An Architecture for Location and Location Privacy in Internet Applications
draft-ietf-geopriv-arch-03

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.

Status of this Memo

Copyright Notice

Table of Contents

1. Introduction

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 use by its ultimate recipient(s) -- Geopriv provides mechanisms to secure 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][3], 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 provides a general introduction to Geopriv and Internet location-based services, and is useful as a good first document for readers new to Geopriv.

1.1. Binding Rules to Data

A central feature of the Geopriv architecture is that location information is always bound to privacy rules to ensure that entities that receive location are informed of how they may use it. These rules can convey simple directives ("do not share my location with others"), or more robust preferences ("allow my spouse to know my exact location all of the time, but only allow my boss to know it

during work hours"). 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 [28] 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.

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). The United States Department of Defense guidelines for
classification [4] provide one example.

1.2. Location-Specific Privacy Risks

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 data may be collected everywhere and at any time, often
without user interaction, and it may potentially describe both what a
person is doing and where he or she is doing it. For example,
location data can reveal the fact that an individual was at a
particular medical clinic at a particular time. 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 is also 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

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 otherwise 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 are informed of relevant privacy rules, and are expected to
only use location in accordance with those rules. The distributed
nature of the architecture inherently limits the degree to which
compliance can be guaranteed and verified by technical means.
Section 5 describes how some security mechanisms can address this to
a limited extent.

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
individuals whose location information has been abused.

2. Terminology Conventions

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].

3. Overview of the Architecture

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 use -- 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.

3.1. Basic Geopriv Scenario

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 (use).

For example, suppose Alice is using a mobile device, she learns of
her location from a wireless location service, and she wishes to
share her location privately with her friends by way of a presence
service. Alice clearly needs to provide the presence server with her
location and rules about which friends can be provided with her
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. To operate in a privacy-protective way, 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  |                          Retrieve
+----------+                      Access Control Rules
     |                      +------------------------------------+
     |                      | +----------------------------+    |
  Location                  | |            Access          |    |
     |                      | |         Control Rules       v    |
     |                      | |                         +-----+   |
     |                      | |                         |     |   |
     |                      | |                         | Bob |--> ...
     |                      | |                  +----->|     |   |
     |                      | |                  |      +-----+   |
     v                      v |                  |         |
  +----------+          +----------+             |
  | |Device| |--Location->| Presence |--Location---->|     +----------+
  | -------- |            | Server   |              |---->| Friend-1 |
  |          |---Rules--->|          |    |---Rules------>|     +----------+
  |  Alice   |            +----------+             |
  +----------+                                     |
                                                   |     +----------+
                                             +---->| Friend-2 |
                                                   +----------+
```

                    Figure 1: Basic Geopriv Scenario

3.2. Roles and Data Formats

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

   Target: An individual or other entity whose location is sought in
      the Geopriv architecture. In many cases the Target will be the
      human user of a Device, but it can also be an object such as a
      vehicle or shipping container to which a Device is attached. In
      some instances the Target will be the Device itself. The Target
      is the entity whose privacy Geopriv seeks to protect. Alice is
      the Target in Figure 1.

   Device: The technical device whose location is tracked as a proxy
      for the location of a Target. Alice's device is the Device in
      Figure 1.

Rule Maker (RM):   Performs the role of creating rules governing
   access to location information for a Target. In some cases the
   Target performs the Rule Maker role (as is the case with Alice),
   and in other cases they are separate. For example, a parent may
   serve as the Rule Maker when the Target is his child, or a
   corporate security officer may serve as the Rule Maker for devices
   owned by the corporation but used by employees. The Rule Maker is
   also not necessarily the owner of the 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):   Performs the roles of initially
   determining or gathering the location of the Device and providing
   it to Location Servers. Location Generators may be any sort of
   software or hardware used to obtain the Device's location
   (examples include GPS chips and cellular networks). A Device may
   even perform the Location Generator role for itself; Devices
   capable of unassisted satellite-based positioning and Devices that
   accept manually entered location information are two examples.
   The wireless location service plays the Location Generator role in
   Figure 1.

Location Server (LS):   Performs the roles of receiving location
   information and rules, applying the rules to the location
   information to determine what other entities, if any, can receive
   location information, and providing the location to Location
   Recpients. Location Servers receive location information from
   Location Generators and rules from Rule Makers, and then apply 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 distributes
   location information. Examples include a server in an access
   network, a presence server, or a Web browser or other software
   running on a Device. The above example includes three Location
   Servers: Alice, the presence service and Bob.

Location Recipient (LR):   Performs the role of receiving location
   information. A Location Recipient may ask for location explicitly
   (by sending a query to a Location Server), or it may receive
   location asynchronously. The presence service, Bob, Friend-1 and
   Friend-2 are Location Recipients in Figure 1.

In general, these roles may or may not be performed by physically
separate entities, as demonstrated by the entities in Figure 1, many
of which perform multiple roles. It is not uncommon for the same
entity to perform both the Location Generator and Location Server
roles, or both the Location Recipient and Location Server roles. A

single entity may take on multiple roles simply by virtue of its own
capabilities and the permissions provided to it.

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.

There is a close relationship between a Target and its Device. The
term "Device" is used when discussing protocol interactions, whereas
the term "Target" is used when discussing generically the person or
object being located and its privacy. While in the example above
there is a one-to-one relationship between the Target and the Device,
Geopriv can also be used to convey location information about a
device that is not directly linked to a single individual or object,
such as a Device shared by multiple individuals.

Two data formats are necessary within this architecture:

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 (see the
   considerations in [5] for LOs containing multiple locations).
   Location Objects typically include some sort of identifier
   associated with the Target.

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 which entities may obtain location
   information in what form (access control rules) and how location
   information may be used by an entity (usage rules).

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

```
+----+
| LG |
+----+
  ^
  |
Positioning
 Data
  |
  |     +------------Privacy Rules------------------>+----+
  |     |                                    +---->| LR |--> ...
  |     |                                    |     | LS |
  v     |                                    |     +----+
+-------+                                    |
|Target |                +----+              |      +----+
|Device |--------------->| LR |--------------+---->| LR |
|  RM   |       LO       | LS |       LO     |      +----+
|  LS   |                +----+              |
+-------+                                    |
                                             |      +----+
                                    +---->| LR |
                                           +----+
```

Figure 2: Basic Geopriv Scenario

4. The Location Life-Cycle

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 Device's
    location.

2. Distribution: Location Servers send location to Location
    Recipients, which may in turn act as Location Servers and further
    distribute location to other Location Recipients (possibly
    several times).

3. Use: 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 and security implications. The
Geopriv roles are mapped onto the location life-cycle in the figure
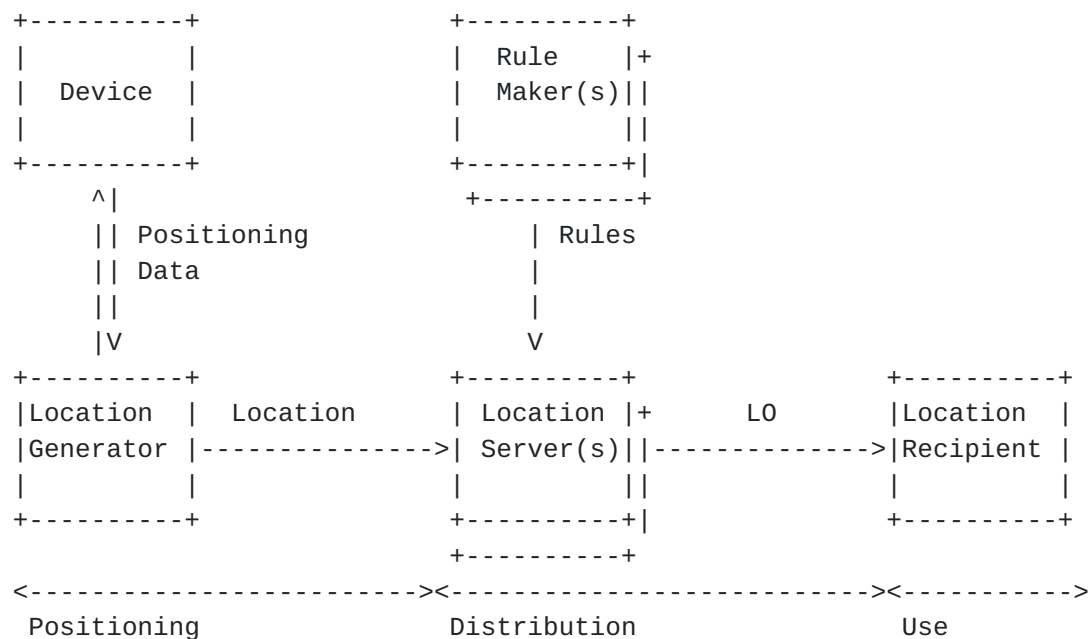below.

```
    +----------+                   +----------+
    |          |                   |  Rule    |+
    |  Device  |                   |  Maker(s)||
    |          |                   |          ||
    +----------+                   +----------+|
        ^|                     +----------+
        || Positioning         | Rules
        || Data                |
        ||                     |
        |V                     V
    +----------+          +----------+                +----------+
    |Location  |  Location| Location |+     LO        |Location  |
    |Generator |--------------->| Server(s)||-------------->|Recipient |
    |          |          |          ||                |          |
    +----------+          +----------+|                +----------+
                          +----------+
    <------------------------><----------------------------><------------>
     Positioning              Distribution                    Use
```

Figure 3: Location Life-Cycle

4.1. Positioning

Positioning is the process by which the physical location of the
Device is computed, based on some observations about the Device'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 Device, and the
outcome is that the LG knows the location of the Device.

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

4.1.1. Determination Mechanisms and Protocols

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

o Device-based

o Network-based

o Network-assisted

As suggested by the above names, a positioning scheme can rely on the
Device, 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 Device-based positioning mechanisms, the Device 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 Device acts as its own LG, and there are no protocols
required to support positioning (since no information needs to be
communicated).

In network-based positioning schemes, an external LG (an Internet
host other than the Device) has access to sufficient information
about the Device, through out-of-band channels, to establish the
position of the Device. The most common examples of this type of LG
are entities that have a physical relationship to the Device (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 Device-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 LGs to determine location, they are not
controlled as private information. Network-based positioning often
provides location to protocols by which the network informs a Device
of its own location (these are known as Location Configuration
Protocols, see Section 4.2.2 for further discussion).

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

o The Device provides measurements to the LG

o The LG provides assistance data to the Device

"Measurements" are understood to be observations about the Device's environment, ranging from wireless signal strengths to the MAC address of a first-hop router. "Assistance" is the complement to measurement, namely the positioning 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 that maps 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 Device sends measurements (e.g., MAC addresses and signal strengths) to an LG, and the LG returns a location to the client. In wired networks, the Device can send its MAC address to the LG, 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 Device provides the LG with GNSS measurements 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 Device 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 Device, the LG, or both) is aware of the Device's location at the end of the protocol interaction.

4.1.2. Privacy Considerations for Positioning

Positioning is the first point at which location may be associated with a particular Target's identity. Local identifiers, unlinked pseudonyms, or private identifiers that are not linked to the real identity of the Target should be used as forms of identity whenever possible. This provides privacy protection by disassociating the location from the Target's identity before it is distributed.

At the conclusion of the positioning process, the entity acting as

the LG has the Device's location (if the Device is performing the LG
role, then they both have it). If the entity acting as the LG also
performs the role of LS, the privacy considerations in Section 4.2.4
apply.

In some deployment scenarios, positioning functions and distribution
functions may need to be provided by separate entities, in which case
the LG and LS roles will not be performed by the same entity. In
this situation, the LG acts as a "dumb," non-privacy-aware
positioning resource, and the LS provides the privacy logic necessary
to support distribution (possibly with multiple LSes using the same
LG). In order to allow the privacy-unaware LG to distribute location
to these LSes while maintaining privacy, the relationship between the
LG and its set of LSes MUST be tightly constrained, effectively
"hard-wired." That is, the LG MUST only provide location to a small
fixed set of LSes, and each of these LSes MUST comply with the
requirements of Section 4.2.4.

4.1.3. Security Considerations for Positioning

Manipulation of the positioning process can expose location through
two mechanisms:

1) A third party could guess or derive measurements about a specific
device and use them to get the location of that Device. To mitigate
this risk, the LG SHOULD be able to authenticate and authorize
devices providing measurements and, if possible, verify that the
presented measurements are likely to be the actual physical values
measured by that client. These security procedures rely on the type
of positioning being done, and may not be technically feasible in all
cases.

2) By eavesdropping, a third party may be able to obtain measurements
sent by the Device itself that indicate the rough position of the
Device. To mitigate this risk, protocols used for positioning MUST
provide confidentiality and integrity protections in order to prevent
observation and modification of transmitted positioning data while en
route between the Target and the LG.

If an LG or a Target chooses to act as an LS, it inherits the
security requirements for an LS, described in Section 4.2.5.

4.2. Location Distribution

When an entity receives location (from an LG or an LS) and
redistributes it to other entities, it acts as an LS. Location
Distribution is the process by which one or more LSes provide LOs to
LRs in a privacy-preserving manner.

The role of an LS is thus two-fold: First, it must collect location information and Rules that control access to that information. Rules can be communicated within an LO, within a protocol that carries LOs, or through a separate protocol that carries Rules. Second, the LS 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.

An LS thus has at least two types of interactions with other hosts, namely receiving and sending LOs. An LS may optionally implement a third interaction, allowing Rule Makers to provision it with Rules. 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 directly from a Rule Maker has such a relationship, while an LS that acquires all its Rules through LOs does not.

4.2.1. Privacy Rules

Privacy Rules are the central mechanism in Geopriv for maintaining a Target's privacy, because they provide a recipient of an LO (an LS or LR) with information on how the LO may be used.

Throughout the Geopriv architecture, Privacy Rules are communicated in rules languages with a defined syntax and semantics. For example, the Common Policy rules language has been defined [6] to provide a framework for broad-based rule specifications. Geopriv Policy [7] defines a language for creating location-specific rules. XCAP [8] can be used as a 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 can only augment privacy protections because all Rules are positive grants of permission.

The following are examples of Privacy Rules governing location distribution:

o Retransmit location when requested from example.com

o Retransmit only city and country

o Retransmit location with no less than a 100 meter radius of uncertainty

o Retransmit location only for the next two weeks

LSes enforce Privacy Rules in two ways: by denying requests for location, or by transforming the location information before retransmitting it.

LSes may also receive Rules governing location retention, such as "Retain location only for 48 hours." Such Rules are simply directives about how long the Target's location information can be retained.

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

Forwarded Rules, on the other hand, travel inside LOs and direct LSes and LRs 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 an LS:

o The LS may retransmit Alice's precise location to Bob, who in turn is permitted to retain the location information for one month

o 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

o The LS may retransmit Alice's country to a photo-sharing website, which in turn is permitted to retain the location information for one year and retransmit it to any requesters

When Steve asks for Alice's location, the LS 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"). A LO may contain a pointer to more
robust Rules, such as those shown in the set of four Rules at the
beginning of this section.

4.2.2. Location Configuration

Some entities performing the LG role are designed only to provide
Targets with their own locations (as opposed to distributing a
Target's location to others). The process of providing a Target with
its own location is known within Geopriv as Location Configuration.
The term Location Information Server (LIS) is often used to describe
the entity that performs this function (although a LIS may also
perform other functions, such as providing a Target's location to
other entities).

A Location Configuration Protocol (LCP) [9] is one mechanism that can
be used by a Device to discover its own location from a LIS. LCPs
provide functions in the way they obtain, transport and deliver
location requests and responses between a LIS and a Device such that
the LIS can trust that the location requests and responses handled
via the LCP are in fact from/to the Target. Several LCPs have been
developed within Geopriv [10][11][12][13].

A LIS whose sole purpose is to perform Location Configuration need
only follow a simple privacy-preserving policy: transmit a Target's
location only to the Target itself. This is known as the "LCP
policy."

Importantly, if an LS is also serving in the role of LG and it has
not been provisioned with Privacy Rules for a particular Target, it
MUST follow the LCP policy, whether it is a LIS or not. In the
positioning phase, an entity serving the roles of both LG and LS that
has not received Privacy Rules must follow this policy. The same is
true for any LS in the distribution phase.

4.2.3. Location References

The location distribution process occurs through a series of
transmissions of LOs: 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," where a reference is represented by a URI that can be dereferenced to obtain the LO. This document summarizes the properties of location-by-reference that are discussed at length in [14].

Distribution of location by reference (distribution of location URIs) offer several benefits. 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").

From a positioning perspective, location by reference can offer the additional benefit of "just in time" positioning. If location is distributed by reference, an entity acting as a combined LG/LS only needs to perform positioning operations when a recipient dereferences a previously distributed URI.

From a privacy perspective, distributing location as a URI instead of as an LO 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. But 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 it enforces the "possession model" defined in [14]), then transmitting a location URI presents the same privacy risks as transmitting the LO itself. Moreover, the use of location URIs without access controls can introduce additional privacy risks: If URIs predictable, 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 need to be constructed so that they contain a random component with sufficient entropy to make guessing infeasible.

4.2.4. Privacy Considerations for Distribution

Location information MUST be accompanied by Rules throughout the distribution process. Otherwise, a recipient will not know what uses are authorized, and will not be able to use the LO. Consequently, LOs MUST be able to express Rules that convey appropriate authorizations.

An LS MUST only accept Rules from authorized Rule Makers. For an LS that receives Rules exclusively in LOs and has no direct relationship with a Rule Maker, this requirement is met by applying the Rules provided in an LO to the distribution of that LO. For an LS with a direct relationship to a Rule Maker, 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 provide 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.

An LS that receives Rules exclusively through LOs MUST examine the
Rules that accompany a given LO in order to determine how the LS may
use the LO (if any Rules are included by reference, the LS SHOULD
attempt to download them). If the LO includes no Rules that allow
the LS to transmit the LO to another entity, then the LS MUST NOT
transmit the LO. If the LO contains no Rules at all (if it is in a
format with no Rules syntax, for example), then the LS MUST delete it
(emergency services provide an exception in that Rules can be implicit,
see [15]). If the LO included Rules by reference, but these Rules
were not obtained for any reason, the LS MUST NOT transmit the LO and
MUST delete it.

An LS that receives Rules both directly from one or more Rule Makers
and through LOs MUST combine the Rules in a given LO with Rules it
has received from the RMs. The strategy the LS 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 an LO-provided rule (regardless of the RM Rules).

Different policies may be applicable in different scenarios. In
cases where an external RM is more trusted than the source of the LO,
the "RM Override" policy may be suitable (for example, if the
external RM is the Target, and the LO is provided by a third party).
Conversely, the "LO Override" policy is better suited to cases where
the LO provider is more trused than the RM (for example, if the RM is

the user of a mobile device LS and the LO contains Rules from the
RM's parents or corporate security office). The "Intersection"
policy takes the strictest view of the permission grants, giving
equal weight to all RMs (including the LO creator).

Each of these policies will also have different privacy consequences.
Following the "Intersection" policy ensures that the most privacy-
protective subset of all RMs' rules will be followed. The "Union"
policy and both "Override" policies may defy the expectations of any
RM (including, potentially, the Target) whose policy is not followed.
For example, if a Target acting as an RM sets Rules and those Rules
are overridden by the application of a more permissive LO Override
policy that has been set by the Target's parent or employer acting as
an RM, the retransmission or retention of the Target's data may come
as a surprise to the Target. For this reason, it is RECOMMENDED that
LSes provide a way for RMs to be able to find out which policy will
be applied to the distribution of a given LO.

4.2.5. Security Considerations for Distribution

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, LSes, and RMs) and other information in the
request.

Protocols to convey LOs and protocols to convey Rules MUST provide
information on the identity of the recipient of location 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, LOs 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 an LO as it travels between
privacy-preserving entities.

Ideally, LOs SHOULD be transmitted with confidentiality protection
end-to-end between an LS that transmits location and the LR that
receives it. In some cases, the protocol conveying an LO provides
confidentiality protection as a built-in security solution for its

signaling (and potentially its data traffic). In this case, carrying
an unprotected LOs within such an encrypted channel is sufficient.
Many protocols, however, are offering communication modes where
messages are either unprotected or protected on a hop-by-hop basis
(for example, between intermediaries in a store-and-forward
protocol). In such a case it is RECOMMENDED that the protocol allows
for the use of encrypted LOs, or for the transmission of a reference
to location in place of an LO [14].

4.3. Location Use

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 an LO in ways that have minimal 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 complete its
interaction with the LS -- then no further action is necessary for
the LR to comply with Geopriv requirements.

As an example of this simplest case, if an LR (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 an 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.

4.3.1. Privacy Considerations for Use

The principal privacy requirement for LRs is to follow usage rules.
Any LR that wants to retransmit or retain the LO 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.

4.3.2. Security Considerations for Use

Since the LR role does not involve transmission of location, there
are no protocol security considerations required to support privacy
(other than ensuring that data does not leak unintentionally caused

by security breaches).

Aside from privacy, LRs often require some assurance that an 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.

5. Security Considerations

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

The life-cycle of an LO often consists of a series of location transmissions. Protocols that carry location can provide strong assurances, but only for a single segment of the LO'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 [16] or TLS [17].

Additionally, if (1) the protocol provides mutual authentication for every segment, and (2) every entity in the location distribution chain exchanges information only with entities with whom it has a trust relationship, entities can transitively obtain assurances regarding the origin and ultimate destination of the LO. 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 receives 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 to the LS (by a previous LS, for example). Likewise, the LR can verify that the LO was transmitted by the LS, but cannot verify the origin of the LO if it did not originate with the LS.

Security mechanisms in protocols are thus unable to provide direct assurances over multiple transmissions of an LO. However, the transmission of location "by reference" can be used to effectively turn multi-hop paths into single-hop paths. If the multiple transmissions of an LO are replaced by multiple transmissions of a URI (a multi-hop dissemination channel), the LO need only traverse a

single hop, namely the dereference transaction between the LR and the dereference server. The requirements for securing location passed by reference [14] are applicable in this case.

The major threats to the security of LOs can be grouped into two categories. First, threats against the integrity and authenticity of LOs can expose entities that rely on LOs. Second, threats against the confidentiality of LOs can allow unauthorized access to location information.

An LO 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, an LO encodes a set of bindings among them. That is, the LO 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 that time for the distribution of his location. Below, we provide a description of the assurances required by each party involved in the location distribution in order to mitigate the possible attacks on these bindings.

Rule Maker: The Rule Maker is responsible for creating the Target's Privacy Rules and for uploading them to the LSes. The primary assurance required by the Rule Maker is that the Target's Privacy Rules are correctly associated with the Target's identity when they are conveyed to each LS that handles the LO. Ensuring the integrity of the Privacy Rules distributed to the LSes prevents rule-tampering attacks. 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 LS.

Location Server: The Location Server is responsible for enforcing the Target's Privacy Rules. The first assurance required by the LS is that the binding between the Target's Privacy Rules and the Target's identity is authentic. Authenticating and authorizing the Rule Maker who creates, updates and deletes the Privacy Rules prevents rule-tampering attacks. The LS has to ensure that the authorization policies are not exposed to third parties, if so desired by the Rule Maker (when the rules themselves are privacy-sensitive).

Location Recipient: The Location Recipient is the consumer of the LO. The LR 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 helps to prevent various attacks, such falsifying the location, modifying

the time-stamp, faking the identity, replaying LOss.

Location Generator: The primary assurance required by the Location
   Generator is that the LS to which the LO is initially published is
   one that is trusted to enforce the Target's Privacy Rules.
   Authenticating the trusted LS mitigates the risk of server
   impersonation attacks. Additionally, the LG is responsible for
   the location determination process, which is also sensible from a
   security perspective because wrong input provided by external
   entites can lead to undesireable disclosure or access to location
   information.

Assurances as to the integrity and confidentiality of a Location
Object can be provided directly through the LO format. RFC 4119 [18]
provides a description for usage of S/MIME to integrity and
confidentiality protection. 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 an LO is impractical. For example, in some deployment
scenarios a direct trust relationship may not exist between the
creator of the Location Object and the recipient. Additionally, in a
scenario where many recipients are authorized to receive a given LO,
the creator of the LO cannot guarantee end-to-end confidentiality
without knowing precisely which recipient will receive the LO. Many
of these cases can, however, be addressed by the usage of a Location-
by-Reference (possibly combined with an LO).

6. Example Scenarios

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 an LO carries are equivalent to those in a PIDF-LO
(namely, that the principal Rules that can be set are limits on the
retransmission and retention of the LO). 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.

6.1. Minimal Scenario

One of the simplest scenarios in the Geopriv architecture is when a
Device determines its own location and uses that LO to request a
service (e.g., by including the LO in an HTTP POST request [19] or

SIP INVITE message [20]), and the server delivers that service
immediately (e.g., in a 200 OK response in HTTP or SIP), without
retaining or retransmitting the Device's location. The Device acts
as an LG by using a Device-based positioning algorithm (e.g., manual
entry) and as an LS by interpreting the rule and transmitting the LO.
The Target acts as a Rule Maker by specifying that the location
should be sent to the server. The server acts as an LR 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, as LR, does not do anything
that creates a privacy risk -- it does 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:

o Positioning: Device-based, Device=LG

o Distribution hop 1: HTTP UA --> Ephemeral web service, privacy via
   user indication

o Use: Ephemeral web service delivers response without retaining or
   retransmitting location

o Key points:

   * LRs that do not behave in ways that risk privacy are Geopriv-
      compliant by default. No further action is necessary.

6.2. Location-based Web Services

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 Device-based positioning
(manual entry and Global Positioning System (GPS)) and network-
assisted positioning (via Assisted GPS (AGPS), and multilateration
with 802.11 and cellular signals), exposing location 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 an LO 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 [21] 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 an LG.

The distribution phase actually occurs entirely within the Target
host. This phase begins when the positioning server, now acting as
LS, follows the LCP policy by providing location only to the Target.
The next hop in distribution occurs when the browser extension (an
entity under the control of the Target) passes an LO to the web page
(an entity under the control of its author). In this phase, the
browser extension acts as an LS, with the Target as the sole Rule
Maker; the user interface for rule-making is effectively a protocol
for conveying Rules, and the extension's API effectively defines a
way to communicate LOs and an LO Format. The web site acts as an LR
when the web page accepts the LO.

The use 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 sending the LO in an XMLHttpRequest to a back-end server
is 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 (a map, for instance) 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:

o Positioning: Network-assisted, positioning server=LG

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

o Distribution hop 1: positioning server=LS --> Target, privacy via
   LCP policy

o Distribution hop 2: Browser=LS --> Web site=LR, privacy via user
   confirmation

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

o 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

6.3. Emergency Calling

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
[22][23]. 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 an LS, preserving the privacy of location information
by only including it in emergency calls.

The second step in an emergency call encompasses location
distribution and use. The entity that is routing the emergency call
sends location though the LoST protocol [15] to a mapping server. In
this role, the routing entity acts as an LS and the LoST server acts
as an LR. 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 [24]. 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 use. The caller (or another entity that inserts the
caller's location) acts as an LS and the PSAP acts as an LR. 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 an 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:

o Positioning: Any

o Distribution/use hop 1: Target=LS --> LoST infrastructure (no
   Rules), privacy via authorization implicit in protocol

o Distribution/use hop 2: Target=LS --> PSAP, privacy via Rules in
   LO

o Use: PSAP uses location to deliver emergency services

o 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

6.4. Combination of Services

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 LSess. 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 [25], 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 allow her location to be shared with those services, and allow
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 an 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 to 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 use.
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
Alice has delegated privacy responsibilities).

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 an LR. 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 acts as an LS, 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 an LO can specify
different Rule sets for different LSes.

Eve is the fourth LS in the chain, and fails to comply with Geopriv
by not checking the Rules in the LO prior to uploading the LO 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:

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

o Distribution/use hop 1: Presence server --> Bob, privacy via
   Alice's access control rules

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

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

o Distribution/use hop 4: Eve --> Social networking site, violates
   privacy by retransmitting

o Use: Social networking site, privacy via checking Rules and
   discarding

o 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

7. **Glossary**

Various security-related terms not defined here are to be understood
in the sense defined in RFC 4949 [26].

$ Access Control Rule

   A rule that describe which entities may receive location
   information and in what form.

$ 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.

$ Device

   The physical device whose location is tracked as a proxy for the
   location of a Target.

$ geodetic location

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

$ 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)

   Performs the role of initially determining or gathering the
   location of a Target. Location Generators may be any sort of
   software or hardware used to obtain a Target's location (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 sought in the
   Geopriv architecture. In many cases the Target will be the human
   user of a Device, or it may be an object such as a vehicle or
   shipping container to which a Device is attached. In some
   instances the Target will be the Device itself. The Target is the
   entity whose privacy Geopriv seeks to protect.

$ Usage Rule

   A rule that describe what uses of location information are
   authorized.

8. Acknowledgements

   Section 5 is largely based on the security investigations conducted
   as part of the Geopriv Layer-7 Location Configuration Protocol design
   team, which produced [9]. We would like to thank all the members of
   the design team.

   We would also like to thank Marc Linsner and Martin Thomson for their
   contributions regarding terminology and LCPs.

9. IANA Considerations

   This document makes no request of IANA.

10. References

10.1. Normative References

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

10.2. Informative References

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

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

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

   [5]   Winterbottom, J., Thomson, M., and H. Tschofenig, "GEOPRIV
         Presence Information Data Format Location Object (PIDF-LO)
         Usage Clarification, Considerations, and Recommendations",
         RFC 5491, March 2009.

   [6]   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.

[7]   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.

[8]   Rosenberg, J., "The Extensible Markup Language (XML)
       Configuration Access Protocol (XCAP)", RFC 4825, May 2007.

[9]   Tschofenig, H. and H. Schulzrinne, "GEOPRIV Layer 7 Location
       Configuration Protocol: Problem Statement and Requirements",
       RFC 5687, March 2010.

[10]   Polk, J., Schnizlein, J., and M. Linsner, "Dynamic Host
       Configuration Protocol Option for Coordinate-based Location
       Configuration Information", RFC 3825, July 2004.

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

[12]   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-08 (work in progress),
       July 2010.

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

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

[15]   Hardie, T., Newton, A., Schulzrinne, H., and H. Tschofenig,
       "LoST: A Location-to-Service Translation Protocol", RFC 5222,
       August 2008.

[16]   Kent, S. and K. Seo, "Security Architecture for the Internet
       Protocol", RFC 4301, December 2005.

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

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

[19]   Fielding, R., Gettys, J., Mogul, J., Frystyk, H., Masinter, L.,
       Leach, P., and T. Berners-Lee, "Hypertext Transfer Protocol --
       HTTP/1.1", RFC 2616, June 1999.

   [20]   Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A.,
          Peterson, J., Sparks, R., Handley, M., and E. Schooler, "SIP:
          Session Initiation Protocol", RFC 3261, June 2002.

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

   [22]   Rosen, B., Schulzrinne, H., Polk, J., and A. Newton, "Framework
          for Emergency Calling using Internet Multimedia",
          draft-ietf-ecrit-framework-11 (work in progress), July 2010.

   [23]   Rosen, B. and J. Polk, "Best Current Practice for
          Communications Services in support of Emergency Calling",
          draft-ietf-ecrit-phonebcp-15 (work in progress), July 2010.

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

   [25]   Saint-Andre, P., Ed., "Extensible Messaging and Presence
          Protocol (XMPP): Core", RFC 3920, October 2004.

   [26]   Shirey, R., "Internet Security Glossary, Version 2", RFC 4949,
          August 2007.

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

URIs

   [28] <http://creativecommons.org/>

Authors' Addresses

   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