

Network Working Group
Request for Comments: 3693
Category: Informational

J. Cuellar
Siemens AG
J. Morris
Center for Democracy & Technology
D. Mulligan
Samuelson Law, Technology & Public Policy Clinic
J. Peterson
NeuStar
J. Polk
Cisco
February 2004

Geopriv Requirements

Status of this Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2004). All Rights Reserved.

Abstract

Location-based services, navigation applications, emergency services, management of equipment in the field, and other location-dependent services need geographic location information about a Target (such as a user, resource or other entity). There is a need to securely gather and transfer location information for location services, while at the same time protect the privacy of the individuals involved.

This document focuses on the authorization, security and privacy requirements for such location-dependent services. Specifically, it describes the requirements for the Geopriv Location Object (LO) and for the protocols that use this Location Object. This LO is envisioned to be the primary data structure used in all Geopriv protocol exchanges to securely transfer location data.

Table of Contents

1.	Overview	3
2.	Conventions Used in this Document.	4
3.	Glossary	4
4.	Primary Geopriv Entities	6
5.	Further Geopriv Terminology.	7
5.1.	Location Information and Sighting.	7
5.2.	The Location Object and Using Protocol	9
5.3.	Trusted vs. Non-trusted Data Flows	10
5.4.	Further Geopriv Principals	10
5.5.	Privacy Rules.	12
5.6.	Identifiers, Authentication and Authorization.	13
6.	Scenarios and Explanatory Discussion	15
7.	Requirements	19
7.1.	Location Object.	19
7.2.	The Using Protocol	21
7.3.	Rule based Location Data Transfer.	21
7.4.	Location Object Privacy and Security	22
7.4.1.	Identity Protection.	22
7.4.2.	Authentication Requirements.	23
7.4.3.	Actions to be secured.	23
7.5.	Non-Requirements	24
8.	Security Considerations.	24
8.1.	Traffic Analysis	24
8.2.	Securing the Privacy Rules	24
8.3.	Emergency Case	24
8.4.	Identities and Anonymity	25
8.5.	Unintended Target.	26
9.	Protocol and LO Issues for later Consideration	26
9.1.	Multiple Locations in one LO	26
9.2.	Translation Fields	26
9.3.	Truth Flag	27
9.4.	Timing Information Format.	27
9.5.	The Name Space of Identifiers.	27
10.	Acknowledgements	28
11.	References	28
11.1.	Normative Reference	28
11.2.	Informative References	28
12.	Authors' Addresses	29
13.	Full Copyright Statement	30

1. Overview

Location-based services (applications that require geographic location information as input) are becoming increasingly common. The collection and transfer of location information about a particular Target can have important privacy implications. A key goal of the protocol described in this document is to facilitate the protection of privacy pursuant to Privacy Rules set by the "user/owner of the Target" (or, more precisely in the terminology of this document given in Section 3 and 5.4 below, the "Rule Maker").

The ability to gather and generate a Target's location, and access to the derived or computed location, are key elements of the location-based services privacy equation. Central to a Target's privacy are (a) the identity of entities that have access to raw location data, derive or compute location, and/or have access to derived or computed location information, and (b) whether those entities can be trusted to know and follow the Privacy Rules of the user.

The main principles guiding the requirements described in this document are:

- 1) Security of the transmission of Location Object is essential to guarantee the integrity and confidentiality of the location information. This includes authenticating the sender and receiver of the Location Object, and securing the Location Object itself.
- 2) A critical role is played by user-controlled Privacy Rules, which describe the restrictions imposed or permissions given by the "user" (or, as defined below, the "Rule Maker"). The Privacy Rules specify the necessary conditions that allow a Location Server to forward Location Information to a Location Recipient, and the conditions and purposes for which the Location Information can be used.
- 3) One type of Privacy Rules specify how location information should be filtered, depending on who the recipient is. Filtering is the process of reducing the precision or resolution of the data. A typical rule may be of the form: "my location can only be disclosed to the owner of such credentials in such precision or resolution" (e.g., "my co-workers can be told the city I am currently in").
- 4) The Location Object should be able to carry a limited but core set of Privacy Rules. The exact form or expressiveness of those Rules in the core set or in the full set is not further discussed in this document, but will be discussed more extensively in future documents produced by this working group.

- 5) Whenever appropriate, the location information should not be linked to the real identity of the user or a static identifier easily linked back to the real identity of the user (i.e., Personally Identifiable Information such as a name, mailing address, phone number, social security number, or email address or username). Rather, the user should be able to specify which local identifier, unlinked pseudonym, or private identifier is to be bound to the location information.
- 6) The user may want to hide the real identities of himself and his partners, not only to eavesdroppers but also to other entities participating in the protocol.

Although complete anonymity may not be appropriate for some applications because of legal constraints or because some location services may in fact need explicit identifications, most often the location services only need some type of authorization information and/or perhaps anonymous identifiers of the entities in question.

2. Conventions Used in this Document

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].

Note that the requirements discussed here are requirements on the generic Location Object and on using protocols for location services. Thus, for the most part, the requirements discussed in this document refer to capabilities that are mandatory-to-implement. For example, requiring that implementations support integrity is not the same thing as requiring that all protocol traffic be authenticated. In contrast, an example of a mandatory-to-use (not just mandatory-to-implement) requirement might be one that states that the user always receives a notice when his location data was not authenticated. This practice is mandatory-to-use, not just to implement.

3. Glossary

For easy reference and readability, below are basic terms that will be defined more formally and fully later in this document.

Location Generator (LG): The entity that initially determines or gathers the location of the Target and creates Location Objects describing the location of the Target.

Location Object (LO): An object conveying location information (and possibly privacy rules) to which Geopriv security mechanisms and privacy rules are to be applied.

Location Recipient (LR): The entity that receives location information. It may have asked for this location explicitly (by sending a query to a location server), or it may receive this location asynchronously.

Location Server (LS): The entity to which a LG publishes location objects, the recipient of queries from location receivers, and the entity that applies rules designed by the rule maker.

Precision: The number of significant digits to which a value has been reliably measured.

Principal: The holder/subject of the credentials, e.g., a workstation user or a network server.

Resolution: The fineness of detail that can be distinguished in a measured area. Applied to Geopriv this means the finite area within provided and closed borders (ex. Latitude and Longitude boundaries).

Rule Holder: The entity that provides the rules associated with a particular target for the distribution of location information. It may either 'push' rules to a location server, or a location server may 'pull' rules from the Rule Holder.

Rule Maker: The authority that creates rules governing access to location information for a target (typically, this is the target themselves).

Rule, or Privacy Rule: A directive that regulates an entity's activities with respect to location information, including the collection, use, disclosure, and retention of location information.

Target: A person or other entity whose location is communicated by a Geopriv Location Object.

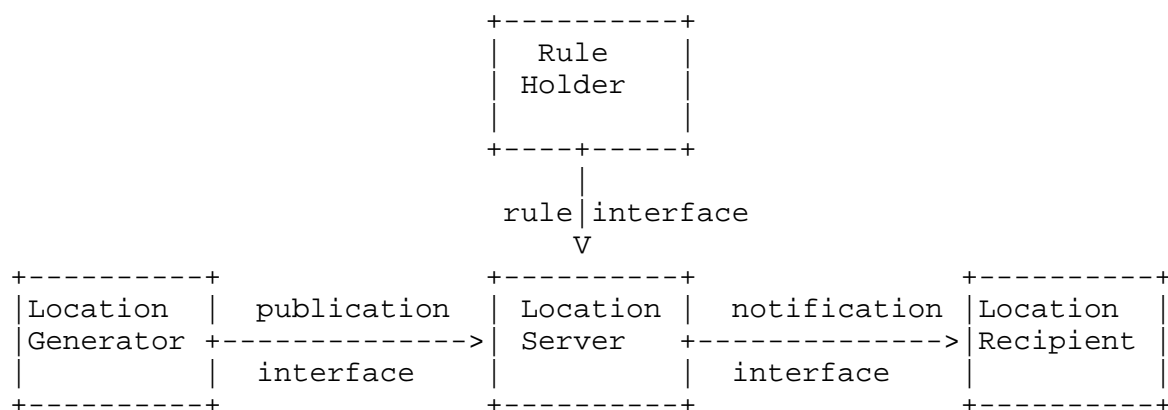
Using Protocol: A protocol that carries a Location Object.

Viewer: A Principal that consumes location information that is communicated by a Geopriv Location Object, but does not pass this information further.

Resolution and Precision are very close terms. Either quality can be 'reduced' to coarsen location information: 'resolution' by defining a off-center perimeter around a user's location or otherwise enlarging the area in consideration (from state to country, say) and 'precision' by discarding significant digits of positioning information (rounding off longitude and latitude from seconds to minutes, say). Another WG document discusses this topic in much more detail.

4. Primary Geopriv Entities

The following picture shows the primary Geopriv entities in a simple and basic architecture, without claim of completeness or any suggestion that the entities identified must in all cases be physically separate entities.



The four primary Entities are described as follows:

Location Generator (LG): The entity that initially determines or gathers the location of the Target and creates Location Objects describing that location. LGs publish Location Objects to Location Servers. The manner in which the Location Generator learns of Location Information is outside the scope of the Geopriv Protocol.

Location Server (LS): The LS is an element that receives publications of Location Objects from Location Generators and may receive subscriptions from Location Recipients. The LS applies the rules (which it learns from the Rule Holder) to LOs it receives from LGs, and then notifies LRs of resulting LOs as necessary.

Location Recipient (LR): The LR is an element that receives notifications of Location Objects from Location Servers. The LR may render these LOs to a user or automaton in some fashion.

Rule Holder (RH): The RH is an element that houses Privacy Rules for receiving, filtering and distributing Location Objects for specific Targets. An LS may query an RH for a set of rules, or rules may be pushed from the RH to an LS. The rules in the Rule Holder are populated by the Rule Maker.

Thus Location Generation is the process of gathering Location Information, perhaps from multiple sources, at an IP-based Geopriv Entity, the LG, which communicates with other Geopriv Entities.

Rules MUST be authenticated and protected. How this is done and in particular how to distribute the keys to the RM and other authorities is outside of the scope of this document. See also Section 8.2, "Securing the Privacy Rules".

The interfaces between the Geopriv entities are not necessarily protocol interfaces; they could be internal interfaces within a single composed device. In some architectures, the Location Generator, Rule Holder, and Location Server might all be implemented in the same device. There may be several Rule Holders that enforce the Privacy Rules at a particular Location Server.

5. Further Geopriv Terminology

The terminology and definitions detailed below include both terms that, besides the primary Geopriv entities, (1) are used in the requirements section of this document, and (2) provide additional detail about the usage model envisioned for the Geopriv Location Object. These latter terms will be utilized in a separate scenarios document and elsewhere.

5.1. Location Information and Sighting

The focus of the Geopriv working group is on information about a Target's location that is NOT based on generally or publicly available sources, but instead on private information provided or created by a Target, a Target's Device, or a Target's network or service provider. Notwithstanding this focus on private location information, the Geopriv Location Object could certainly be used to convey location information from publicly available sources.

Location Information: A relatively specific way of describing where a Device is located.

This Location Information may have been determined in many different ways, including:

(a) derived or computed from information generally not available to the general public (such as information mainly available to a network or service provider), (b) determined by a Device that may not be generally publicly addressable or accessible, or (c) input or otherwise provided by a Target.

As examples, the Location Information could include (a) information calculated by triangulating on a wireless signal with respect to cell phone towers, (b) longitude and latitude information determined by a Device with GPS (global positioning satellite) capabilities, (c) information manually entered into a cell phone or laptop by a Target in response to a query, or (d) automatically delivered by some other IP protocol, such as at device configuration via DHCP.

Excluded from this definition is the determination of location information wholly without the knowledge or consent of the Target (or the Target's network or access service provider), based on generally available information such as an IP or e-mail address. In some cases, information like IP address can enable someone to estimate (at least roughly) a location. Commercial services exist that provide rough location information based on IP addresses. Currently, this type of location information is typically less precise than the type of location information addressed in this document. Although this type of location computation still raises significant potential privacy and public privacy concerns, such scenarios are generally outside the scope of this document.

Within any given location-based transaction, the INITIAL determination of location (and thus the initial creation of Location Information) is termed a Sighting:

Sighting:

The initial determination of location based on non-public information (as discussed in the definition of Location Information), and the initial creation of Location Information.

Some variant of the sighting information is included in the Location Object. Abstractly, it consists of two separate data fields:

(Identifier, Location)

where Identifier is the identifier assigned to a Target being sighted, and Location is the current position of that Target being sighted. Not all entities may have access to exactly the same piece of sighting information. A sighting may be transformed to a new sighting pair:

(Identifier-1, Location-1)

before it is provided by a Location Generator or Location Server to Location Recipient. In this case, Identifier-1 may be a Pseudonym, and Location-1 may have less precision or resolution than the original value.

5.2. The Location Object and Using Protocol

A main goal of the Geopriv working group is to define a Location Object (LO), to be used to convey both Location Information and basic privacy-protecting instructions:

Location Object (LO): This data contains the Location Information of the Target, and other fields including an identity or pseudonym of the Target, time information, core Privacy Rules, authenticators, etc. Most of the fields are optional, including the Location Information itself.

Nothing is said about the semantics of a missing field. For instance, a partially filled object MAY be understood implicitly as a request to complete it. Or, if no time information is included, this MAY implicitly mean "at the current time" or "at a very recent time", but it could be interpreted in a different way, depending on the context.

The "using protocol" is the protocol that uses (reads or modifies) the Location Object. A protocol that just transports the LO as a string of bits, without looking at them (like an IP storage protocol could do), is not a using protocol, but only a transport protocol. Nevertheless, the entity or protocol that caused the transport protocol to move the LO is responsible for the appropriate distribution, protection, usage, retention, and storage of the LO based on the rules that apply to that LO.

The security and privacy enhancing mechanisms used to protect the LO are of two types: First, the Location Object definition MUST include the fields or mechanisms used to secure the LO as such. The LO MAY be secured, for example, using cryptographic checksums or encryption as part of the LO itself. Second, the using protocol may also provide security mechanisms to securely transport the Location Object.

When defining the LO, the design should observe that the security mechanisms of the Location Object itself are to be preferred. Thus the definition of the LO MUST include some minimal crypto functionality (Req. 14 and 15). Moreover, if the RM specifies the use of a particular LO security mechanism, it MUST be used (Req. 4).

5.3. Trusted vs. Non-trusted Data Flows

Location information can be used in very different environments. In some cases, the participants will have longstanding relationships, while in others the participants may have discrete interactions with no prior contractual or other contact.

The different relationships raise different concerns for the implementation of privacy rules, including the need to communicate Privacy Rules. A public Rule Holder, for example, may be unnecessary in a trusted environment where more efficient methods of addressing privacy issues exist. The following terms distinguish between the two basic types of data flows:

Trusted Data Flow:

A data flow that is governed by a pre-existing contractual relationship that addresses location privacy.

Non-trusted Data Flow:

The data flow is not governed by a pre-existing contractual relationship that addresses location privacy.

5.4. Further Geopriv Principals

Target:

The entity whose location is desired by the Location Recipient. In many cases the Target will be the human "user" of a Device or an object such as a vehicle or shipping container to which the Device is attached. In some instances the Target will be the Device itself.

Device:

The technical device whereby the location is tracked as a proxy for the location of a Target.

A Device might, for example, be a cell phone, a Global Positioning Satellite (GPS) receiver, a laptop equipped with a wireless access Device, or a transmitter that emits a signal that can be tracked or located. In some situations, such as when a Target manually inputs location information (perhaps with a web browser), the Target is effectively performing the function of a Device.

Rule Maker (RM):

The individual or entity that has the authorization to set the applicable Privacy Rules for a potential Geopriv Target. In many cases this will be the owner of the Device, and in other cases this may be the user who is in possession of the Device. For example, parents may control what happens to the location information derived from a child's cell phone. A company, in contrast, may own and provide a cell phone to an employee but permit the employee to set the privacy rules.

There are four scenarios in which some form of constraint or override might be placed on the Privacy Rules of the Rule Maker:

1. In the case of emergency services (such as E911 within the United States), local or national laws may require that accurate location information be transmitted in certain defined emergency call situations. The Geopriv Working Group MUST facilitate this situation.
2. In the case of legal interception, the RM may not be aware of an override directive imposed by a legal authority. It is not the expectation of the Working Group that a particular accommodation will be made to facilitate this situation.
3. In the context of an employment relationship or other contractual relationship, the owner of a particular location (such as a corporate campus) may impose constraints on the use of Privacy Rules by a Rule Maker. It is not the expectation of the Working Group that a particular accommodation will be made to facilitate this situation.
4. It is conceivable that a governmental authority may seek to impose constraints on the use of Privacy Rules by a Rule Maker in non-emergency situations. It is not the expectation of the Working Group that a particular accommodation will be made to facilitate this situation.

Viewer:

An individual or entity who receives location data about a Target and does not transmit the location information or information based on the Target's location (such as driving directions to or from the Target) to any party OTHER than the Target or the Rule Maker.

Data Transporter:

An entity or network that receives and forwards data without processing or altering it. A Data Transporter could theoretically be involved in almost any transmission between a Device and a Location Server, a Location Server and a second Location Server, or a Location Server and a Viewer. Some location tracking scenarios may not involve a Data Transporter.

Access Provider (AP):

The domain that provides the initial network access or other data communications services essential for the operation of communications functions of the Device or computer equipment in which the Device operates. Often, the AP -- which will be a wireless carrier, an Internet Service Provider, or an internal corporate network -- contains the LG. Sometimes the AP has a "dumb" LG, one that transmits Geopriv LOs but does not use any part of the Geopriv Location Object. Other cases may not involve any AP, or the AP may only act as a Data Transporter.

Location Storage:

A Device or entity that stores raw or processed Location Information, such as a database, for any period of time longer than the duration necessary to complete an immediate transaction regarding the Location Information.

The existence and data storage practices of Location Storage is crucial to privacy considerations, because this may influence what Location Information could eventually be revealed (through later distribution, technical breach, or legal processes).

5.5. Privacy Rules

Privacy Rules are rules that regulate an entity's activities with respect to location and other information, including, but not limited to, the collection, use, disclosure, and retention of location information. Such rules are generally based on fair information practices, as detailed in (for example) the OECD Guidelines on the Protection of Privacy and Transporter Flows of Personal Data [OECD].

Privacy Rule:

A rule or set of rules that regulate an entity's activities with respect to location information, including the collection, use, disclosure, and retention of location information. In particular, the Rule describes how location information may be used by an entity and which transformed location information may be released to which entities under which conditions. Rules must be obeyed; they are not advisory.

A full set of Privacy Rules will likely include both rules that have only one possible technical meaning, and rules that will be affected by a locality's prevailing laws and customs. For example, a distribution rule of the form "my location can only be disclosed to the owner of such credentials and in such precision or resolution" has clear-cut implications for the protocol that uses the LO. But other rules, like retention or usage Rules, may have unclear technical consequences for the protocol or for the involved entities. For example, the precise scope of a retention rule stating "you may not store my location for more than 2 days" may in part turn on local laws or customs.

5.6. Identifiers, Authentication and Authorization

Anonymity is the property of being not identifiable (within a set of subjects). Anonymity serves as the base case for privacy: without the ability to remain anonymous, individuals may be unable to control their own privacy. Unlinkability ensures that a user may make multiple uses of resources or services without others being able to link these uses to each other. Unlinkability requires that entities be unable to determine whether the same user caused certain specific operations in the system. [ISO99] A pseudonym is simply a bit string which is unique as an ID and is suitable to be used for end-point authentication.

Unlinked Pseudonym:

A pseudonym where the linking between the pseudonym and its holder is, at least initially, not known to anybody with the possible exception of the holder himself or a trusted server of the user. See [Pfi01] (there the term is called Initially Unlinked Pseudonym).

The word authentication is used in different manners. Some require that authentication associates an entity with a more or less well-known identity. This basically means that if A authenticates another entity B as being "id-B", then the label "id-B" is a well-known, or at least a linkable identity of the entity. In this case, the label "id-B" is called a publicly known identifier, and the authentication is "explicit":

Explicit Authentication:

The act of verifying a claimed identity as the sole originator of a message (message authentication) or as the end-point of a channel (entity authentication). Moreover, this identity is easily linked back to the real identity of the entity in question, for instance being a pre-existing static label from a predefined name space (telephone number, name, etc.)

Authorization:

The act of determining if a particular right, such as access to some resource, can be granted to the presenter of a particular credential.

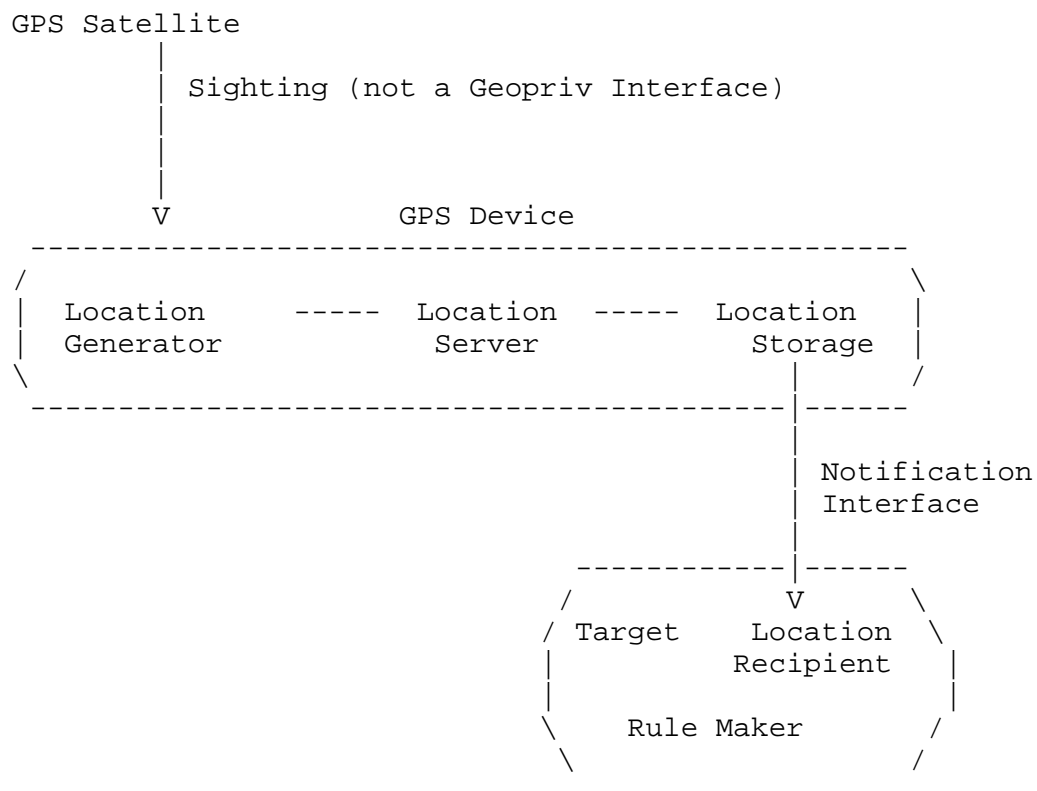
Depending on the type of credential, authorization may or may not imply Explicit Authentication.

6. Scenarios and Explanatory Discussion

In this subsection we introduce short scenarios to illustrate how these terms and attributes describe location information transactions. Additional illustrative scenarios are discussed in a separate document.

SCENARIO 1: GPS Device with Internal Computing Power: Closed System

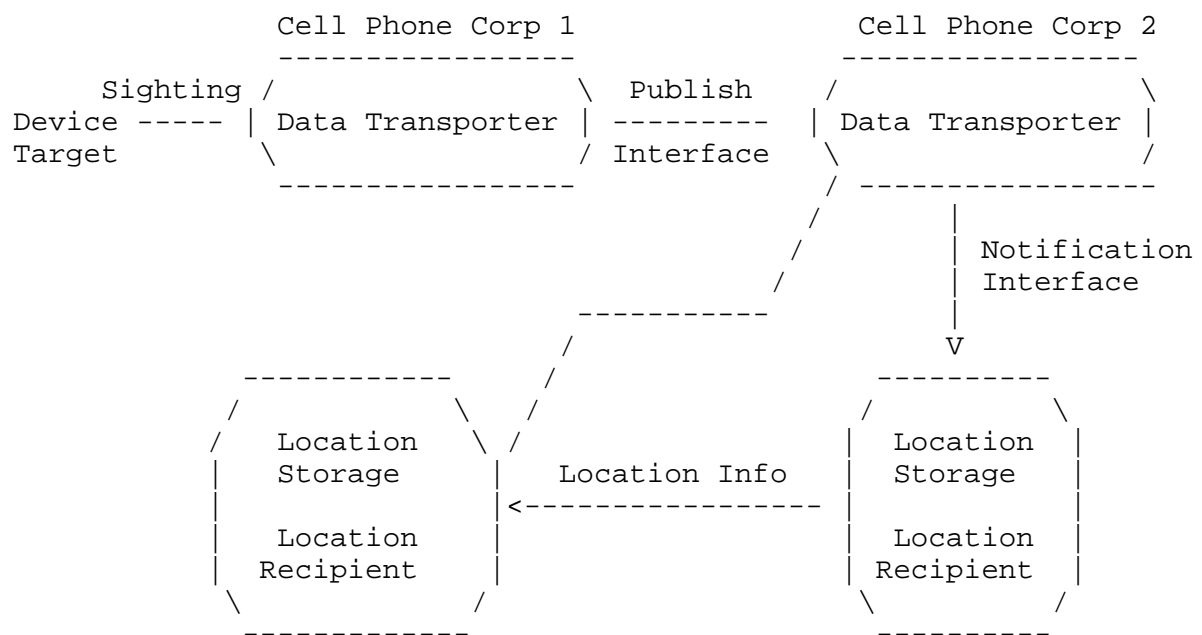
In this example, the Target wishes to know his/her location using the Global Positioning System (GPS) and the Device is capable of independently processing the raw data to determine its location. The location is derived as follows: the Device receives transmissions from the GPS satellites, internally computes and displays location. This is a closed system. For the purpose of this and subsequent examples, it is assumed that the GPS satellite broadcasts some signal, and has no information about the identity or whereabouts of Devices using the signal.



In this scenario the GPS Device is both the AP and the LG. The interaction occurs in a Trusted environment because it occurs in the Rule Maker's Device.

SCENARIO 2: Cell Phone Roaming

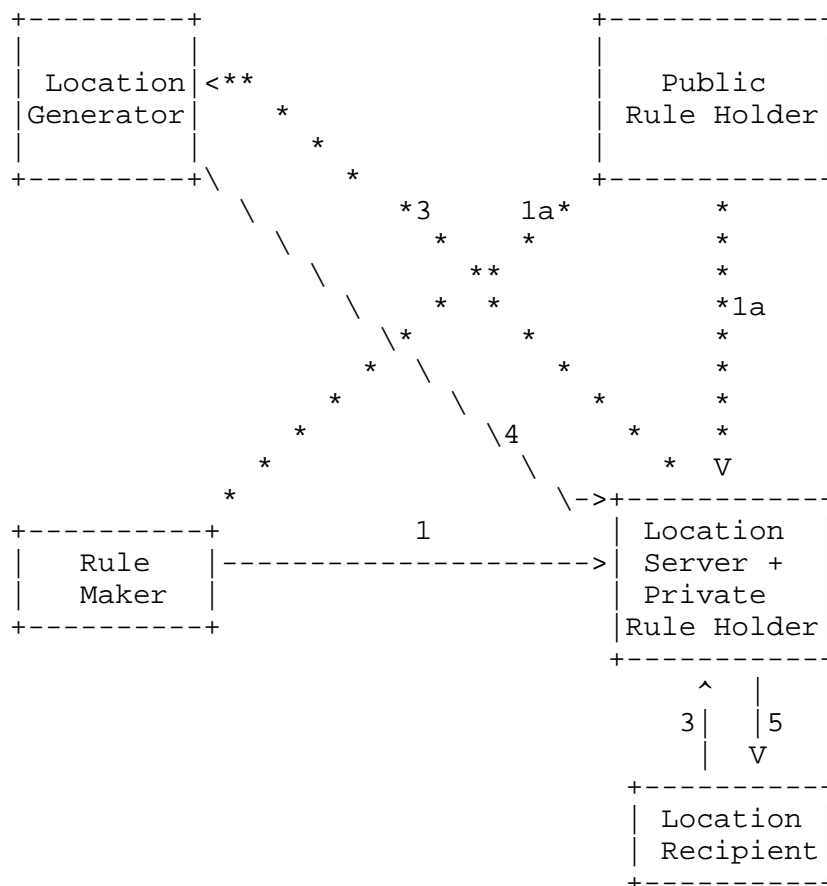
In this example, a cell phone is used outside its home service area (roaming). Also, the cell phone service provider (cell phone Corp 2) outsourced the accounting of cell phone usage. The cell phone is not GPS-enabled. Location is derived by the cell phone network in which the Target and Device are roaming. When the Target wishes to use the cell phone, cell phone Corp 1 (AP) provides the roaming service for the Target, which sends the raw data about usage (e.g., duration of call, location in the roaming network, etc.) to cell phone Corp 2, the home service provider. Cell phone Corp 2 submits the raw data to the accounting company, which processes the raw data for the accounting statements. Finally, the raw data is sent to a data warehouse where the raw data is stored in a Location Server (e.g., computer server).



Here, cell phone Corp 1 is the AP and the LG. In this scenario, Cell phone Corp 2 is likely to be a Trusted entity, but cell phone Corp 1 may be Non-trusted.

SCENARIO 3: Mobile Communities and Location-Based Services

The figure below shows a common scenario, where a user wants to find his friends or colleagues or wants to share his position with them or with a Location-Based Service Provider. Some of the messages use a Location Object to carry, for instance, identities or pseudonyms, credentials and proof-of-possession of them, Rules and Location Data Information, including Data Types and Precision or Resolution. Messages that do not use the Location Object and are outside of the scope of the Geopriv WG, but should be mentioned for understandability, are shown in the figure as starred arrows ("***>").



Assume that the Rule Maker and the Target are registered with the Location Server. The RM has somehow proven to the LS that he indeed is the owner of the privacy rights of the Target (the Target is usually a Device owned by the Rule Maker). The Rule Maker and the Location Server have agreed on the set of keys or credentials and cryptographic material that they will use to authenticate each other.

and in particular, to authenticate or sign the Rules. How this has been done is outside of the scope of the document.

1: Rule Transfer:

The Rule Maker sends a Rule to the Location Server. This Rule may or may not be a field in a Location Object.

1a: Signed Rule:

As an alternative, the Rule Maker may write a Rule and place it in a Public Rule Holder. The entities access the repository to read the signed Rules.

2: Location Information Request:

The Location Recipient requests location information for a Target. In this request, the Location Recipient may select which location information data type it prefers. One way of requesting Location Information MAY be sending a partially filled Location Object, including only the identities of the Target and Location Recipient and the desired Data Type and precision or resolution, and providing proof of possession of the required credentials. But whether or not the using protocol understands this partially filled object as a request MAY depend on the using protocol or on the context. The Location Recipient could also specify the need for periodic location information updates, but this is probably out of the scope of Geopriv.

3: Locate:

When a Location Server receives a Location Information Request for a Target which has no current location information, the server may ask the Location Generator to locate the Target.

4: Location Information:

The Location Generator sends the "full" location information to the Location Server. This Location Information may or may not be embedded in a Location Object.

5: Filtered Location Information:

The Location Server sends the location information to the Location Recipient. The information may be filtered in the sense that in general a less precise or a computed version of the information is being delivered.

7. Requirements

7.1. Location Object

Remember that this document is primarily specifying requirements on the definition of the LO. Some Requirements read like this: "The LO definition MUST contain Field 'A' as an optional field." This requirement states that

- o the document that defines the LO MUST define the LO field 'A',
- o the field 'A' MUST be defined as optional to use (an instance of a LO MAY or may not contain the field 'A').

Some Requirements read like this: "The LO definition MUST contain Field 'A', which MAY be an optional field." This requirement states that

- o the document that defines the LO MUST define the LO field 'A',
- o the field 'A' MAY be defined as optional or not to use. If it is defined as optional to use, any instance of an LO MAY or may not contain the field 'A'; if it is not optional, all instances of LOs MUST contain the field 'A'.

Req. 1. (Location Object generalities)

1.1) Geopriv MUST define one Location Object (LO) -- both in syntax and semantics -- that must be supported by all Geopriv entities.

1.2) Some fields of the Location Object MAY be optional. This means that an instance of a Location Object MAY or may not contain the fields.

1.3) Some fields of the Location Object MAY be defined as "extensions". This means that the syntax or semantics of these fields is not fully defined in the basic Location Object definition, but their use may be private to one or more of the using protocols.

1.4) The Location Object MUST be extensible, allowing the definition of new attributes or fields.

1.5) The object MUST be suitable for requesting and receiving a location.

1.6) The object MUST permit (but not require) the Privacy Rules to be enforced by a third party.

1.7) The object MUST be usable in a variety of protocols, such as HTTP and SIP, as well as local APIs.

1.8) The object MUST be usable in a secure manner even by applications on constrained devices.

Req. 2. (Location Object fields) The Location Object definition MUST contain the following Fields, which MAY be optional to use:

2.1) Target Identifier

2.2) Location Recipient Identity

This identity may be a multicast or group identity, used to include the Location Object in multicast-based using protocols.

2.3) Location Recipient Credential

2.4) Location Recipient Proof-of-Possession of the Credential

2.5) Location Field

2.5.1) Motion and direction vectors. This field MUST be optional.

2.6) Location Data Type

When transmitting the Location Object, the sender and the receiver must agree on the data type of the location information. The using protocol may specify that the data type information is part of the Location Object or that the sender and receiver have agreed on it before the actual data transfer.

2.7) Timing information:

(a) When was the Location Information accurate? (sighting time)

(b) Until when considered current? TTL (Time-to-live) (This is different than a privacy rule setting a limit on data retention)

2.8) Rule Field: this field MAY be a referral to an applicable Rule (for instance, a URI to a full Rule), or it MAY contain a Limited Rule (see Req. 11), or both.

2.9) Security-headers and -trailers (for instance encryption information, hashes, or signatures) (see Req. 14 and 15).

2.10) Version number

Req. 3. (Location Data Types)

3.1) The Location Object MUST define at least one Location Data Type to be supported by all Geopriv receivers (entities that receive LOs).

3.2) The Location Object SHOULD define two Location Data Types: one for latitude / longitude / altitude coordinates and one for civil locations (City, Street, Number) supported by all Geopriv receivers (entities that receive LOs).

3.3) The latitude / longitude / altitude Data Type SHOULD also support a delta format in addition to an absolute one, used for the purpose of reducing the size of the packages or the security and confidentiality needs.

3.4) The Location Object definition SHOULD agree on further Location Data Types supported by some Geopriv entities and defined by other organizations.

7.2. The Using Protocol

Req. 4. The using protocol has to obey the privacy and security instructions coded in the Location Object and in the corresponding Rules regarding the transmission and storage of the LO.

Req. 5. The using protocol will typically facilitate that the keys associated with the credentials are transported to the respective parties, that is, key establishment is the responsibility of the using protocol.

Req. 6. (Single Message Transfer) In particular, for tracking of small target devices, the design should allow a single message/packet transmission of location as a complete transaction.

Other requirements on the using protocol are out of the scope of this document, but might be the subject of future efforts from this working group. See also Section 9 (Protocol and LO Issues for later Consideration).

7.3. Rule based Location Data Transfer

Req. 7. (LS Rules) The decision of a Location Server to provide a Location Recipient access to Location Information MUST be based on Rule Maker-defined Privacy Rules.

It is outside of our scope how Privacy Rules are managed and how a Location Server has access to the Privacy Rules. Note that it might be that some rules contain private information not intended for untrusted parties.

Req. 8. (LG Rules) Even if a Location Generator is unaware of and lacks access to the full Privacy Rules defined by the Rule Maker, the Location Generator MUST transmit Location Information in compliance with instructions set by the Rule Maker. Such compliance MAY be accomplished by the Location Generator transmitting the LO only to a URI designated by the Rule Maker.

Req. 9. (Viewer Rules) A Viewer does not need to be aware of the full Rules defined by the Rule Maker (because a Viewer SHOULD NOT retransmit Location Information), and thus a Viewer SHOULD receive only the subset of Privacy Rules necessary for the Viewer to handle the LO in compliance with the full Privacy Rules (such as, instruction on the time period for which the LO can be retained).

Req. 10. (Full Rule language) Geopriv MAY specify a Rule language capable of expressing a wide range of privacy rules concerning location information. This Rule language MAY be an existing one, an adaptation of an existing one or a new Rule language, and it SHOULD be as simple as possible.

Req. 11. (Limited Rule language) Geopriv MUST specify a limited Rule language capable of expressing a limited set of privacy rules concerning location information. This Rule language MAY be an existing one, an adaptation of an existing one or a new Rule language. The Location Object MUST include sufficient fields and data to express the limited set of privacy rules.

7.4. Location Object Privacy and Security

7.4.1. Identity Protection

Req. 12. (Identity Protection) The Location Object MUST support use of Unlinked Pseudonyms in the corresponding identification fields of Rule Maker, Target, Device, and Location Recipient. Since Unlinked Pseudonyms are simply bit strings that are not linked initially to a well-known identity, this requirement boils down to saying that the name space for Identifiers used in the LO has to be large enough to contain many unused strings.

7.4.2. Authentication Requirements

Req. 13. (Credential Requirements) The using protocol and the Location Object SHOULD allow the use of different credential types, including privacy-enhancing credentials (for instance those described in [Bra00] or [Cha85]).

7.4.3. Actions to be secured

Req. 14. (Security Features) The Location Object MUST support fields suitable for protecting the Object to provide the following security features:

14.1) Mutual end-point authentication: the using protocol is able to authenticate both parties in a Location Object transmission,

14.2) Data object integrity: the LO is secured from modification by unauthorized entities during transmission and storage,

14.3) Data object confidentiality: the LO is secured from eavesdropping (unauthorized reading) during transmission and storage, and

14.4) Replay protection: an old LO may not be replayed by an adversary or by the same entity that used the LO itself (except perhaps during a small window of time that is configurable or accepted by the Rule Maker).

Req. 15. (Minimal Crypto)

15.1) Geopriv MUST specify a minimum mandatory to implement Location Object security, including mandatory to implement crypto algorithms for digital signature algorithms and encryption algorithms.

15.2) It MAY also define further mandatory to implement Location Object security mechanisms for message authentication codes (MACs) or other purposes.

15.3) The protocol SHOULD allow a bypass if authentication fails in an emergency call.

The issue addressed in the last point is that an emergency call in some unfavorable situations may not be completed if the minimal authentication fails. This is probably not what the user would like to happen. The user may prefer an unauthenticated call to an

unauthenticated emergency server over no call completion at all, even at the risk that he is talking to an attacker or that his information is not secured.

7.5. Non-Requirements

Non-Req. 1. (Bridging to non-IP networks) The Geopriv specification SHOULD NOT specify the bridging to non-IP networks (PSTN, etc).

8. Security Considerations

The purpose of the Geopriv Location Object and the requirements on the using protocol are to allow a Privacy Rule-controlled disclosure of location information for location services.

8.1. Traffic Analysis

The information carried within the Location Object is secured in a way compliant with the privacy and security Rules of the Rule Maker, but other information, carried in other objects or headers are in general not secured in the same way. This means that Geopriv may not as a general matter, secure the Target against general traffic analysis attacks or other forms of privacy violations.

8.2. Securing the Privacy Rules

The Privacy Rules of the Rule Maker regarding the location of the Target may be accessible to a Location Server in a public or non-public Rule Holder, or they may be carried by the Location Object, or they may be presented by the Location Recipient as capabilities or tokens. Each type of Rule has to be secured its own particular way.

The rules in a non-public Rule Holder are typically authenticated using a MAC (Message Authentication Code) or a signature, depending on the type of keys used. The rules in a public Rule Holder (one that in principle may be accessed directly by several entities, for instance several Location Servers) are typically digitally signed. Rule Fields in an LO are secured as part of the LO itself. A Geopriv Token (a token or ticket issued by the Rule Maker to a Location Recipient, expressing the explicit consent of the Rule Maker to access his location information) is authenticated or signed.

8.3. Emergency Case

Let us consider the situation where the authentication fails in an emergency call because the authentication center fails to authenticate itself. In this case, one way of implementing the

authentication bypass for emergency calls (mentioned in Req 15.3) is to let the user have the choice of writing a Rule that says:

- "If the emergency server does not authenticate itself, send the location information anyway", or
- "If the emergency server does not authenticate itself, let the call fail".

Second, in the case where the authentication of the emergency call fails because the user may not authenticate itself, the question arises: whose Rule to use? It is reasonable to use a default one: this location information can only be sent to an emergency center.

The third situation, which should be studied in more detail, is: what to do if not only the user fails to authenticate itself, but also the emergency center is not authenticable? It is reasonable to send the Location Information anyway, but are there any security threats that must be considered?

8.4. Identities and Anonymity

The use of Unlinked Pseudonyms is necessary to obtain anonymity.

The purpose of the use of Unlinked Pseudonyms is the following: the using protocol should be able to hide the real identity of the Rule Maker, the Target, and the Device, from Location Servers or Location Recipients, if required by the RM. Also, the using protocol SHOULD be able to hide the real identity of the Location Recipient from the Location Server.

In this last case, the Target is not concerned about the Server identifying him and knowing his location, but identifying his business partners, and therefore his habits, etc. Reasons for hiding the real identities of the Location Recipients include (a) that this knowledge may be used to infer the identity of the Target, (b) that knowledge of the identity of the Location Recipient may embarrass the Target or breach confidential information, and (c) that the dossier telling who has obtained a Target's location information over a long period of time can give information on habits, movements, etc. Even if the location service providers agree to respect the privacy of the user, are compelled by laws or regulations to protect the privacy of the user, and misbehavior or negligence of the Location Server can be ruled out, there is still risk that personal data may become available to unauthorized persons through attacks from outsiders, unauthorized access from insiders, technical or human errors, or legal processes.

On some occasions, a Location Server has to know who is supplying the Privacy Rules for a particular Target, while in other situations it could be enough to know that the supplier of the Rules is authorized to do so.

8.5. Unintended Target

An Unintended Target is a person or object tracked by proximity to the Target. This special case most frequently occurs if the Target is not a person. For example, the Target may be a rental car equipped with a GPS Device, used to track car inventory. The rental company may not care about the driver's location, but the driver's privacy is implicitly affected.

Geopriv may or may not protect or affect the privacy of Unintended Targets, but the impact on Unintended Targets should be acknowledged.

9. Protocol and LO Issues for later Consideration

This section briefly discusses issues relating to the Location Object or the protocol that have emerged during the discussion of earlier versions of this document.

9.1. Multiple Locations in one LO

A location Field is intended to represent one point or one region in space (either 1, 2, or 3 dimensionally). The possibility of inclusion of multiple locations is discussed in another document. The current rough consensus is the following: the LO definition MAY allow the Location Field to be optional, to appear exactly one time or to occur several times. Each Location Field may contain one or more "Location Representations", each of which is intended to represent a different measurement or a different formatting of the same position. But there are other possibilities for using multiple Location Fields and multiple representations: maybe several Location Fields would be used to report the same sighting in different formats, or multiple sightings at different times, or multiple sensor locations for the same device, or other purposes, which could also depend on the using protocol. This is all for further discussion.

9.2. Translation Fields

It is possible to include fields to indicate that one of the locations is a translation of another. If this is done, it is also possible to have a field to identify the translator, as identity and method.

9.3. Truth Flag

Geopriv MUST be silent on the truth or lack-of-truth of the location information contained in the LO. Thus, the LO MUST NOT provide an attribute in object saying "I am (or am not) telling you the whole truth."

9.4. Timing Information Format

The format of timing information is out of the scope of this document.

9.5. The Name Space of Identifiers

Who defines the Identities: can the using protocol define the Identifiers or must the using protocol use and authenticate Pseudonyms proposed by the Rules, chosen independently of the using protocol? Of course, if the using protocol has an appropriate namespace, containing many unused names that may be used as pseudonyms and may be replaced by new ones regularly, then the Location Object may be able to use the name space. For this purpose, the user would probably have to write his Rules using this name space. Note that it is necessary to change the used pseudonyms regularly, because identifying the user behind an unlinked pseudonym can be very simple.

There are several advantages in letting the using protocol define the name space:

- o the embedded authentication would be easier, as the using protocol often already has the credentials for the authentication identity in place and the "embedded" authentication would be independent on the form of Identifiers,
- o the size of the names would be fixed.

On the other hand, the benefits of the Rule choosing the identifiers are:

- o the user has a control of his anonymity, and
- o the interworking of multiple systems with Location object across protocol boundaries is facilitated.

10. Acknowledgements

We wish to thank the members of the IETF Geopriv WG for their comments and suggestions. Aaron Burstein, Mehmet Ersue, Allison Mankin, Randall Gellens, and the participants of the Geopriv meetings in San Diego and Yokohama provided detailed comments or text.

11. References

11.1. Normative Reference

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

11.2. Informative References

[Bra00] Stefan A.: Rethinking Public Key Infrastructures and Digital Certificates : Building in Privacy, MIT Press; ISBN: 0262024918; 1st edition, August, 2000

[Cha85] Chaum, David: Security without Identification, Card Computers to make Big Brother Obsolete. Original Version appeared in: Communications of the ACM, vol. 28 no. 10, October 1985 pp. 1030-1044. Revised version available at <http://www.chaum.com/articles/>

[ISO99] ISO99: ISO IS 15408, 1999, <http://www.commoncriteria.org/>.

[OECD] OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data, <http://www.oecd.org>.

[Pfi01] Pfitzmann, Andreas; Koehntopp, Marit: Anonymity, Unobservability, and Pseudonymity - A Proposal for Terminology; in: H Federrath (Ed.): Designing Privacy Enhancing Technologies; Proc. Workshop on Design Issues in Anonymity and Unobservability; LNCS 2009; 2001; 1-9. Newer versions available at <http://www.koehntopp.de/marit/pub/anon>

12. Authors' Addresses

Jorge R Cuellar
Siemens AG
Corporate Technology
CT IC 3
81730 Munich, Germany

EMail: Jorge.Cuellar@siemens.com

John B. Morris, Jr.
Director, Internet Standards, Technology & Privacy Project
Center for Democracy & Technology
1634 I Street NW, Suite 1100
Washington, D.C. 20006 USA

EMail: jmorris@cdt.org
URI: <http://www.cdt.org>

Deirdre K. Mulligan
Samuelson Law, Technology & Public Policy Clinic
Boalt Hall School of Law
University of California
Berkeley, CA 94720 USA

EMail: dmulligan@law.berkeley.edu
URI: <http://www.law.berkeley.edu/cenpro/samuelson/>

Jon Peterson
NeuStar, Inc.
1800 Sutter St
Suite 5707
Concord, CA 94520 USA

EMail: jon.peterson@neustar.biz
URI: <http://www.neustar.biz/>

James M. Polk
Cisco Systems
2200 East President George Bush Turnpike
Richardson, Texas 75082 USA

EMail: jmpolk@cisco.com

13. Full Copyright Statement

Copyright (C) The Internet Society (2004). 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 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.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

