

Network Working Group
Request for Comments: 3694

Category: Informational Samuelson Law, Technology & Public Policy Clinic
M. Danley
D. Mulligan
J. Morris
Center for Democracy & Technology
J. Peterson
NeuStar
February 2004

Threat Analysis of the Geopriv Protocol

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

This document provides some analysis of threats against the Geopriv protocol architecture. It focuses on protocol threats, threats that result from the storage of data by entities in the architecture, and threats posed by the abuse of information yielded by Geopriv. Some security properties that meet these threats are enumerated as a reference for Geopriv requirements.

Table of Contents

1.	Introduction	3
2.	Habitat of the Geopriv Protocol	3
3.	Motivations of Attackers of Geopriv	4
4.	Representative Attacks on Geopriv	5
4.1.	Protocol Attacks	5
4.1.1.	Eavesdropping and/or Interception	5
4.1.2.	Identity Spoofing	6
4.1.3.	Information Gathering	7
4.1.4.	Denial of Service	8
4.2.	Host Attacks	9
4.2.1.	Data Stored at Servers	9
4.2.2.	Data Stored in Devices	9
4.2.3.	Data Stored with the Viewer	10
4.2.4.	Information Contained in Rules	10
4.3.	Usage Attacks	11
4.3.1.	Threats Posed by Overcollection	11
5.	Countermeasures for Usage Violations	12
5.1.	Fair Information Practices	12
6.	Security Properties of the Geopriv Protocol	13
6.1.	Rules as Countermeasures	13
6.1.1.	Rule Maker Should Define Rules	13
6.1.2.	Geopriv Should Have Default Rules	14
6.1.3.	Location Recipient Should Not Be Aware of All Rules.	14
6.1.4.	Certain Rules Should Travel With the LO	14
6.2.	Protection of Identities	14
6.2.1.	Short-Lived Identifiers May Protect Target's Identity	15
6.2.2.	Unlinked Pseudonyms May Protect the Location Recipients' Identity	15
6.3.	Security During Transmission of Data	15
6.3.1.	Rules May Disallow a Certain Frequency of Requests	15
6.3.2.	Mutual End-Point Authentication	16
6.3.3.	Data Object Integrity & Confidentiality	16
6.3.4.	Replay Protection	16
7.	Security Considerations	16
8.	IANA Considerations	16
9.	Informative References	16
10.	Authors' Addresses	17
11.	Full Copyright Statement	18

1. Introduction

The proliferation of location-based services that integrate tracking and navigation capabilities gives rise to significant privacy and security concerns. Such services allow users to identify their own location as well as determine the location of others. In certain peer-to-peer exchanges, device identification takes place automatically within a defined location perimeter, informing peer devices of a given user's identity and availability. Additionally, records of location exchanges can reveal significant information about the habits, whereabouts, and associations of individual users.

The Geopriv requirements allow the Location Object (LO) to support a wide variety of uses of Location Information (LI); the Geopriv object itself is intended to be technology-neutral, allowing a wide variety of devices to provide LI in the form of an LO. Geopriv also requires that many classes of Viewers be capable of requesting LI from a Location Server. The Geopriv requirements account for circumstances in which the Target has a contractual relationship with the entities that transmit and receive LI and those in which no contract exists. Requiring the Geopriv object to support any technology, Target-Viewer relationship, or underlying legal framework governing LI, complicates the protection of privacy and the security of LI.

This document analyzes threats to LI in transmission and storage. The possibility that the LI will be compromised by these threats varies depending on the circumstances. A server selling location information to potential marketers poses a distinctly lower risk than an outside individual intercepting a Target's present location to commit a physical attack. It is important that these threats are considered as we work towards defining the LO.

Some of the threats discussed in this document may be outside the scope of the Geopriv charter, e.g., threats arising from failure to meet contractual obligations. Nevertheless, a comprehensive discussion of threats is necessary to identify desirable security properties and counter-measures that will improve the security of the LO, and thereby better protect LI.

2. Habitat of the Geopriv Protocol

The Geopriv architecture will be deployed in the open Internet - in a security environment in which potential attackers can inspect packets on the wire, spoof Internet addresses, and launch large-scale denial-of-service attacks. In some architectures, portions of Geopriv traffic (especially traffic between the Location Generator and an initial Location Server) may occur over managed networks that do not interface with the public Internet.

The protocol itself assumes interaction between a number of logical roles, many of which will commonly be implemented in distributed network devices (for a full list of Geopriv roles and entities with definitions, see [1]). The endpoints of the common Geopriv transactions are the Location Generator (the source of location information from the perspective of the network) and the Location Recipient. Both a Location Generator and a Location Recipient may have a relationship with a Location Server; the Location Generator publishes data to a Location Server (which may provide a grooming/filtration function for location information), and the Location Recipient requests and/or receives information from the Location Server. This provides two points where Geopriv information could require protection across the wire. Rules can also be passed over the network from a Rule Holder to a Location Server; this provides another point where the architecture requires security.

It is important to note that Location Generators and Location Recipients may be implemented on low-cost devices for which strong cryptographic security is currently prohibitively expensive computationally.

3. Motivations of Attackers of Geopriv

The most obvious motivation for an attacker of Geopriv is to learn the location of a subject who wishes to keep their position private, or even for authorized Viewers to ascertain location information with a greater degree of precision than the Rule Maker desires. However, there are several other potential motivations that cause concern. Attackers might also wish to prevent a Target's location from being distributed, or to modify or corrupt location information in order to misrepresent the location of the Target, or to redirect the Target's location information to a third party that is not authorized to know this information. Attackers may want to identify the associates of a Target, or learn the habit or routines of a Target. Attackers might want to learn the identity of all of the parties that are in a certain location. Finally, some attackers may simply want to halt the operation of an entire Geopriv system through denial-of-service attacks.

There is also a class of attackers who may be authorized as legitimate participants in a Geopriv protocol exchange but who abuse location information. This includes the distribution or accumulation of location information outside the parameters of agreements between the principals, possibly for commercial purposes or as an act of unlawful surveillance.

4. Representative Attacks on Geopriv

4.1. Protocol Attacks

4.1.1. Eavesdropping and/or Interception

Imagine a location-based computer game, based on traditional hide-and-seek, in which a centralized server provides hints as to the location of the 'hider' to a set of 'seekers'. Seekers are given access to very coarse location data, whereas a single referee is given access to unfiltered and precise location information of the hider. Each seeker has a wireless device (in the Geopriv architecture, a Location Recipient) that feeds them coarse positioning data from the Location Server. The hider carries a device (a Location Generator employing GPS) that transmits location information to the Location Server.

If one of the seekers wished to cheat by attacking the Geopriv protocol, there are a number of ways they could mount such an attack in order to learn the precise location of the hider. They might eavesdrop on one of two network connections - either the connection between the Location Generator and the Location Server, or the connection between the Location Server and the referee's Location Recipient (which receives precise information). They might also attempt to impersonate the referee to the Location Server, in order to receive unfiltered Location Information. Alternatively, they could impersonate the Location Server to the Location Generator carried by the hider, which would also give them access to precise location information. Finally, the cheater could attempt to act as the Rule Maker, whereby providing Rules to the Location Server would enable the cheater's Location Recipient access to uncoarsened location information.

From these threats, we can derive a need for several security properties of the architecture.

- o Confidentiality is required on both the connection between the Location Generator and the Location Server, as well as the connection between the Location Server and any given Location Recipient.
- o Location Servers must be capable of authenticating and authorizing Location Recipients to prevent impersonation.
- o Similarly, Location Generators must be capable of authenticating and authorizing Location Servers in order to prevent impersonation.

- o Finally, the Location Server must be able to authenticate Rule Makers, to make sure that unauthorized parties cannot change rules.

4.1.1.2. Identity Spoofing

Consider a case in which the same boss employs two rivals. One goes on a business trip to Cleveland. Both rivals carry devices that are tracked by a Location Generator (such as cell phones which the cell carrier can triangulate), and both rivals allow their boss access to their (coarse) location information. The rival that remained home wants to hack the Geopriv protocol to make it appear that the traveling rival is actually goofing off in South Beach rather than attending a dull technology conference in Cleveland. How would such an attack be mounted?

The attacker might attempt to spoof network traffic from the Location Generator to the Location Server (especially if, through some other means such as a denial-of-service attack, the Location Generator became unable to issue its own reports). The goal of the attacker may be to provide falsified location information appropriate for someone in Miami, or perhaps even to replay a genuine location object from a previous visit of the rival to Miami. The attacker might also try to spoof traffic from the Location Server to the boss' Location Recipient.

From these threats we can derive a need for several security properties of the architecture.

- o There is a need for the Location Server to authenticate Location Generators.
- o Location Recipients must be capable of authenticating Location Servers.
- o Location information must be protected from replay attacks.

Identity spoofing may create additional threats when the protocol is attacked. In many circumstances, the identity of the Viewer is the basis for controlling whether LI is revealed and, if so, how that LI is filtered. If the identity of that entity is compromised, privacy is threatened. Anyone inside or outside the transaction that is capable of impersonating an authorized entity can gain access to confidential information, or initiate false transmissions in the authorized entity's name. The ability to spoof the identity of the Location Recipient, for example, would create the risk of an unauthorized entity accessing both the identity and the location of the Target at the moment the LO was sent.

4.1.3. Information Gathering

Eavesdropping and interception can also create traffic analysis threats as the interceptor collects more data over time. Traffic analysis threats are leveraged by an eavesdropper to determine, from the very fact of a network transmission, the relationship between the various entities involved. If an employer sends the location of an employee to a customer, an eavesdropper could determine that these three entities are somehow interacting with one another. If eavesdropping continues over time, the collection of interactions would involve the employer, employees, and all of their customers. Such a log of information would reveal that the employer and employee frequently were associated with one another, and would reveal which clients more frequently dealt with the pair. Thus, the traffic analysis threat creates the risk of eavesdroppers determining the Target's associates.

Traffic analysis might also allow an eavesdropper to ascertain the identity or characteristics of targets in a particular location. By observing transmissions between Location Generators in a particular location and Location Servers (perhaps by eavesdropping on a wireless or wireline LAN scoped to the location in question), and then possibly following the data to various Location Recipients, an attacker may be able to learn the associates, including the employer, of targets in that location, and perhaps to extrapolate further identity information.

If the eavesdropper is able to intercept not only an encrypted LO, but the plaintext LI itself, other threats are raised. Let's return to the above example of the employer requesting an employee's location information. In this instance, the interception of one such past transaction may reveal the identities and/or locations of all three parties involved, in addition to revealing their association. In circumstances where there is a log of this data, however, analysis could reveal any regular route that the employee may travel in visiting customers, a general area that the employee works in, the identities and location of the employee's entire customer base, and information about how the entities relate.

Threats based on traffic analysis are difficult to meet with protocol security measures, but they are important to note.

From these threats we can derive a need for several security properties of the architecture.

- o The Rule Maker must be able to define Rules regarding the use of their LI.

- o The connection between the Location Generator and Location Server, as well as the connection between the Location Server and Location Recipient must remain confidential.
- o Location Servers must be capable of authenticating Location Recipients to prevent impersonation.
- o Location Servers must be able to authenticate Rule Makers to ensure that unauthorized entities cannot change rules.

4.1.4. Denial of Service

Parties who wish to deprive entire networks of Geopriv service, rather than just targeting particular users, would probably focus their efforts on the Location Server. Since in many scenarios the Location Server plays the central role of managing access to location information for many devices, it is in such architectures a natural single point of failure.

The Geopriv protocol appears to have some opportunities for amplification attacks. When the Location Generator publishes location information, the Location Server acts as an exploder, potentially delivering this information to numerous targets. If the Location Generator were to provide very rapid updates of position (as many as link speed could accommodate, especially in high-bandwidth wireless environments), then were the Location Server to proxy information to Seekers at a similar rate, this could become problematic when large numbers of Seekers are tracking the same user.

Also note that most operations associated with the Location Server probably require cryptographic authentication. Cryptographic operations entail a computational expense on the part of the Location Server. This could provide an attractive means for attackers to flood the Location Server with dummied Geopriv information that is spoofed to appear to come from a Location Generator, Location Recipient, or the Rule Maker. Because the Location Server has to expend resources to verify credentials presented by these Geopriv messages, floods of Geopriv information could have greater impact than denial-of-service attacks based on generic packet flooding.

From these threats we can derive a need for several security properties of the architecture.

- o Location Servers must use stateless authentication challenges and similar measures to ensure that authentication attempts will not unnecessarily consume system resources.

- o The Rule Maker must be able to provision policies that limit the rate at which Location Information is sent to prevent amplification attacks.

4.2. Host Attacks

4.2.1. Data Stored at Servers

LI maintained at a server is subject to many potential risks. First, there may be accidental misuse of LI by the server. Whether by negligence, carelessness, or lack of knowledge, the server may accidentally release LI to the wrong Location Recipients, or fail to properly filter the LI that is sent out. Second, the server may intentionally misuse LI. A server may decide to sell a "profile" it has compiled of a Target or Location Recipient despite provisions to the contrary in the Rule Maker's Rule. Alternatively, an individual working for the server may, for personal gain, misuse access to the server to obtain LI. Third, even with the most secure and trusted server, there is the risk that someone outside the system will hack into it in order to retrieve LI. Last, there is always the potential that someone would use the legal system to subpoena an individual's records from a Server. Such a process would likely result in the revelation of the Target's location information without notice to the Target or the Target's consent.

Data stored at the server may reveal the Target's present location if the data is used or intercepted at or near the moment of transmission. If a Target requests a map from their present location to a nearby store, and the Location Server sends that information to the wrong Location Recipient, the Viewer could know the identity of the Target, the Target's current location, and the location where the Target might be headed.

Data stored at the Location Server can also create many of the traffic analysis threats discussed in Section 4.1 above. If access is gained not only to the fact of the LO transmission, but also to the LI transmitted, anyone with access to that information can put together a history of where that Target has been, for how long, and with whom.

4.2.2. Data Stored in Devices

Because Geopriv is required to work with any given type of technology or Device, it is difficult to determine the particular threat potential of individual devices. For example, any device that maintains a log of location requests sent, or LOs received, would

pose a similar threat to the information maintained at a Location Server, discussed above. A court subpoena or warrant for an individual's device could additionally reveal a similar log.

Additionally, depending on the device, there is always the potential for data to be compromised in some way. For a Device with a screen, there is always the potential that another individual will have the opportunity to view the Device display without the user's knowledge. A Device that provides verbal feedback (i.e., to give directions to the blind) creates additional potential for LI to be compromised. If the Target/Viewer is sitting in a public place and requests directions from the Target's home to another location, anyone who can hear the Device output may be able to determine the Target's identity, their residence, and possibly the location to which they are headed.

In addition, if the device retained location information and the Device were lost or stolen, someone other than the Rule Maker could potentially access information regarding who LI was sent to and when, as well as potentially the location of the Target during each transaction. Such information could enable an entity to determine significant private information based on who the owner of the Device has associated with in the past, as well as each location where the Target has been and for how long.

4.2.3. Data Stored with the Viewer

The threats posed here are similar to those discussed above in relation to Location Servers and Devices. The main purpose of separating out threats posed by data stored at the Viewer is to show that, depending on the complexity of the transaction and the other entities involved, data storage at various points in the transaction can bring rise to the same types of privacy risks.

4.2.4. Information Contained in Rules

In many instances, the Rules a Rule Maker creates will reveal information either about the Rule Maker or the Target. A rule that degrades all information sent out by approximately 25 miles might tell an interceptor how to determine the Target's true location. A Rule that states, "Tell my boss what room I'm in when I'm in the building, but when I'm outside the building between 9 a.m. and 5 p.m. tell him I'm in the building," would reveal a lot more information than most employees would desire. Any boss who was the Location Recipient who received LI that specified "in the building" would then realize that the employee was elsewhere.

In addition, if an entity had access to a log of data at the Location Server or at a Device, knowledge of the content of Rules would enable a sort of "decoding" of the location information of the device to something more accurate. Thus, my boss could not only tell where I am at this minute, but could tell how many times over the last year I had been outside the building between 9 a.m. and 5 p.m.

The Rules themselves may also reveal information about the Target. A rule such as the one above would clearly reveal the employment relationship between the two individuals, as well as the fact that the employee was hiding something from the employer.

In combination with other information, the location information may enable the identification of the Target.

4.3. Usage Attacks

4.3.1. Threats Posed by Overcollection

Weak or absent default privacy rules would also compromise LI. Without default Rules for LOs, it is likely that a large number of Devices would reveal LI by default. Privacy rules should control the collection, use, disclosure, and retention of Location Information. These rules must comply with fair information practices - these practices are further discussed in Section 5.1.

While technically savvy Device users may create privacy rules to protect their LI, many individuals will lack the skill or motivation to do so. Thus, left to their own devices many individuals would likely be left without privacy rules for their LI. This in turn would leave these users' LI entirely vulnerable to various attacks. Default rules are necessary to address this problem.

Without default rules, for example, a device might signal out to anyone nearby at regular intervals, respond to anyone nearby who queried it, or send signals out to unknown entities.

The lack of a default rule of "Do not re-distribute," would allow the Location Server to pass the Target's location information on to others. Lack of a default rule limiting the retention of LI could increase the risk posed by inappropriate use and access to stored data.

While defining default privacy rules is beyond the scope of this document, default rules are necessary to limit the privacy risks posed by the use of services and devices using LI.

5. Countermeasures for Usage Violations

5.1. Fair Information Practices

Principles of fair information practices require entities that handle personal information to meet certain obligations with respect to its collection, use, maintenance and security, and give individuals whose personal information is collected certain due process-like rights in the handling of their information. Fair information practices are designed to prevent specific threats posed by the collection of personal information about individuals. For this reason, fair information practices are "countermeasures" that should be reflected in technical systems that handle personal information and the Rules that govern their use. A brief discussion of fair information practices may be beneficial in formulating requirements for the LO.

There are seven main principles of fair information practices:

1. **Openness:** The existence of a record-keeping system for personal information must be known, along with a description of the main purpose and uses of the data. Thus, any entity that collects LI should inform individuals that this information is being collected and inform them about what the LI is being used for. Openness is designed to prevent the creation of secret systems.
2. **Individual Participation:** Individuals should have a right to view all information collected about them, and to be able to correct or remove data that is not timely, accurate, relevant, or complete. The practice of individual participation acknowledges that sometimes information that is collected may be inaccurate or inappropriate.
3. **Collection Limitation:** Data should be collected by lawful and fair means and should be collected, where appropriate, with the knowledge or consent of the subject. Data collection should be minimized to that which is necessary to support the transaction. Placing limits on collection helps protect individuals from the dangers of overcollection - both in terms of collecting too much information, or of collecting information for too long of a time period.
4. **Data Quality:** Personal data should be relevant to the purposes for which it is collected and used; personal information should be accurate, complete, and timely. The requirement of data quality is designed to prevent particular kinds of harms that can flow from the use (appropriate or inappropriate) of personal information.

5. Finality: There should be limits to the use and disclosure of personal data: data should be used only for purposes specified at the time of collection; data should not be otherwise used or disclosed without the consent of the data subject or other legal authority. A consumer who provides LI to a business in order to receive directions, for example, does not provide that information for any other purpose. The business should then only use that LI to provide directions, and not for other purposes.
6. Security: Personal Data should be protected by reasonable security safeguards against such risks as loss, unauthorized access, destruction, use, modification, or disclosure. While some security measures may take place outside of the LO (i.e., limiting employee access to Location Servers), other measures may be done through the LO or LO applications.
7. Accountability: Record keepers should be accountable for complying with fair information practices. It will typically be easier for an individual to enforce these practices if they are explicitly written - either in the Rules written by the Rule Maker, or in contracts between the individual and a trusted entity.

6. Security Properties of the Geopriv Protocol

The countermeasures suggested below reflect the threats discussed in this document. There is thus some overlap between the proposed security properties listed below, and the requirements in [1].

6.1. Rules as Countermeasures

The sections below are designed to illustrate that in many instances threats to LI can be limited through clear, unavoidable rules determined by Rule Makers.

6.1.1. Rule Maker Should Define Rules

The Rule Maker for a given Device will generally be either the user of, or owner of, the Device. In certain circumstances, the Rule Maker may be both of these entities. Depending on the device, the Rule Maker may or may not be the individual most closely aligned with the Target. For instance, a child carrying a cell phone may be the Target, but the parent of that child would likely be the Rule Maker for the Device. Giving the Rule Maker control is a potential opportunity to buttress the consent component of the collection limitation and finality principles discussed above.

6.1.2. Geopriv Should Have Default Rules

Because some Rule Makers may not be informed about the role Rules play in the disclosure of their LI, Geopriv should include default Rules. The Rule Maker is, of course, always free to change his or her Rules to provide more or less protection. To protect privacy and physical safety, default Rules should, at a minimum, limit disclosure and retention of LI.

Default Rules are also necessary for so-called "dumb" Location Generators (LG). If a LG is unable to determine the Rules set by the Rule Maker before publishing the LO on to a Location Server, it is important that some default Rules protect that LO in transit, and ensure that the LO is eventually only sent to authorized Location Recipients. These default LG Rules would help prevent many of the threats discussed in this document. The Rule Maker should be able to determine the content of these default Rules at any time.

6.1.3. Location Recipient Should Not Be Aware of All Rules

A Viewer should not be aware of the full Rules defined by the Rule Maker. The Viewer will only need to be aware of those Rules it must obey (i.e., those regarding its use and retention of the LI). Other Rules, such as those specifying the accuracy or filtering of the LI, or rules that do not cover the given interaction should not be revealed to the Viewer. This countermeasure is consistent with the minimization component of the collection limitation principle and ensures that the Rule Maker reveals only what he intends to reveal.

6.1.4. Certain Rules Should Travel With the LO

Security of LI at the device level is a bit complicated, as the Rule Maker has no real control over what is done with the LI once it arrives at the Location Recipient. If certain Rules travel with the LO, the Rule Maker can encourage Viewer compliance with its Rules. Potentially, a Rule could travel with the LO indicating when it was time to purge the data, preventing the compilation of a "log" of the Target's LI on any Device involved in the transmission of the LO. Allowing Rules to travel with the LO has the potential to limit the opportunity for traffic analysis attacks.

6.2. Protection of Identities

Identities are an extremely important component of the LO. While, in many instances, some form of identification of the Target, Rule Maker, and Viewer will be necessary for authentication, there are various methods to separate these authentication "credentials" from the true identity of these devices. These countermeasures are

particularly useful in that compromise of a log of LI, no matter where the source, is less threatening to privacy when the Target's identity is stripped.

6.2.1. Short-Lived Identifiers May Protect Target's Identity

Short-Lived identifiers would allow the using protocol to hide the true identity of the Rule Maker and the Target from Location Servers or Location Recipients. These identifiers would still allow authentication, ensuring that only appropriate Location Recipients received the LO. At the same time, however, making these identifiers short-lived helps prevent any association of a true identity of a Target with particular habits and associates.

6.2.2. Unlinked Pseudonyms May Protect the Location Recipients' Identity

Unlinked pseudonyms would protect the identity of the Location Recipients in much the same manner as short-lived identifiers would protect the Target's identity. When using both, any record that a Location Server had of a transaction would have two "credentials" associated with an LI transmission: one linked to the Target and one linked to the Location Recipient. These credentials would allow the Location Server to authenticate the transmission without ever acquiring knowledge of the true identities of the individuals associated with each side of the transaction.

6.3. Security During Transmission of Data

The attacks described in this document motivate the following security properties for the connections between the Location Generator and Location Server, the Location Server and Rule Maker, and the Location Server and Location Recipient:

6.3.1. Rules May Disallow a Certain Frequency of Requests

The Rule Maker might be able to set a Rule that disallows a certain number of requests made within a specific period of time. This type of arrangement would allow the Rule Maker to somewhat prevent attackers from detecting patterns in randomly coarsened data. To an "untrusted" Location Recipient, for example, to whom the Rule Maker only wants to reveal LI that is coarsened to the level of a city, only one request might be honored every 2 hours. This would prevent Location Recipients from sending repeated requests to gain more accurate presence information.

Similarly, thresholds on notifications of location information can help to combat amplification attacks.

6.3.2. Mutual End-Point Authentication

Authentication is crucial to the security of LI during transmission. The Location Server must be capable of authenticating Location Recipients to prevent impersonation. Location Generators must be capable of authenticating Location Servers to ensure that raw location information is not sent to improper entities. Additionally, Location Servers must be able to authenticate Rule Makers to ensure that unauthorized entities cannot change Rules.

6.3.3. Data Object Integrity & Confidentiality

The LO must maintain integrity at all points of communication between Location Servers and Location Recipients. Confidentiality is required on both the connection between the Location Generator and the Location Server, as well as on the connection between the Location Server and any given Location Recipient. Confidentiality of Rules sent over the network to the Location Server is of comparable importance.

6.3.4. Replay Protection

Replay protection prevents an attacker from capturing a particular piece of location information and replaying it at a later time in order to convince Viewers of an erroneous location for the target. Both Location Recipients and Location Servers, depending on their capabilities, may need replay protection.

7. Security Considerations

This informational document characterizes potential security threats targeting the Geopriv architecture.

8. IANA Considerations

This document introduces no additional considerations for IANA.

9. Informative References

- [1] Cuellar, J., Morris, J., Mulligan, D., Peterson, J. and J. Polk, "Geopriv Requirements", RFC 3693, January 2004.

10. Authors' Addresses

Michelle Engelhardt Danley
Samuelson Law, Technology & Public Policy Clinic
Boalt Hall School of Law
University of California
Berkeley, CA 94720
USA

Email: mre213@nyu.edu
URI: <http://www.law.berkeley.edu/cenpro/samuelson/>

Deirdre 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/>

John B. Morris, Jr.
Center for Democracy & Technology
1634 I Street NW
Suite 1100
Washington, DC 20006
USA

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

Jon Peterson
NeuStar, Inc.
1800 Sutter St
Suite 570
Concord, CA 94520
USA

Phone: +1 925/363-8720
Email: jon.peterson@neustar.biz
URI: <http://www.neustar.biz/>

11. 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.

