conveyed in Location Objects throughout the Geopriv architecture.

$ **Rule** See Privacy Rule.

$ **Rule Maker (RM)** An individual or entity that is authorized to set Privacy Rules for a Target. In some cases a Rule Maker and a Target will be the same individual or entity, and in other cases they will be separate. For example, a parent may serve as the Rule Maker when the Target is his child. The Rule Maker is also not necessarily the owner of a Target device. For example, a corporation may own a device that it provides to an employee but permit the employee to serve as the Rule Maker and set her own Privacy Rules. Rule Makers provide the Privacy Rules associated with a Target to Location Servers.

$ **Forwarded Rule** A Privacy Rule that travels inside a Location Object. Forwarded Rules direct Location Recipients about how to handle the location information they receive. Because the Forwarded Rules themselves may reveal potentially sensitive information about a Target, only the minimal subset of Forwarded Rules necessary for a Location Recipient to handle a Location Object is distributed to the Location Recipient.

$ **Target** An individual or other entity whose location is described by a Location Object. The Target is the entity whose privacy Geopriv seeks to protect.

$ **Usage Rule** A rule that describe what uses of location information are authorized.

7. Acknowledgements

TOC

This work was largely based on the security investigations conducted as part of the Geopriv Layer-7 Location Configuration Protocol design team, which produced [5] (Tschofenig, H. and H. Schulzrinne, “GEOPRIV Layer 7 Location Configuration Protocol; Problem Statement and Requirements,” July 2009.). We would like to thank all the members of the design team.

8. IANA Considerations

TOC

This document makes no request of IANA.

9. References

TOC

9.1. Normative References

TOC

[1] Bradner, S., “Key words for use in RFCs to Indicate Requirement Levels,” BCP 14, RFC 2119, March 1997 (TXT, HTML, XML).

9.2. Informative References

TOC

[2] Cuellar, J., Morris, J., Mulligan, D., Peterson, J., and J. Polk, “Geopriv Requirements,” RFC 3693, February 2004 (TXT).

[3] Danley, M., Mulligan, D., Morris, J., and J. Peterson, “Threat Analysis of the Geopriv Protocol,” RFC 3694, February 2004 (TXT).

[4] U.S. Department of Defense, “National Industrial Security Program Operating Manual,” DoD 5220-22M, January 1995.

[5] Tschofenig, H. and H. Schulzrinne, “GEOPRIV Layer 7 Location Configuration Protocol; Problem Statement and Requirements,” draft-ietf-geopriv-l7-lcp-ps-10 (work in progress), July 2009 (TXT).

[6] Polk, J., Schnizlein, J., and M. Linsner, “Dynamic Host Configuration Protocol Option for Coordinate-based Location Configuration Information,” RFC 3825, July 2004 (TXT).

[7] Schulzrinne, H., “Dynamic Host Configuration Protocol (DHCPv4 and DHCPv6) Option for Civic Addresses Configuration Information,” RFC 4776, November 2006 (TXT).

[8] Polk, J., “Dynamic Host Configuration Protocol (DHCP) IPv4 and IPv6 Option for a Location Uniform Resource Identifier (URI),” draft-ietf-geopriv-dhcp-lbyr-uri-option-07 (work in progress), March 2010 (TXT).

[9] 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).

[10] Schulzrinne, H., Tschofenig, H., Morris, J., Cuellar, J., Polk, J., and J. Rosenberg, “Common Policy: A Document Format for Expressing Privacy Preferences,” RFC 4745, February 2007 (TXT).

[11] Schulzrinne, H., Tschofenig, H., Morris, J., Cuellar, J., and J. Polk, “Geolocation Policy: A Document Format for Expressing Privacy Preferences for Location Information,” draft-ietf-geopriv-policy-21 (work in progress), January 2010 (TXT).

[12] Rosenberg, J., “The Extensible Markup Language (XML) Configuration Access Protocol (XCAP),” RFC 4825, May 2007 (TXT).

[13] Marshall, R., “Requirements for a Location-by-Reference Mechanism,” draft-ietf-geopriv-lbyr-requirements-09 (work in progress), November 2009 (TXT).

[14] Peterson, J., “A Presence-based GEOPRIV Location Object Format,” RFC 4119, December 2005 (TXT).

[15] World Wide Web Consortium, “The XMLHttpRequest Object,” W3C document http://www.w3.org/TR/XMLHttpRequest/, April 2008.

[16]

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).

[17] Rosen, B. and J. Polk, “Best Current Practice for Communications Services in support of Emergency Calling,” draft-ietf-ecrit-phonebcp-14 (work in progress), January 2010 (TXT).

[18] Hardie, T., Newton, A., Schulzrinne, H., and H. Tschofenig, “LoST: A Location-to-Service Translation Protocol,” RFC 5222, August 2008 (TXT).

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

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

Authors' Addresses

TOC

Richard Barnes
BBN Technologies
9861 Broken Land Pkwy, Suite 400
Columbia, MD 21046
USA
Phone: +1 410 290 6169
Email: rbarnes@bbn.com

Matt Lepinski
BBN Technologies
10 Moulton St
Cambridge, MA 02138
USA
Phone: +1 617 873 5939
Email: mlepinski@bbn.com

Alissa Cooper
Center for Democracy & Technology
1634 I Street NW, Suite 1100
Washington, DC
USA
Email: acooper@cdt.org

John Morris
Center for Democracy & Technology
1634 I Street NW, Suite 1100

Washington, DC
USA
Email: jmorris@cdt.org

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