

GEOPRIV
Internet-Draft
Updates: [3693](#), [3694](#)
(if approved)
Intended status: Informational
Expires: August 28, 2008

R. Barnes
M. Lepinski
BBN Technologies
H. Tschofenig
Nokia Siemens Networks
H. Schulzrinne
Columbia University
February 25, 2008

Security Requirements for the Geopriv Location System
draft-barnes-geopriv-lo-sec-02

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on August 28, 2008.

Copyright Notice

Copyright (C) The IETF Trust (2008).

Abstract

Internet protocols that deal with presence-based location objects support a wide variety of applications. However, the dissemination of location objects from sources of location to consumers is a common feature of all location-based applications. In order to enable the

Internet-Draft

Location Object Security

February 2008

development of broadly-applicable security and privacy mechanisms for dissemination of location objects, this document describes an end-to-end architecture for policy-constrained location distribution. In this architecture, location distribution is accomplished by a set of distributed actors. We describe the assurances that these actors require from the architecture, and derive more a more detailed description of the security features required to provide those assurances.

Internet-Draft

Location Object Security

February 2008

Table of Contents

1.	Introduction	4
2.	Terminology	5
3.	An End-to-end Location Architecture	6
3.1.	Structure of a Location Transmission	6
3.1.1.	Structure of a Location Request	8
3.1.2.	Location References	9
3.1.3.	LS Processing of Location Requests	10
3.2.	The End-to-end Model and Global Roles	11
3.3.	Usage Scenarios for this Model	13
3.3.1.	RFC 3693 model of transmission	13
3.3.2.	Location Configuration	14
3.3.3.	Location Conveyance by Value	14
3.3.4.	Location Conveyance by Reference	14
3.3.5.	Presence Server	14
4.	Required Assurances	15
4.1.	Location Transmission	15
4.1.1.	Rule Maker	15
4.1.2.	Location Server	16
4.1.3.	Location Recipient	16
4.2.	End-to-end distribution	16
4.2.1.	Location Generator	16
4.2.2.	Viewer	17
4.2.3.	Target	17
4.3.	Summary of Required Assurances	18
5.	Security Requirements	18
5.1.	Unauthorized Modification of Rules	19
5.2.	Unauthorized Exposure of Rules	20
5.3.	Acceptance of Rules from Unauthorized Rule Makers	20
5.4.	Unauthorized Exposure of Location Objects	20
5.5.	Unauthorized Modification of Location Objects	23
5.6.	Assertion of Location Object Origins	24
5.7.	Summary of Security Requirements	24
6.	Security Considerations	26
7.	Acknowledgements	26

8.	IANA Considerations	26
9.	References	26
9.1.	Normative References	26
9.2.	Informative References	26
	Authors' Addresses	27
	Intellectual Property and Copyright Statements	29

[1.](#) Introduction

Demand for location-based Internet applications, especially location-based Internet calling [[I-D.ietf-ecrit-framework](#)], has driven the creation of Internet protocols for communicating information about the location of Internet end hosts or other entities. Of interest, for example, are protocols for informing hosts of their own location (location configuration protocols), transmitting location information from one host to another (location conveyance protocols), and requesting location information from a server (location dereference protocols).

The first goal of this document is to describe how location information is used by these protocols over its entire "life-cycle". This life-cycle begins when location information is introduced into an IP network via a location configuration protocol, continues through one or more transmissions by way of location conveyance and dereference protocols, and ultimately ends when the location is delivered to an application consumer.

The Location Objects (LO) described in [RFC 3693](#) and [RFC 3694](#) are usually encoded as XML documents in the Presence Information Data Format - Location Object (PIDF-LO) schema [[RFC4119](#)]. While the general trend in the IETF is to require that LOs be in this format, certain protocols do not use PIDF-LO, most notably the DHCP extensions to carry location in civic [[RFC4776](#)] or geospatial [[RFC3825](#)] format. In this document, such formats for location information are also regarded as LO formats, even though they do not comply with the requirements for LO formats in [RFC 3693](#).

The expansion of scope to include location object formats other than those in compliance [RFC 3693](#) is not meant to in any way deprecate or supercede the requirements of [RFC 3693](#). This document is intended to treat security aspects of location communication independent of the other considerations that [RFC 3693](#) addresses. Where the two documents overlap, we aim to provide greater specificity in guidance and requirements.

A model for the use of Internet protocols to transmit location information via a store-and-forward network of Location Servers has been described in [RFC 3693](#) [[RFC3693](#)]. Privacy concerns and privacy-relevant security concerns are described in [RFC 3694](#) [[RFC3694](#)]. This document extends those documents in three ways: First, we explicitly take into account end-to-end properties of the system, through multiple location transmissions. Second, we address security concerns not directly related to the privacy of location information (of concern for Viewers), such as location integrity and access control (of concern to Location Generators). Third, and most

importantly, we extend these considerations beyond a presence-based model to create a general framework for policy-based dissemination of location objects.

Similarly, several policy languages have been developed in the context of presence authorization (and for location within that context). [RFC 4745](#) [[RFC4745](#)] defines a general framework for expressing privacy policies, and [RFC 5025](#) [[RFC5025](#)] specializes this framework to the case of presence documents (of which PIDF-LO location objects are considered a subset). This document considers these sorts of authorization rules in the context of a broader location request authorization framework.

The remainder of this document is structured as follows: After relevant terminology is introduced in [Section 2](#), [Section 3](#) describes an architecture for the end-to-end distribution of location over the Internet. In particular, this architecture describes a set of entities that work together to move location information from source to consumer. Based on the roles they play in the architecture, these entities may require certain assurances, and these are described in [Section 4](#). Finally, in [Section 5](#), the technical properties and mechanisms required to enable these assurances are reflected in a set

of requirements for Geopriv security mechanisms.

2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [[RFC2119](#)].

The focus of this document is the security properties of two types of protocols and two types of data formats:

- o Policy Conveyance Protocols communicate policy information between rule makers and location servers. These can be dedicated protocols (e.g., XCAP [[RFC4825](#)]), or, when rules are carried within a location object, the location conveyance protocol can act as a policy conveyance protocol.
- o Location Conveyance Protocols communicate location requests and responses between the location server and the location recipient (e.g., SIP Geolocation, HELD, and Location Dereference Protocols). Location configuration protocols [[I-D.ietf-geopriv-l7-lcp-ps](#)] and location dereference protocols [[I-D.ietf-geopriv-lbyr-requirements](#)] are special cases of location conveyance protocols.

- o Location Object Formats define how location information and ancillary data are encoded; information is passed between distant points in the distribution chain by being carried in the LO.
- o Location Reference Formats define how location references (i.e., request parameters) are encoded for dissemination from an LS to LRs.

The roles played by these protocols are described in [Section 3.1](#), and corresponding security requirements are described in [Section 5](#).

This document re-defines the following terms from [RFC 3693](#) in an effort to refine their scope: Rule Maker, Location Server, Location Recipient, Location Generator, Viewer. Full definitions are given in [Section 3.1](#) and [Section 3.2](#).

3. An End-to-end Location Architecture

In this section, we present an architecture for the end-to-end communication of location information. The overall pattern of transmissions involved in this communication is often complex and thus such systems are modeled as the composition of atomic building blocks.

In [Section 3.1](#) we describe a single location transmission, and the roles played by parties in such a transmission. A location transmission is an atomic unit that models a single movement of location information. In [Section 3.2](#) we describe how multiple location transmissions can fit together within an end-to-end system and the global roles played by entities in such a composite system. Finally, in [Section 3.3](#) we demonstrate how this model maps to common location use-cases such as location configuration and point-to-point location conveyance.

3.1. Structure of a Location Transmission

Location transmission is the basic building block for policy-constrained location distribution. The model we describe here for a location transmission is based on the one described for a presence server in [RFC 4745](#). The protocol interactions involved in a location transmission are illustrated in Figure 1:

1. A Rule Maker informs the Location Server about Privacy Rules governing the distribution of Location Objects.
2. In some cases, the LR will acquire a location reference (e.g., a URI or a domain name for the LS) through an external

dissemination channel; a specification of this channel is outside the scope of this document.

3. The transmission is initiated either when the LR sends a request to the LS, or when the LS is directed to transmit location by some other mechanism. (These two cases roughly correspond to Passive and Active Request-Response modes of [RFC 4745](#), respectively.)

4. The LS determines whether the transmission is permitted by currently available policy, and if so, transmits location to the LR. Note that in addition to rules installed by the RM, the LS also uses policies contained in the LO itself and policies defined by local configuration.

The policy transaction in step (1) is conducted using a policy conveyance protocol. The reference communicated in step (2) is communicated through an unspecified dissemination channel in a given location reference format. The transmission in step (4) is conducted using a location conveyance protocol, and when the transmission is initiated by the LR, the request uses the location conveyance protocol as well. The LO is transmitted in some location object format.

This model makes two important simplifying assumptions. First, multiple asynchronous responses to a single request are considered part of the same transmission. That is, we do not distinguish between the Passive Request-Response and Asynchronous modes of [RFC 4745](#). Second, multiple LOs contained within a single response are considered as a single response. (A response containing multiple LOs is authorized if and only if all of the LOs in the response would be authorized independently.)




Figure 1: A single location transmission

There are three roles involved in this transaction, a Location Server

(LS), a Location Recipient (LR), and a Rule Maker (RM). A single entity may play multiple of these roles within a single transmission (see Section [Section 3.3](#) for examples). The only two roles that are necessarily separate are that of the LS and the LR.

Rule Maker The Rule Maker is the party who produces the rules governing whether a Location Recipient is allowed to receive location information and what precision of location information a Location Recipient is allowed to receive. (Formats for these rules are described in [[RFC4745](#)] and [[I-D.ietf-geopriv-policy](#)].) The Rule Maker may send rules directly to the Location Server, or the Location Server may receive the rules as part of a location object as per [[RFC4119](#)]. Note that some transmissions may occur without a Rule Maker, in which cases the transmission is constrained only by policy contained in the LO itself and LS-internal policy.

Location Server The Location Server is the party who possesses the location information at the beginning of the transmission. The Location Server receives rules governing the location information as received from the Rule Maker, as part of the location object containing the location information, or as part of its internal configuration. The Location Server is responsible for applying these rules and as such he may need to reduce the precision of the location information or terminate the location transmission if the Location Recipient is not authorized to receive the location information. After applying the appropriate rules, the Location Server sends the location information to the Location Recipient.

Location Recipient The Location Recipient receives the location from the Location Server, either by making a request to the LS or as a result of an LS-initiated transmission.

The distinction between LS-initiated and LR-initiated transfers is significant, because in the latter case, the LR can influence which LO is transmitted. Additional concerns related to the dissemination of references and the interaction between requests and policy make the LS policy decision process considerably more complex when transmissions are initiated by the LR. Thus, we treat that case in more detail in the below.

[3.1.1](#). Structure of a Location Request

Logically, a location request is a message sent from an LR to an LS that requests that the LS send an LO (or set of LOs) to the LR. This means that the request must contain at least two types of data:

1. A description of the LO to be returned
2. An identifier for the LR to which it should be delivered

Depending on the individual protocol and the individual request, the internal structure of these data can vary. For example, the identifier for the LR can be a source IP address or a SIP URI. The description of the LO to be returned could be a detailed set of parameters, or an opaque identifier; it could even be implicit, being inferred from the LR's identity. In general, we consider the identifier for the LR as a single datum, while the description of the LO is considered as logically consisting of a set of parameters, e.g:

- o Identity of the target
- o Time of sighting / timestamp
- o Format of desired LO
- o Positioning mechanism used in sighting

The LS may accept these parameters in "clear" or "opaque" form, i.e., in the form that can be readily matched against authorization rules or in the form of a random token that maps to a clear value in a way known only to the LS). In order to be considered "opaque", the values assigned by the LS MUST have sufficient entropy that they are difficult to guess without prior knowledge. Note also that the LS may choose to map a single opaque token to a collection of clear values.

Implicit in the representation of parameter values by opaque tokens is that these tokens have a lifetime, namely, the period of time for which the LS retains a mapping between the opaque token and one or more clear parameter values.

[3.1.2.](#) Location References

A location reference is a data structure that provides information on how to make a request for location. In order to be useful at all, a reference must contain contact information (e.g., a domain name) for an LS. Additionally, the reference may contain parameter values that describe an LO. The request that is generated from a reference has the indicated parameter values filled into appropriate fields, and is sent to the indicated LS.

References are the mechanism whereby values for opaque parameters are

distributed. An LS constructs a reference containing opaque values which is then distributed to LRs through some dissemination channel.

Every reference that conveys opaque parameter values has a validity lifetime, which is the intersection of the validity intervals of the opaque values it conveys.

To say this another way: Suppose that whenever an LS creates a reference it creates a new set of values for all opaque parameters (or, equivalently, creates a single opaque token that maps to a set of clear values), all with the same validity interval. Then the reference is valid over the same interval as the opaque tokens, and the LS can render the reference unusable by deleting the associated mapping(s).

More concretely, location references are often encoded as URIs. For example, if there were an HTTP request protocol defined, the URI `<http://ls.example.net/134245>` would indicate that an HTTP request should be sent to `ls.example.net`, with the value `"/134245"` (or `"http://lis.example.net/134245"`) as the Request URI (and in other fields as specified by the protocol). The validity lifetime of this URI is the lifetime for which the LS will store a mapping between the opaque value `"134245"` and a set of clear parameters.

A location reference logically refers to a set of LOs, namely the set of LOs that the indicated LS will return to authorized requestors in response to requests with the indicated parameter values. When the reference does not specify all available parameters, this set contains LOs for all possible parameter values. Even when all parameters are set, the set of referenced LOs contains all values that are returned over time.

The size of the referenced LO set determines the sensitivity of the reference. A reference that refers to a single LO can only expose that LO; i.e., its sensitivity is at most the sensitivity of the referenced LO (less if the LS applies access control). On the other hand, a reference that can be used to obtain a large set of locations can allow the holder of the reference track a target over time or to gather the LOs for many targets.

[3.1.3](#). LS Processing of Location Requests

An LS determines whether to return an L0 in response to a request, and which L0 to return, based on three types of policy:

1. A policy specifying which parameters are accepted in clear form (and how these should be formatted) and which are accepted in opaque form (these sets need not be disjoint). (The LS also maintains list of mappings of opaque tokens to clear values, which acts as a validation of opaque tokens.)

Barnes, et al.

Expires August 28, 2008

[Page 10]

Internet-Draft

Location Object Security

February 2008

2. A set of authorization rules of the form specified in [RFC 4745](#).
3. A decision function for choosing which among multiple L0s to return.

The second of these three, can be populated from any of three sources: (1) Rule Makers, (2) Received L0s, and (3) internal configuration. The first and last are internal policies of the LS. When the LS receives a request, it applies these policies in the same order they are presented above:

1. The LS verifies that clear parameters are properly formatted and that the values of opaque parameters are known tokens (i.e., tokens with currently valid mappings to clear values). Valid opaque parameters are translated into clear values.
2. The LS applies authorization rules to information provided in the request to determine the set of L0s that it is authorized to return. (Note: this set may not be explicitly enumerated, but rather expressed as a set of criteria.)
3. If any of the authorized L0s are compliant with the request, then the LS applies the decision function to decide which L0(s) to return to the LR.

[3.2](#). The End-to-end Model and Global Roles

The life-cycle of a Location Object typically consists of multiple location transmissions. For example, location might first be acquired via a location configuration protocol and then conveyed via a location conveyance protocol. This end-to-end distribution process can be described as a "chaining together" of the individual

transmissions described above; different transmissions are connected by an entity that acts as an LR in one transmission and an LS in the next. This process is illustrated in Figure 2. Note that although Figure 2 depicts a single "path", a single LS may transmit location to multiple LRs over time; grouping these paths together forms a logical distribution tree, with the LG as the root node and Viewers as leaf nodes.




Figure 2: End-to-end location distribution

In addition to the roles within a particular location transmission, there are also three additional global roles within the larger composite system. As described in [Section 3](#), a given party may need particular assurances based on the global role that it plays.

Location Generator The Location Generator is the party that initially introduces location information into the Internet. The LG may be (but need not be) the entity that performs the sighting of the Target. The LG may be the same as the target when mechanisms such as GPS are used, but in many settings the location generator is a separate entity.

Viewer The Viewer is the party that ultimately makes use of the location information; in particular, the Viewer does not transmit

location further. The Viewer is the Location Recipient in the final location transmission.

Target The Target is the party whose location is described by the transmitted L0. Although the Target does not explicitly play a role in the model above, every L0 has a Target, and the Target can participate in the distribution process by playing other roles.

It is common for a party to play different roles within the different transmissions. For example, the Target might be the Location Recipient during location configuration, then act as Location Server when transmitting the L0 to a presence server, then act as a Rule Maker by providing the presence server rules for further dissemination of the L0. In some cases, the Target may be a Location Generator or a Viewer; obviously, we assume that the roles of the LG and the Viewer are played by different entities.

It is assumed that the only information passed from one transmission to another is the L0 itself, so that information that is communicated across multiple hops is encoded in the L0. In particular, mechanisms for providing security across multiple location transmissions must define a new L0 format, e.g., a PIDF-L0 document encapsulated with the Cryptographic Message Syntax instead of an unencapsulated

PIDF-L0.

[3.3](#). Usage Scenarios for this Model

In order to make the meaning of the above model clearer, this section describes how several common use cases can be described using the model. In addition, we describe how the transmission model described in [RFC 3693](#) maps into the model described above.

[3.3.1](#). [RFC 3693](#) model of transmission

[Section 4 of RFC 3693](#) depicts the relationships of the primary Geopriv entities, in which the Location Server acts as a relay between a Location Generator and a Location Recipient, with rules provided by a Rule Holder. In this document, we take a more limited view of three of these roles: Here an Location Server is simply an entity that transmits location (relying on a separate associated Location Recipient for input), a Location Generator is simply a

server is the Location Server and has been internally configured with policy (there is no independent Rule Maker in this scenario). The LCP server applies the rules and returns a location object to the endpoint.

[3.3.3.](#) Location Conveyance by Value

A protocol, such as SIP, that conveys a location object by value can be modeled as a location transmission as follows. The calling device is the Location Server, and initiates the transmission. The calling device possesses a location object that contains the rules governing the location information. The called device is the Location Recipient. The calling device applies the rules within the location object and then sends the location object to the called device (e.g. in the body of a SIP INVITE).

[3.3.4.](#) Location Conveyance by Reference

A protocol, such as SIP, that conveys a location object by reference can be modeled as a location transmission as follows. In this case, SIP is the dissemination channel, over which a URI pointing to location is conveyed. The calling device sends the location reference, containing both the identity of an LS and the identity of the location information, to the called party (e.g. in the header of a SIP INVITE). The called party is the Location Recipient. Upon receiving the location reference, the called party sends a dereference request, containing a description of the desired LO, to the IS. The LS is naturally the Location Server, and has been previously provisioned with rules by the Rule Maker (likely the Target). The LS applies the appropriate rules for the location information and returns a location object to the called party.

[3.3.5.](#) Presence Server

The subscription to location information on a presence server can be modeled as a location transmission as follows. The presence watcher is the Location Recipient, and initiates the transmission. Through

some dissemination mechanism (e.g. a business card) the watcher learns the identity of a presence server and the identity of a target whose presence is stored on the server. The presence subscriber sends a subscription request, containing the identity of the target

(which identifies the desired LO), to the presence server. The presence server has previously been provisioned with rules by the Rule Maker (in this case, most likely the target). The presence server applies the rules and constructs a location object which it then sends to the presence subscriber.

[4.](#) Required Assurances

Each of the entities in the above model has expectations about how the system works, which may or may not be valid in a given situation. Depending on the needs of the entities, they may require assurances that their expectations are valid in a given situation. The goal of Geopriv security and privacy mechanisms is to provide such assurances. In order to determine requirements for Geopriv security mechanisms, then, we need to understand the assurances required by participants in the architecture.

[4.1.](#) Location Transmission

As described above, there are generally three logical roles in a single location transmission. In some cases, the same entity may play multiple roles within a transmission. In that case, the set of assurances required by that entity is the union of the assurances required by the roles it fulfills.

[4.1.1.](#) Rule Maker

The goal of the Rule Maker is provide the LS with policies to apply to transmissions, and to ensure that the rules clearly specify how the LS should execute them. The first assurance of relevance to the RM is to assure that rules are faithfully transmitted to the correct destination: Since policy documents can themselves be sensitive, the RM must verify that they are delivered only to the LS it intends, i.e., it must authenticate the identity of the LS and verify that the authenticated identity belongs to the desired LS. And since changes to a policy document can affect many subsequent transmissions, the RM requires assurance that rules are not modified en route to the LS. Second, in order to assure that policy is correctly executed, the transmitted rules must define an unambiguous mapping from requests to allowable LOs. The RM is further assured that the rules will be executed when the LS can provide confirmation that it is able to process the supplied rules at the time that the RM transmits them.

[4.1.2.](#) Location Server

The goal of the Location Server is to transmit location in compliance with relevant policy. Thus, the primary assurances the LS requires are related to the question of whether a given transmission is authorized by policy. The LS must determine whether the LR is authorized to receive location. As a pre-requisite, the LS must also verify that policy is valid, i.e., that the Rule Maker is authorized to dictate policy. All of these authorization decisions require that the server authenticate the identities of the parties requesting access, the identity of the LR requesting location and the RM requesting modifications to policy. Note that the manner in which these authorization policies are installed on the LS and applied to specific transmissions is a matter of local configuration.

[4.1.3.](#) Location Recipient

The goal of the Location Recipient is to acquire the desired L0. In general, this assurance can be decomposed into assuring that the LS is the one intended to deliver the L0 and that the L0 is faithfully transmitted from the LS to the LR; the LS is trusted by the LR to deliver the proper L0. When the transmission is initiated by the LS, the LR may not have information about the LS or the L0 prior to the transmission, so it must also trust that the LS delivering location was properly instructed to do so. When the LR initiates a transmission, the LR knows the identity of the LS and relevant properties of the L0, so the LR can be assured that the LS from which it receives location is the proper one by authenticating the identity of the LS. In both cases, the LR requires assurance that the L0 is not modified while en route between the LS and the LR.

[4.2.](#) End-to-end distribution

In addition to the three transmission entities described above, we also consider three distinguished entities in an end-to-end distribution scenario. They require assurances about the entire distribution chain, or the entire distribution tree.

[4.2.1.](#) Location Generator

The Location Generator is the Location Server at the root of the distribution tree. The LG thus offers the valuable service of acting as an Internet-accessible source of location information, and its primary interest is in controlling the use of this service, especially controlling access to it. In terms of the model, the LG is interested in controlling the set of Viewers that are able to interpret and use the L0. There are two basic approaches to

achieving this control: First, the LG may distribute LOs that are

encrypted in such a way that only the Viewers that are authorized to access the location encoded in the LO. Second, the LG may distribute location references (i.e., it may support a dissemination channel), and only provide an LO in response to a dereference query by an authorized LR. (Because these references are not valuable by themselves, the LG can allow them to be distributed by parties that may not be authorized to access the location they refer to.) In order for the Viewer to obtain the referenced location, it has to engage in a transmission in which the LG is the LS; as part of this transmission, the LG can authenticate the LR and verify directly that the Viewer is authorized to receive the location.

[4.2.2.](#) Viewer

The Viewer is the ultimate consumer of a Location Object. As a consumer, the Viewer requires assurance that the LO it receives is correct. In most situations, it is not possible to verify the correctness of location directly. Rather, a Viewer can receive assurance that a location is correct by virtue of assurance as to the identity of the source of the LO (i.e., the LG that provided it), and assurance that the LO was not modified en route to the Viewer. That is, the Viewer can have more confidence in the correctness of an LO when it can verify that the location was provided by a source that it trusts to provide correct location. As with LG access control, this verification can be done either through the object itself or through the LS that provides it. If the LO itself provides a verifiable assertion as to its origin, then the Viewer receives assurance about its correctness even if it receives the LO via an untrusted channel. On the other hand, if the LS that provides the LO is trusted to provide correct location, then the Viewer receives assurance about the LO's correctness even if the ultimate origin of the LO (i.e., the LG) remains unknown to the Viewer.

[4.2.3.](#) Target

The interests of the Target are discussed at length in [RFC 3693](#) and [RFC 3694](#). The Target by itself has no technical involvement in the distribution process; in order to affect how its location is distributed, it must take on one of the roles described above. For instance, the Target will commonly act as an LS to explicitly control

how location is transmitted, or as a Rule Maker to control distribution by a third-party LS. Much like the LG, the main concern of the Target is controlling access to its location. If the Target acts as an LS, then the assurances and mechanisms available to it are essentially the same as those available to the LG.

[4.3.](#) Summary of Required Assurances

- o Rules must be protected against unauthorized modification en route.
- o Rules must be protected against exposure to unauthorized parties.
- o Location servers must accept rules only from authorized rule makers, as determined by local policy.
- o Location objects must be exposed only to location recipients authorized by the associated rules (i.e. the rules contained in the location object, obtained from an authorized rule maker, or pre-configured on the location server in accordance with local policy).
- o Location objects must be protected against unauthorized modification en route.
- o Location generators must be able to assert that they have created a particular location object.
- o Location viewers must be able to determine that a location object has passed unchanged through a chain of location servers and (intermediate) location recipients.

[5.](#) Security Requirements

In order to enable the GEOPRIV location distribution system to provide assurances discussed in [Section 4](#), the constituent protocols must make certain security mechanisms available to the parties involved. In this section, we describe which security mechanisms are

necessary to achieve each the assurances described in [Section 4.3](#), and then provide requirements for such security mechanisms. In so doing, we provide requirements for three types of protocols:

- o Policy Conveyance Protocols communicate policy information between rule makers and location servers. These can be dedicated protocols (e.g., XCAP), or, when rules are carried within a location object, the location conveyance protocol can act as a policy conveyance protocol.
- o Location Conveyance Protocols communicate location requests and responses between the location server and the location recipient (e.g., SIP Geolocation, HELD, and Location Dereference Protocols)

- o Location Object Formats define how location information and ancillary data are encoded; information is passed between distant points in the distribution chain by being carried in the LO.
- o Location Reference Formats define how location references (i.e., request parameters) are encoded for dissemination from an LS to LRs.

The term "Location Conveyance Protocol" is similar to the term "using protocol" introduced in [RFC 3693](#), and used in [RFC 4745](#), et al. The distinction between a Location Conveyance Protocol and other protocols that may incidentally carry location information (e.g., IP or TCP) is that a Location Conveyance Protocol makes a normative requirement on the LS (i.e., the party that transmits the LO) to apply policy. Note that in some cases, the LOs carried by a location conveyance protocol will themselves carry rules. When a location conveyance protocol supports the transmission of such LOs, it is also considered a policy conveyance protocol.

The requirements described below are not strict requirements: They are lists of security features that must be present in order to support a certain set of assurances. A protocol specification can be in compliance with this document either by explaining how the protocol meets the security requirements for each assurance, or by explicitly disclaiming its ability to provide assurances for which it does not fulfill the requirements.

Note also that the security features listed below need not be provided by cryptographic means in all cases. The primary example of non-cryptographic protection is the use of appropriate policy at an LS. As an additional example, protocols that are restricted to a local network (such as DHCP or LLDP-MED) may derive security properties from the physical security of the network.

[5.1.](#) Unauthorized Modification of Rules

Rules are exposed to the risk to unauthorized modification en route when they are transmitted from a rule maker to a location server. A policy conveyance protocol can protect rules from unauthorized modification in two ways. First, the policy conveyance protocol can allow rules to be transmitted within an integrity-preserving encapsulation, such as CMS or S/MIME; this includes the use of an integrity-preserving LO format. Second, the policy conveyance protocol can allow a mode of operation in which it is carried over an integrity-protected channel, such as TLS.

REQ-1 A policy conveyance protocol MUST either support the provision of rules in an integrity-preserving encapsulation, or else it must offer a mode of operation in which rules are only transmitted over an integrity-protected channel.

[5.2.](#) Unauthorized Exposure of Rules

Rules can be exposed to unauthorized parties in two ways. Either the RM transmits the rules to a party who is not authorized by the RM to act as a location server; or else an unauthorized party is able to access the rules en route from the RM to a location server. A mechanism that addresses both risks of exposure is to encapsulate the rules inside an encrypted object that can only be read by an authorized location server (e.g. CMS or S/MIME). Alternatively, the first risk can be mitigated by authenticating the location server to the rule maker; and the second risk can be mitigated by transmitting the rules only over confidentiality-protected channel.

REQ-2 A policy conveyance protocol MUST either support the

encapsulation of rules in an encrypted object format, or else it must provide mechanisms for the RM to authenticate the LS, and to the RM to transmit rules only over a confidentiality-protected channel.

5.3. Acceptance of Rules from Unauthorized Rule Makers

A location server must not accept rules from parties who are not authorized by local policy to update the set of rules used by the location server. This risk can be mitigated in two ways. By encapsulating the rules inside an object that is signed by the authorized rule maker, or by allowing authenticate the LS to authenticate the RM within the policy-handling protocol.

REQ-3 A policy conveyance protocol MUST either support the signing of rules by the rule maker, or else the policy conveyance protocol must provide a mechanism for the LS to authenticate the identity of the RM.

5.4. Unauthorized Exposure of Location Objects

A location object can be exposed to unauthorized parties via a location transmission in two ways. Either a party other than the location recipient in a transmission is able to access the location object en route between the LS and the LR; or else the location server transmits the location object to an unauthorized location recipient. The former risk can be mitigated by transmitting the location object only over an encrypted channel. Mitigation of the latter risk differs depending on whether it is the location server or

the location recipient who initiates the location transmission (i.e., the "push" case vs. the "pull" case). However, note that the mechanisms discussed in this section need not be applied in the case where distribution of a location object is unconstrained, i.e., when the authorization policy of the location server indicates that all possible location recipients are authorized to receive a particular location object.

When the location server initiates the transmission, the LS must apply the authorization policy contained in the appropriate rules (i.e. the rules contained in the location object, obtained from an authorized rule maker, or locally configured on the LS) to determine

if the location recipient is authorized to receive the given location object. In order to have a reliable identity on which to base these authorization decisions, the location server must either authenticate the location recipient within the location conveyance protocol, or else encapsulate the location object in a secure format so that it is accessible only to the authorized recipient.

The case where the location recipient initiates the transaction is further sub-divided depending on whether the location server receives parameters in a "clear" or "opaque" form, as discussed in [Section 3.1.2](#). If the location server receives parameters in a "clear" form, then parameters themselves cannot provide any authentication. In this case, the location server MUST authenticate the identity of the location recipient and then apply authorization policy to determine if the location recipient is authorized to receive the requested location object.

If the location server receives parameters in "opaque" form, then the location server may be able to derive some assurance about the location recipient based on the fact that the location recipient possesses the opaque token(s) presented in a request. In some cases, policy may indicate that the possession of these tokens is sufficient for the location server to determine that the location recipient is authorized to receive a given location object. Note however, that because these opaque parameters are intended to be used by multiple requestors, they are not bound to the identity of any given watcher and thus MUST NOT be used to satisfy a requirement to authenticate the LR via a shared secret (as in [RFC 5025](#)). When policy does not allow the use of these opaque tokens as authorization credentials, the location server MUST authenticate and explicitly authorize the location recipient as in the "clear" case above.

Opaque parameters are to LRs via a dissemination channel (the dotted line in Figure 1), in the form of location references in a location reference format. In formulating policy that determines whether an opaque token suffices for authentication, rule makers and LS operators

should keep in mind that the utility of opaque parameters for authentication is inherently limited by the security of the dissemination channel. An opaque token is a reliable authenticator only if it is only known to authorized location recipients. So a token used as authenticator MUST be provided confidentiality

protection by the dissemination channel, and it MUST contain enough entropy that it is difficult to guess, a minimum of 128 bits.

Dissemination channels can take many different forms, from the SIP Geolocation header to SMTP message bodies to business cards. Because of this diversity, this document does not place requirements on the security features of dissemination protocols, but instead provides recommendations for which protocols should be used as dissemination channels. In particular, it is RECOMMENDED protocols used as dissemination channels provide confidentiality, authenticity, and integrity protection. Conversely, because these properties cannot be guaranteed, it is RECOMMENDED that an LS minimize the risk introduced by this exposure by minimizing the set of LOs to which a location reference refers, when that reference is not subject to authentication and access control.

Finally, it may be the case that some LSs along a distribution path are unauthorized to access an LO that they transmit. In this case, an LO must be encapsulated in an encrypted LO format so that it is only accessible by authorized viewers. This encapsulation may be applied by the LG or by an intermediate LS.

Second, opaque tokens can be retransmitted. Therefore, unless an opaque token format is able to encode retransmission rules, possession of an opaque token is never sufficient to authorize a party to receive a location object for which retransmission is forbidden.

REQ-4 A location conveyance protocol MUST either support the encapsulation of LOs in an encrypted object format, or else it must provide mechanisms for the LS to authorize the LR, and to the LS to transmit LOs only over a confidentiality-protected channel. Input to the authorization process might be the authenticated identity or an opaque token (as a form of proof of possession).

REQ-5 An LS MUST apply authentication and authorization policy to requests in which all parameters are in clear form. When a request contains opaque parameters, it is RECOMMENDED that the same process be followed.

REQ-6 A location reference format MUST define a format for references that requires a cryptographically random component with a minimum entropy of 128 bits.

REQ-7 An LS that does not apply identity-based authorization policy to requests for some references (e.g., for opaque references) MUST minimize the set of locations to which those references refer by setting a restrictive default policy. Additional rules provided by RMs MAY modify this default policy to make it more or less permissive.

REQ-8 In order to support the prevention of unauthorized exposure to intermediate LSs, a location object format MUST include a format in which the LO is encrypted.

[5.5.](#) Unauthorized Modification of Location Objects

Location Objects are at risk of unauthorized modification en route when they are transmitted from the location server to the location recipient. Location objects can be protected against such unauthorized modification if the location conveyance protocol transmits location objects in an integrity-protected format or over an integrity-protected channel. Additionally, a Location Recipient risks receipt of a modified (or fabricated) location object if it does not authenticate that the location object was transmitted by an entity that is authorized (by local policy) to act as a location server.

Note that these protections by the location conveyance protocol need not be used if the location object itself is signed (either by the location generator or by the location server); provided that the location recipient is able to verify this signature. However, when the location recipient is not the location viewer, then the location recipient may be unable to verify a signature intended to provide end-to-end (or middle-to-end) integrity.

REQ-9 A location conveyance protocol MUST either support the conveyance of LOs in an integrity-preserving encapsulation, or else it must offer a mode of operation in which LOs are only transmitted over an integrity-protected channel.

REQ-10 A location conveyance protocol MUST allow the LR to authenticate the LS.

[5.6.](#) Assertion of Location Object Origins

A location generator can assert that it created a particular location object by generating a cryptographic signature over the location object. Such an assertion allows a location viewer to identify the party that created a location object or that participated in its distribution, and thus make local policy decisions based on the origin or intermediate provenance of a location object. Additionally, such a signature can provide end-to-end integrity protection for the portion of the location object covered by the signature.

Likewise, an LS can sign an object to assert that it was included along the distribution path of the LO. The mechanisms discussed in [Section 5.5](#) enable a location recipient to determine that a location object was not modified en route from the most recent location server. However, in a setting where a location object traverses a chain of multiple location servers and location recipients, the ultimate location viewer may not trust every location server in the chain. When a location viewer has a trust relationship with a particular location server in the chain, that server can sign the object to assure the integrity of the location object through multiple transmissions (i.e., to provide middle-to-end integrity protection).

REQ-11 In order to support assertion of the origin and distribution of LOs, and end-to-end or middle-to-end integrity protection, a location object format must enable an LG or LS to cryptographically sign a location object.

[5.7.](#) Summary of Security Requirements

The following security requirements apply to a policy conveyance protocol:

REQ-1 A policy conveyance protocol MUST either support the provision of rules in an integrity-preserving encapsulation, or else it must offer a mode of operation in which rules are only transmitted over an integrity-protected channel.

REQ-2 A policy conveyance protocol MUST either support the encapsulation of rules in an encrypted object format, or else it

must provide mechanisms for the RM to authenticate the LS, and to the RM to transmit rules only over a confidentiality-protected channel.

REQ-3 A policy conveyance protocol MUST either support the signing of rules by the rule maker, or else the policy conveyance protocol must provide a mechanism for the LS to authenticate the identity of the RM.

The following security requirements apply to a location conveyance protocol:

REQ-4 A location conveyance protocol MUST either support the encapsulation of LOs in an encrypted object format, or else it must provide mechanisms for the LS to authenticate the LR, and to the LS to transmit LOs only over a confidentiality-protected channel.

REQ-9 A location conveyance protocol MUST either support the conveyance of LOs in an integrity-preserving encapsulation, or else it must offer a mode of operation in which LOs are only transmitted over an integrity-protected channel.

REQ-10 A location conveyance protocol MUST allow the LR to authenticate the LS.

The following security requirements apply to a secure location object format:

REQ-8 In order to support the prevention of unauthorized exposure to intermediate LSSs, a location object format MUST include a format in which the LO is encrypted.

REQ-11 In order to support assertion of the origin and distribution of LOs, and end-to-end or middle-to-end integrity protection, a location object format must enable an LG or LS to cryptographically sign a location object.

The following security requirements apply to a location reference

format:

REQ-6 A location reference format MUST define a format for references that requires a cryptographically random component with a minimum entropy of 128 bits.

In addition, the following are recommended practices for LS policy:

REQ-4 An LS MUST apply authentication and authorization policy to requests in which all parameters are in clear form. When a request contains opaque parameters, it is RECOMMENDED that the same process be followed.

Barnes, et al.

Expires August 28, 2008

[Page 25]

Internet-Draft

Location Object Security

February 2008

REQ-7 An LS that does not apply authentication and authorization policy to requests for some references MUST minimize the set of locations to which those references refer.

[6.](#) Security Considerations

The focus of this document is the security of location objects. As such, security concerns are discussed throughout.

[7.](#) Acknowledgements

This work was based on the security investigations conducted as part of the GEOPRIV Layer-7 Location Configuration Protocol design team, which produced [[I-D.ietf-geopriv-l7-lcp-ps](#)]. We would like to thank all the members of the design team.

[8.](#) IANA Considerations

This document makes no request of IANA.

[9.](#) References

[9.1.](#) Normative References

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

[9.2](#). Informative References

[I-D.ietf-ecrit-framework]

Rosen, B., Schulzrinne, H., Polk, J., and A. Newton, "Framework for Emergency Calling using Internet Multimedia", [draft-ietf-ecrit-framework-04](#) (work in progress), November 2007.

[I-D.ietf-geopriv-http-location-delivery]

Barnes, M., Winterbottom, J., Thomson, M., and B. Stark, "HTTP Enabled Location Delivery (HELD)", [draft-ietf-geopriv-http-location-delivery-05](#) (work in progress), February 2008.

[I-D.ietf-geopriv-l7-lcp-ps]

Tschofenig, H. and H. Schulzrinne, "GEOPRIV Layer 7 Location Configuration Protocol; Problem Statement and

Barnes, et al.

Expires August 28, 2008

[Page 26]

Internet-Draft

Location Object Security

February 2008

Requirements", [draft-ietf-geopriv-l7-lcp-ps-06](#) (work in progress), November 2007.

[I-D.ietf-geopriv-lbyr-requirements]

Marshall, R., "Requirements for a Location-by-Reference Mechanism", [draft-ietf-geopriv-lbyr-requirements-01](#) (work in progress), October 2007.

[I-D.ietf-geopriv-policy]

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-14](#) (work in progress), February 2008.

[RFC3693]

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

[RFC3694]

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

- [RFC3825] Polk, J., Schnizlein, J., and M. Linsner, "Dynamic Host Configuration Protocol Option for Coordinate-based Location Configuration Information", [RFC 3825](#), July 2004.
- [RFC4119] Peterson, J., "A Presence-based GEOPRIV Location Object Format", [RFC 4119](#), December 2005.
- [RFC4745] 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.
- [RFC4776] Schulzrinne, H., "Dynamic Host Configuration Protocol (DHCPv4 and DHCPv6) Option for Civic Addresses Configuration Information", [RFC 4776](#), November 2006.
- [RFC4825] Rosenberg, J., "The Extensible Markup Language (XML) Configuration Access Protocol (XCAP)", [RFC 4825](#), May 2007.
- [RFC5025] Rosenberg, J., "Presence Authorization Rules", [RFC 5025](#), December 2007.

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

Hannes Tschofenig
Nokia Siemens Networks
Otto-Hahn-Ring 6
Munich, Bavaria 81739
Germany

Email: Hannes.Tschofenig@nsn.com
URI: <http://www.tschofenig.com>

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

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

Full Copyright Statement

Copyright (C) The IETF Trust (2008).

This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgment

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