

Internet Engineering Task Force (IETF)
Request for Comments: 6772
Category: Standards Track
ISSN: 2070-1721

H. Schulzrinne, Ed.
Columbia University
H. Tschofenig, Ed.
Nokia Siemens Networks
J. Cuellar
Siemens
J. Polk
Cisco
J. Morris

M. Thomson
Microsoft
January 2013

Geolocation Policy: A Document Format for
Expressing Privacy Preferences for Location Information

Abstract

This document defines an authorization policy language for controlling access to location information. It extends the Common Policy authorization framework to provide location-specific access control. More specifically, this document defines condition elements specific to location information in order to restrict access to data based on the current location of the Target.

Furthermore, this document defines two algorithms for reducing the granularity of returned location information. The first algorithm is defined for usage with civic location information, whereas the other one applies to geodetic location information. Both algorithms come with limitations. There are circumstances where the amount of location obfuscation provided is less than what is desired. These algorithms might not be appropriate for all application domains.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc6772>.

Copyright Notice

Copyright (c) 2013 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	4
2.	Terminology	5
3.	Generic Processing	7
3.1.	Structure of Geolocation Authorization Documents	7
3.2.	Rule Transport	7
4.	Location-Specific Conditions	7
4.1.	Geodetic Location Condition Profile	8
4.2.	Civic Location Condition Profile	9
5.	Actions	9
6.	Transformations	9
6.1.	Set Retransmission-Allowed	9
6.2.	Set Retention-Expiry	10
6.3.	Set Note-Well	10
6.4.	Keep Ruleset Reference	10
6.5.	Provide Location	11
6.5.1.	Civic Location Profile	12
6.5.2.	Geodetic Location Profile	13
7.	Examples	15
7.1.	Rule Example with Civic Location Condition	15
7.2.	Rule Example with Geodetic Location Condition	16
7.3.	Rule Example with Civic and Geodetic Location Condition	17
7.4.	Rule Example with Location-Based Transformations	18
7.5.	Location Obfuscation Example	19
8.	XML Schema for Basic Location Profiles	23
9.	XML Schema for Geolocation Policy	24
10.	XCAP Usage	25
10.1.	Application Unique ID	26
10.2.	XML Schema	26
10.3.	Default Namespace	26
10.4.	MIME Media Type	26
10.5.	Validation Constraints	26
10.6.	Data Semantics	26
10.7.	Naming Conventions	26
10.8.	Resource Interdependencies	26
10.9.	Authorization Policies	27
11.	IANA Considerations	27
11.1.	Geolocation Policy XML Schema Registration	27
11.2.	Geolocation Policy Namespace Registration	27
11.3.	Geolocation Policy Location Profile Registry	28
11.4.	Basic Location Profile XML Schema Registration	28
11.5.	Basic Location Profile Namespace Registration	29
11.6.	XCAP Application Usage ID	29
12.	Internationalization Considerations	30
13.	Security Considerations	30
13.1.	Introduction	30
13.2.	Obfuscation	31

13.3. Algorithm Limitations 32
 13.4. Usability 33
 13.5. Limitations of Obscuring Locations 33
 14. References 35
 14.1. Normative References 35
 14.2. Informative References 35
 Appendix A. Acknowledgments 38
 Appendix B. Pseudocode 39

1. Introduction

Location information needs to be protected against unauthorized access to preserve the privacy of humans. In RFC 6280 [RFC6280], a protocol-independent model for access to geographic information is defined. The model includes a Location Generator (LG) that determines location information, a Location Server (LS) that authorizes access to location information, a Location Recipient (LR) that requests and receives location information, and a Rule Maker (RM) that writes authorization policies. An authorization policy is a set of rules that regulates an entity’s activities with respect to privacy-sensitive information, such as location information.

The data object containing location information in the context of this document is referred to as a Location Object (LO). The basic rule set defined in the Presence Information Data Format Location Object (PIDF-LO) [RFC4119] can restrict how long the Location Recipient is allowed to retain the information, and it can prohibit further distribution. It also contains a reference to an enhanced rule set and a human-readable privacy policy. The basic rule set does not protect access to location information. It only conveys the user’s privacy preferences. This document describes an enhanced rule set that provides richer constraints on the distribution of LOs.

The enhanced rule set allows the entity that uses the rules defined in this document to restrict the retention and to enforce access restrictions on location data, including prohibiting any dissemination to particular individuals, during particular times or when the Target is located in a specific region. The RM can also stipulate that only certain parts of the Location Object are to be distributed to recipients or that the resolution is reduced for parts of the Location Object.

In the typical sequence of operations, a Location Server receives a query for location information for a particular Target. The authenticated identity of the Location Recipient, together with other information provided with the request or generally available to the server, is then used for searching through the rule set. If more than one rule matches the condition element, then the combined

permission is evaluated according to the description in Section 10 of [RFC4745]. The result of the rule evaluation is applied to the location information, yielding a possibly modified Location Object that is delivered to the Location Recipient.

This document does not describe the protocol used to convey location information from the Location Server to the Location Recipient.

This document extends the Common Policy framework defined in [RFC4745]. That document provides an abstract framework for expressing authorization rules. As specified there, each such rule consists of conditions, actions, and transformations. Conditions determine under which circumstances the entity executing the rules, such as a Location Server, is permitted to apply actions and transformations. In a location information context, transformations regulate how a Location Server modifies the information elements that are returned to the requestor by, for example, reducing the granularity of returned location information.

This document defines two algorithms for reducing the granularity of returned location information. The first algorithm is defined for usage with civic location information (see Section 6.5.1) while the other one applies to geodetic location information (see Section 6.5.2). Both algorithms come with limitations, i.e., they provide location obfuscation under certain conditions and may therefore not be appropriate for all application domains. These limitations are documented within the Security Consideration section (see Section 13). The geodetic transformation algorithm in Section 6.5.2 mitigates privacy risks for both stationary and moving Targets. However, moving Targets will reveal additional information to an adversary. To cover applications that have more sophisticated privacy requirements, additional algorithms may need to be defined. This document foresees extensions in the form of new algorithms and therefore defines a registry (see Section 11.3).

The XML schema defined in Section 9 extends the Common Policy schema by introducing new child elements to the condition and transformation elements. This document does not define child elements for the action part of a rule.

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 RFC 2119 [RFC2119].

This document reuses the terminology of RFC 6280 [RFC6280], such as Location Server (LS), Location Recipient (LR), Rule Maker (RM), Target, Location Generator (LG), and Location Object (LO). This document uses the following terminology:

Presentity or Target:

RFC 6280 [RFC6280] uses the term "Target" to identify the object or person of which location information is required. The presence model described in RFC 2778 [RFC2778] uses the term "presentity" to describe the entity that provides presence information to a presence service. A presentity in a presence system is a Target in a location information system.

Watcher or Location Recipient:

The receiver of location information is the Location Recipient (LR) in the terminology of RFC 6280 [RFC6280]. A watcher in a presence system, i.e., an entity that requests presence information about a presentity, is a Location Recipient in a location information system.

Authorization policy:

An authorization policy is given by a rule set. A rule set contains an unordered list of (policy) rules. Each rule has a condition, an action, and a transformation component.

Permission:

The term "permission" refers to the action and transformation components of a rule.

Location Servers:

Entities that evaluate the geolocation authorization policies.

Presence Servers:

The geolocation privacy architecture is, as described in RFC 4079 [RFC4079], aligned with the presence architecture, and a "Presence Server" is therefore an entity that distributes location information along with other presence-specific XML data elements.

3. Generic Processing

3.1. Structure of Geolocation Authorization Documents

A geolocation authorization document is an XML document, formatted according to the schema defined in [RFC4745]. Geolocation authorization documents inherit the media type of Common Policy documents, application/auth-policy+xml. As described in [RFC4745], this document is composed of rules that contain three parts: conditions, actions, and transformations. Each action or transformation, which is also called a permission, has the property of being a positive grant of information to the Location Recipient. As a result, there is a well-defined mechanism for combining actions and transformations obtained from several sources. This mechanism is privacy enabling, since the lack of any action or transformation can only result in less information being presented to a Location Recipient.

3.2. Rule Transport

There are two ways the authorization rules described in this document may be conveyed between different parties:

- o RFC 4119 [RFC4119] allows enhanced authorization policies to be referenced via a Uniform Resource Locator (URL) in the 'ruleset-reference' element. The 'ruleset-reference' element is part of the basic rules that always travel with the Location Object.
- o Authorization policies might, for example, also be stored at a Location Server / Presence Server. The Rule Maker therefore needs to use a protocol to create, modify, and delete the authorization policies defined in this document. Such a protocol is available with the Extensible Markup Language (XML) Configuration Access Protocol (XCAP) [RFC4825].

4. Location-Specific Conditions

This section describes the location-specific conditions of a rule. The <conditions> element contains zero or more <location-condition> child element(s). The <conditions> element only evaluates to TRUE if all child elements evaluate to TRUE; therefore, multiple <location-condition> elements are not normally useful.

The <location-condition> element MUST contain at least one <location> child element. The <location-condition> element evaluates to TRUE if any of its child <location> elements matches the location of the Target, i.e., <location> elements are combined using a logical OR.

The three attributes of <location> are 'profile', 'xml:lang', and 'label'. The 'profile' indicates the location profile that is included as child elements in the <location> element. Two location profiles, geodetic and civic, are defined in Sections 4.1 and 4.2. Each profile describes under what conditions a <location> element evaluates to TRUE.

The 'label' attribute allows a human-readable description to be added to each <location> element. The 'xml:lang' attribute contains a language tag providing further information for rendering of the content of the 'label' attribute.

The <location-condition> and the <location> elements provide extension points. If an extension is not understood by the entity evaluating the rules, then this rule evaluates to FALSE. This causes a <conditions> element to evaluate to FALSE if a <location-condition> element is unsupported. A <location-condition> is considered TRUE if any of the <location> elements understood by the rule evaluator is TRUE.

4.1. Geodetic Location Condition Profile

The geodetic location profile is identified by the token 'geodetic-condition'. Rule Makers use this profile by placing a Geography Markup Language [GML] <Circle> element within the <location> element (as described in Section 5.2.3 of [RFC5491]).

The <location> element containing the information for the geodetic location profile evaluates to TRUE if the current location of the Target is completely within the described location (see Section 6.1.15.3 of [OGC-06-103r4]). Note that the Target's actual location might be represented by any of the location shapes described in [RFC5491]. If the geodetic location of the Target is unknown, then the <location> element containing the information for the geodetic location profile evaluates to FALSE.

Implementations MUST support the World Geodetic System 1984 (WGS 84) [NIMA.TR8350.2-3e] coordinate reference system using the formal identifier from the European Petroleum Survey Group (EPSG) Geodetic Parameter Dataset (as formalized by the Open Geospatial Consortium (OGC)):

2D: WGS 84 (latitude, longitude), as identified by the URN "urn:ogc:def:crs:EPSG::4326". This is a two-dimensional CRS.

A Coordinate Reference System (CRS) MUST be specified using the above URN notation only; implementations do not need to support user-defined CRSs.

Implementations MUST specify the CRS using the "srsName" attribute on the outermost geometry element. The CRS MUST NOT be changed for any sub-elements. The "srsDimension" attribute MUST be omitted, since the number of dimensions in these CRSs is known.

4.2. Civic Location Condition Profile

The civic location profile is identified by the token 'civic-condition'. Rule Makers use this profile by placing a <civicAddress> element, defined in [RFC5139], within the <location> element.

All child elements of a <location> element that carry <civicAddress> elements MUST evaluate to TRUE (i.e., logical AND) in order for the <location> element to evaluate to TRUE. For each child element, the value of that element is compared to the value of the same element in the Target's civic location. The child element evaluates to TRUE if the two values are identical based on an octet-by-octet comparison.

A <location> element containing a <civic-condition> profile evaluates to FALSE if a civic address is not present for the Target. For example, this could occur if location information has been removed by other rules or other transmitters of location information or if only the geodetic location is known. In general, it is RECOMMENDED behavior for an LS not to apply a translation from geodetic location to civic location (i.e., geocode the location).

5. Actions

This document does not define location-specific actions.

6. Transformations

This document defines several elements that allow Rule Makers to specify transformations that

- o reduce the accuracy of the returned location information, and
- o set the basic authorization policies carried inside the PIDF-LO.

6.1. Set Retransmission-Allowed

This element specifies a change to or the creation of a value for the <retransmission-allowed> element in the PIDF-LO. The data type of the <set-retransmission-allowed> element is a boolean.

If the value of the <set-retransmission-allowed> element is set to TRUE, then the <retransmission-allowed> element in the PIDF-LO MUST be set to TRUE. If the value of the <set-retransmission-allowed>

element is set to FALSE, then the <retransmission-allowed> element in the PIDF-LO MUST be set to FALSE.

If the <set-retransmission-allowed> element is absent, then the value of the <retransmission-allowed> element in the PIDF-LO MUST be kept unchanged, or if the PIDF-LO is created for the first time, then the value MUST be set to FALSE.

6.2. Set Retention-Expiry

This transformation asks the LS to change or set the value of the <retention-expiry> element in the PIDF-LO. The data type of the <set-retention-expiry> element is a non-negative integer.

The value provided with the <set-retention-expiry> element indicates seconds, and these seconds are added to the time that the LS provides location. A value of zero requests that the information is not retained.

If the <set-retention-expiry> element is absent, then the value of the <retention-expiry> element in the PIDF-LO is kept unchanged, or if the PIDF-LO is created for the first time, then the value MUST be set to the current date.

6.3. Set Note-Well

This transformation asks the LS to change or set the value of the <note-well> element in the PIDF-LO. The data type of the <set-note-well> element is a string.

The value provided with the <set-note-well> element contains a privacy statement as a human-readable text string, and an 'xml:lang' attribute denotes the language of the human-readable text.

If the <set-note-well> element is absent, then the value of the <note-well> element in the PIDF-LO is kept unchanged, or if the PIDF-LO is created for the first time, then no content is provided for the <note-well> element.

6.4. Keep Ruleset Reference

This transformation specifies whether the <external-ruleset> element in the PIDF-LO carries the extended authorization rules defined in [RFC4745]. The data type of the <keep-rule-reference> element is boolean.

If the value of the <keep-rule-reference> element is set to TRUE, then the <external-ruleset> element in the PIDF-LO is kept unchanged

when included. If the value of the <keep-rule-reference> element is set to FALSE, then the <external-ruleset> element in the PIDF-LO MUST NOT contain a reference to an external rule set. The reference to the ruleset is removed, and no rules are carried as MIME bodies (in case of Content-ID (cid:) URIs [RFC2392]).

If the <keep-rule-reference> element is absent, then the value of the <external-ruleset> element in the PIDF-LO is kept unchanged when available, or if the PIDF-LO is created for the first time, then the <external-ruleset> element MUST NOT be included.

6.5. Provide Location

The <provide-location> element contains child elements of a specific location profile that controls the granularity of returned location information. This form of location granularity reduction is also called 'obfuscation' and is defined in [DUCKHAM05] as

the means of deliberately degrading the quality of information about an individual's location in order to protect that individual's location privacy.

Location obscuring presents a number of technical challenges. The algorithms provided in this document are provided as examples only. A discussion of the technical constraints on location obscuring is included in Section 13.5.

The functionality of location granularity reduction depends on the type of location provided as input. This document defines two profiles for reduction, namely:

- o civic-transformation: If the <provide-location> element has a <provide-civic> child element, then civic location information is disclosed as described in Section 6.5.1, subject to availability.
- o geodetic-transformation: If the <provide-location> element has a <provide-geo> child element, then geodetic location information is disclosed as described in Section 6.5.2, subject to availability.

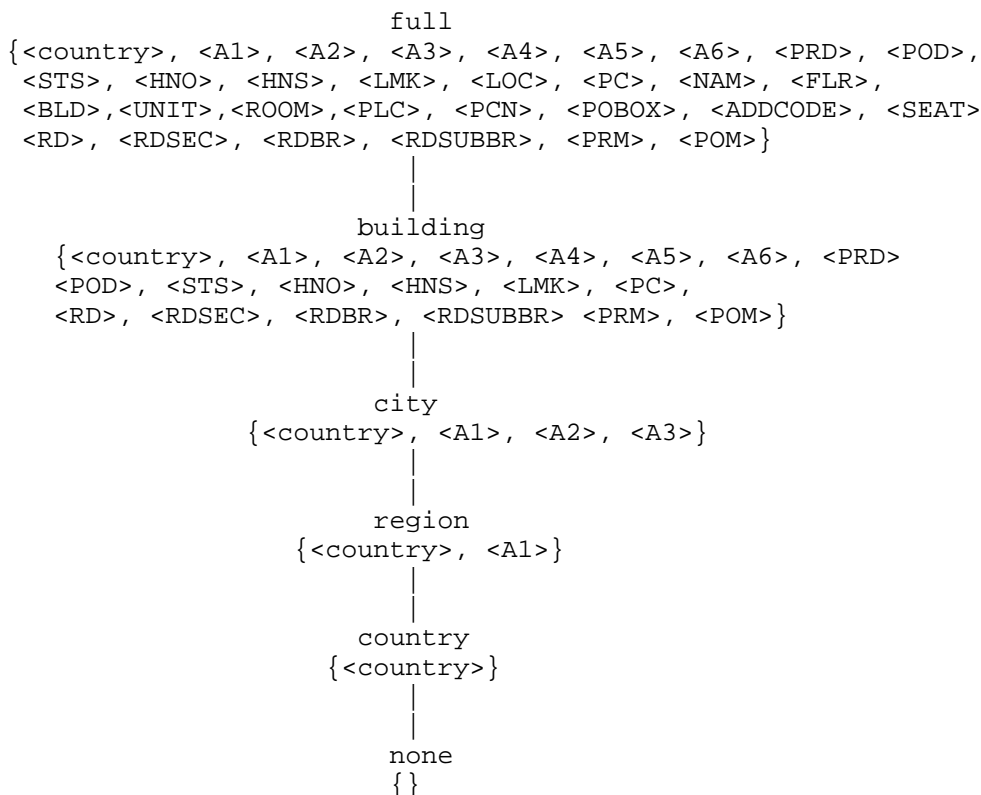
The <provide-location> element MUST contain the 'profile' attribute if it contains child elements, and the child elements MUST be appropriate for the profile.

If the <provide-location> element has no child elements, then civic as well as geodetic location information is disclosed without reducing its granularity, subject to availability. In this case, the profile attribute MUST NOT be included.

6.5.1. Civic Location Profile

This profile uses the token 'civic-transformation'. This profile allows civic location transformations to be specified by means of the <provide-civic> element that restricts the level of civic location information the LS is permitted to disclose. The symbols of these levels are: 'country', 'region', 'city', 'building', and 'full'. Each level is given by a set of civic location data items such as <country> and <A1>, ..., <POM>, as defined in [RFC5139]. Each level includes all elements included by the lower levels.

The 'country' level includes only the <country> element; the 'region' level adds the <A1> element; the 'city' level adds the <A2> and <A3> elements; the 'building' level and the 'full' level add further civic location data as shown below.



The default value is "none".

The schema of the <provide-civic> element is defined in Section 8.

6.5.2. Geodetic Location Profile

This profile uses the token 'geodetic-transformation' and refers only to the Coordinate Reference System (CRS) WGS 84 (urn:ogc:def:crs:EPSG::4326, 2D). This profile allows geodetic location transformations to be specified by means of the <provide-geo> element that may restrict the returned geodetic location information based on the value provided in the 'radius' attribute. The value of the 'radius' attribute expresses the radius in meters.

The schema of the <provide-geo> element is defined in Section 8.

The algorithm proceeds in six steps. The first two steps are independent of the measured position to be obscured and should be run only once or very infrequently for each region and desired uncertainty. The steps are:

1. Choose a geodesic projection with Cartesian coordinates and a surface you want to cover. Limit the worst-case distortion of the map as noted below.
2. Given a desired uncertainty radius "d", choose a grid of so-called "landmarks" at a distance of at least d units apart from each other.
3. Given a measured location $M=(m,n)$ on the surface, calculate its 4 closest landmarks on the grid, with coordinates: SW = (l,b), SE=(r,b), NW=(l,t), NE=(r,t). Thus, $l \leq m < r$ and $b \leq n < t$. See notes below.
4. Let $x=(m-l)/(r-l)$ and $y=(n-b)/(t-b)$.

x and y are thus the scaled local coordinates of the point M in the small grid square that contains it, where x and y range between 0 and 1.

5. Let $p = 0.2887 (= \sqrt{3}/6)$ and $q = 0.7113 (= 1-p)$. Determine which of the following eight cases holds:

- C1. $x < p$ and $y < p$
- C2. $p \leq x < q$ and $y < x$ and $y < 1-x$
- C3. $q \leq x$ and $y < p$
- C4. $p \leq y < q$ and $x \leq y$ and $y < 1-x$
- C5. $p \leq y < q$ and $y < x$ and $1-x \leq y$

C6. $x < p$ and $q \leq y$
C7. $p \leq x < q$ and $x \leq y$ and $1-x \leq y$
C8. $q \leq x$ and $q \leq y$

6. Depending on the case, let C (=Center) be

C1: SW
C2: SW or SE
C3: SE

C4: SW or NW
C5: SE or NE

C6: NW
C7: NW or NE
C8: NE

Return the circle with center C and radius d.

Notes:

Regarding Step 1:

The scale of a map is the ratio of a distance (a straight line) on the map to the corresponding air distance on the ground. For maps covering larger areas, a map projection from a sphere (or ellipsoid) to the plane will introduce distortion, and the scale of the map is not constant. Also, note that the real distance on the ground is taken along great circles, which may not correspond to straight lines on the map, depending on the projection used. Let us measure the (length) distortion of the map as the quotient between the maximal and the minimal scales on the map. The distortion MUST be below 1.5. (The minimum distortion is 1.0: if the region of the map is small, then the scale may be taken as a constant over the whole map).

Regarding Step 3:

SW is mnemonic for southwest, b for bottom, l for left (SW=(l,b)), etc., but the directions of the geodesic projection may be arbitrary, and thus SW may not be southwest of M, but it will be left and below M *on the map*.

7. Examples

This section provides a few examples for authorization rules using the extensions defined in this document.

7.1. Rule Example with Civic Location Condition

This example illustrates a single rule that employs the civic location condition. It matches if the current location of the Target equals the content of the child elements of the <location> element. Requests match only if the Target is at a civic location with country set to 'Germany', state (A1) set to 'Bavaria', city (A3) set to 'Munich', city division (A4) set to 'Perlach', street name (A6) set to 'Otto-Hahn-Ring', and house number (HNO) set to '6'.

No actions and transformation child elements are provided in this rule example. The actions and transformation could include presence-specific information when the Geolocation Policy framework is applied to the Presence Policy framework (see [RFC5025]).

```
<?xml version="1.0" encoding="UTF-8"?>
<ruleset xmlns="urn:ietf:params:xml:ns:common-policy"
  xmlns:gp="urn:ietf:params:xml:ns:geolocation-policy">

  <rule id="AA56i09">
    <conditions>
      <gp:location-condition>
        <gp:location
          profile="civic-condition"
          xml:lang="en"
          label="Siemens Neuperlach site 'Legoland'"
          xmlns="urn:ietf:params:xml:ns:pidf:geopriv10:civicAddr">
          <country>DE</country>
          <A1>Bavaria</A1>
          <A3>Munich</A3>
          <A4>Perlach</A4>
          <A6>Otto-Hahn-Ring</A6>
          <HNO>6</HNO>
        </gp:location>
      </gp:location-condition>
    </conditions>
    <actions/>
    <transformations/>
  </rule>
</ruleset>
```

7.2. Rule Example with Geodetic Location Condition

This example illustrates a rule that employs the geodetic location condition. The rule matches if the current location of the Target is inside the area specified by the polygon. The polygon uses the EPSG 4326 coordinate reference system. No altitude is included in this example.

```
<?xml version="1.0" encoding="UTF-8"?>
<ruleset
  xmlns="urn:ietf:params:xml:ns:common-policy"
  xmlns:gp="urn:ietf:params:xml:ns:geolocation-policy"
  xmlns:gml="http://www.opengis.net/gml"
  xmlns:gs="http://www.opengis.net/pidflo/1.0">

  <rule id="BB56A19">
    <conditions>
      <gp:location-condition>
        <gp:location
          xml:lang="en"
          label="Sydney Opera House"
          profile="geodetic-condition">
          <gs:Circle srsName="urn:ogc:def:crs:EPSG::4326">
            <gml:pos>-33.8570029378 151.2150070761</gml:pos>
            <gs:radius uom="urn:ogc:def:uom:EPSG::9001">1500
            </gs:radius>
          </gs:Circle>
        </gp:location>
      </gp:location-condition>
    </conditions>
    <transformations/>
  </rule>
</ruleset>
```


7.3. Rule Example with Civic and Geodetic Location Condition

This example illustrates a rule that employs a mixed civic and geodetic location condition. Depending on the available type of location information, namely civic or geodetic location information, one of the location elements may match.

```
<?xml version="1.0" encoding="UTF-8"?>
<ruleset
  xmlns="urn:ietf:params:xml:ns:common-policy"
  xmlns:gp="urn:ietf:params:xml:ns:geolocation-policy"
  xmlns:gml="http://www.opengis.net/gml"
  xmlns:gs="http://www.opengis.net/pidflo/1.0">

  <rule id="AA56i09">
    <conditions>
      <gp:location-condition>
        <gp:location profile="civic-condition"
          xmlns="urn:ietf:params:xml:ns:pidf:geopriv10:civicAddr">
          <country>DE</country>
          <A1>Bavaria</A1>
          <A3>Munich</A3>
          <A4>Perlach</A4>
          <A6>Otto-Hahn-Ring</A6>
          <HNO>6</HNO>
        </gp:location>
        <gp:location profile="geodetic-condition">
          <gs:Circle srsName="urn:ogc:def:crs:EPSG::4326">
            <gml:pos>-34.410649 150.87651</gml:pos>
            <gs:radius uom="urn:ogc:def:uom:EPSG::9001">1500
            </gs:radius>
          </gs:Circle>
        </gp:location>
      </gp:location-condition>
    </conditions>
    <actions/>
    <transformations/>
  </rule>
</ruleset>
```

7.4. Rule Example with Location-Based Transformations

This example shows the transformations specified in this document. The `<provide-civic>` element indicates that the available civic location information is reduced to building level granularity. If geodetic location information is requested, then a granularity reduction is provided as well.

```
<?xml version="1.0" encoding="UTF-8"?>
<ruleset xmlns="urn:ietf:params:xml:ns:common-policy"
  xmlns:gp="urn:ietf:params:xml:ns:geolocation-policy"
  xmlns:lp="urn:ietf:params:xml:ns:basic-location-profiles">

  <rule id="AA56i09">
    <conditions/>
    <actions/>
    <transformations>
      <gp:set-retransmission-allowed>>false
      </gp:set-retransmission-allowed>
      <gp:set-retention-expiry>86400</gp:set-retention-expiry>
      <gp:set-note-well xml:lang="en">My privacy policy goes here.
      </gp:set-note-well>
      <gp:keep-rule-reference>>false
      </gp:keep-rule-reference>

      <gp:provide-location
        profile="civic-transformation">
        <lp:provide-civic>building</lp:provide-civic>
      </gp:provide-location>

      <gp:provide-location
        profile="geodetic-transformation">
        <lp:provide-geo radius="500"/>
      </gp:provide-location>

    </transformations>
  </rule>
</ruleset>
```

The following rule describes the shorthand notation for making the current location of the Target available to Location Recipients without granularity reduction.

```
<?xml version="1.0" encoding="UTF-8"?>
<ruleset xmlns="urn:ietf:params:xml:ns:common-policy"
  xmlns:gp="urn:ietf:params:xml:ns:geolocation-policy">

  <rule id="AA56ia9">
    <conditions/>
    <actions/>
    <transformations>
      <gp:provide-location/>
    </transformations>
  </rule>
</ruleset>
```

7.5. Location Obfuscation Example

Suppose you want to obscure positions in the continental USA.

Step 1:

First, you choose a geodesic projection. If you are measuring location as latitude and longitude, a natural choice is to take a rectangular projection. One latitudinal degree corresponds to approximately 110.6 kilometers, while a good approximation of a longitudinal degree at latitude ϕ is $(\pi/180) * M * \cos(\phi)$, where π is approximately 3.1415, and M is the Earth's average meridional radius, approximately 6,367.5 km. For instance, one longitudinal degree at 30 degrees (say, New Orleans) is 96.39 km, while the formula given offers an estimation of 96.24, which is good enough for our purposes.

We will set up a grid not only for the continental USA, but for the whole earth between latitudes 25 and 50 degrees, and thus will cover also the Mediterranean, South Europe, Japan, and the north of China. As will be seen below, the grid distortion (for not too large grids in this region) is approx $\cos(25)/\cos(50)$, which is 1.4099.

As origin of our grid, we choose the point at latitude 25 degrees and longitude 0 (Greenwich). The latitude 25 degrees is chosen to be just south of Florida and thus south of the continental USA. (On the Southern Hemisphere, the origin should be north of the region to be covered; if the region crosses the Equator, the

origin should be on the Equator. In this way, it is guaranteed that the latitudinal degree has the largest distance at the latitude of the origin).

At 25 degrees, one degree in east-west direction corresponds to approximately $(\pi/180)*M*\cos(25) = 100.72$ km.

The same procedure, basically, produces grids for

- * 45 degrees south to 45 degrees north: Tropics and subtropics, Africa, Australia
- * 25 to 50 degrees (both north or south): Continental United States, Mediterranean, most of China; most of Chile and Argentina, New Zealand
- * 35 to 55 degrees (both north or south): Southern and Central Europe
- * 45 to 60 degrees (both north or south): Central and Northern Europe, Canada
- * 55 to 65 degrees (both north or south): most of Scandinavia
- * 60 to 70 degrees (both north or south): Alaska

Since we do not want to change the grid system often (this would leak more information about obscured locations when they are repeatedly visited), the algorithm should prefer to use the grids discussed above, with origin at the Greenwich meridian and at latitudes $o=0$, $o=25$, $o=35$, $o=45$, $o=55$, and $o=60$ degrees (north) or at latitudes $o=-25$, $o=-35$, $o=-45$, $o=-55$, and $o=-60$ degrees (the minus to indicate "south").

Our choice for the continental USA is $o=25$.

For locations close to the poles, a different projection should be used (not discussed here).

Step 2:

To construct the grid, we start with our chosen origin and place grid points at regular intervals along each of the axes (north-south and east-west) with a distance d between each.

We will now construct a grid for a desired uncertainty of $d = 100\text{km}$. At our origin, 100 km correspond roughly to $d_1 = 100/100.72 = 0.993$ degrees in an east-west direction and to $d_2 = 100/110.6 = 0.904$ degrees in a north-south direction.

The (i,j) -point in the grid (i and j are integers) has longitude d_1*i and latitude $25+d_2*j$, measured in degrees. More generally, if the grid has origin at coordinates $(0,o)$, measured in degrees, the (i,j) -point in the grid has coordinates (longitude = d_1*i , latitude = $o+d_2*j$). The grid has almost no distortion at the latitude of the origin, but it does as we go further away from it.

The distance between two points in the grid at 25 degrees latitude is indeed approximately 100 km, but just above the Canadian border, on the 50th degree, it is $0.993*(\pi/180)*M*\cos(50) = 70.92\text{km}$. Thus, the grid distortion is $100/70.92 = 1.41$, which is acceptable (<1.5). (In the north-south direction, the grid has roughly no distortion; the vertical distance between two neighboring grid points is approximately 100 km).

Step 3:

Now suppose you measure a position at M , with longitude -105 (the minus sign is used to denote 105 degrees *west*; without minus, the point is in China, 105 degrees east) and latitude 40 degrees (just north of Denver, CO). The point M is 105 degrees west and 15 degrees north of our origin (which has longitude 0 and latitude 25).

Let "floor" be the function that returns the largest integer smaller or equal to a floating point number. To calculate SW, the closest point of the grid on the southwest of $M=(m,n)$, we calculate

$$i = \text{floor}(m/d_1) = \text{floor}(-105/0.993) = -106$$

$$j = \text{floor}(n-o/d_2) = \text{floor}(15/0.904) = 16$$

Those are the indexes of SW on the grid. The coordinates of SW are then: $(d_1*i, 25+d_2*j) = (-105.242, 39.467)$.

Thus:

$$l = d_1 * \text{floor}(m/d_1) = -105.243$$

$$r = l + d_1 = -105.243 + 0.993 = -104.250$$

$$b = o + d_2 * \text{floor}(n-o/d_2) = 39.467$$

$$t=b+d2 = 39.467+0.904 = 40.371$$

These are the formulas for l , r , b , and t in the general case of Cartesian projections based on latitude and longitude.

Step 4:

Calculate x and y , the local coordinates of the point M in the small grid square that contains it. This is easy:

$$x=(m-l)/(r-l) = [-105 -(-105.243)]/0.993 = 0.245$$

$$y=(n-b)/(t-b) = [40 - 39.467]/0.904 = 0.590$$

Step 5:

First, compare x with p (0.2887) and $1-p$ (0.7113). x is smaller than p . Therefore, only cases 1, 4, or 6 could hold.

Also, compare y with p (0.2887) and $1-p$ (0.7113). y is between them: $p \leq y < q$. Thus, we must be in case 4. To check, compare y (0.59) with x (0.245) and $1-x$. y is larger than x and smaller than $1-x$. We are in case C4 ($p \leq y < q$ and $x \leq y$ and $y < 1-x$).

Step 6:

Now we choose either SW or NW as the center of the circle.

The obscured location is the circle with radius 100 km and center in SW (coordinates: -105.243, 39.467) or NW (coordinates: -105.243, 40.371).

8. XML Schema for Basic Location Profiles

This section defines the location profiles used as child elements of the transformation element.

```
<?xml version="1.0" encoding="UTF-8"?>
<xs:schema
  targetNamespace="urn:ietf:params:xml:ns:basic-location-profiles"
  xmlns:xs="http://www.w3.org/2001/XMLSchema"
  elementFormDefault="qualified"
  attributeFormDefault="unqualified">

  <!-- profile="civic-transformation" -->

  <xs:element name="provide-civic" default="none">
    <xs:simpleType>
      <xs:restriction base="xs:string">
        <xs:enumeration value="full"/>
        <xs:enumeration value="building"/>
        <xs:enumeration value="city"/>
        <xs:enumeration value="region"/>
        <xs:enumeration value="country"/>
        <xs:enumeration value="none"/>
      </xs:restriction>
    </xs:simpleType>
  </xs:element>

  <!-- profile="geodetic-transformation" -->

  <xs:element name="provide-geo">
    <xs:complexType>
      <xs:attribute name="radius" type="xs:integer"/>
    </xs:complexType>
  </xs:element>

</xs:schema>
```

9. XML Schema for Geolocation Policy

This section presents the XML schema that defines the Geolocation Policy schema described in this document. The Geolocation Policy schema extends the Common Policy schema (see [RFC4745]).

```
<?xml version="1.0" encoding="UTF-8"?>
<xs:schema
  targetNamespace="urn:ietf:params:xml:ns:geolocation-policy"
  xmlns:gp="urn:ietf:params:xml:ns:geolocation-policy"
  xmlns:xs="http://www.w3.org/2001/XMLSchema"
  elementFormDefault="qualified"
  attributeFormDefault="unqualified">

  <!-- Import Common Policy-->
  <xs:import namespace="urn:ietf:params:xml:ns:common-policy"/>

  <!-- This import brings in the XML language attribute xml:lang-->
  <xs:import namespace="http://www.w3.org/XML/1998/namespace"
    schemaLocation="http://www.w3.org/2001/xml.xsd"/>

  <!-- Geopriv Conditions -->

  <xs:element name="location-condition"
    type="gp:locationconditionType"/>

  <xs:complexType name="locationconditionType">
    <xs:complexContent>
      <xs:restriction base="xs:anyType">
        <xs:choice minOccurs="1" maxOccurs="unbounded">
          <xs:element name="location" type="gp:locationType"
            minOccurs="1" maxOccurs="unbounded"/>
          <xs:any namespace="##other" processContents="lax"
            minOccurs="0" maxOccurs="unbounded"/>
        </xs:choice>
      </xs:restriction>
    </xs:complexContent>
  </xs:complexType>

  <xs:complexType name="locationType">
    <xs:complexContent>
      <xs:restriction base="xs:anyType">
        <xs:choice minOccurs="1" maxOccurs="unbounded">
          <xs:any namespace="##other" processContents="lax"
            minOccurs="0" maxOccurs="unbounded"/>
        </xs:choice>
        <xs:attribute name="profile" type="xs:string"/>
        <xs:attribute name="label" type="xs:string"/>
      </xs:restriction>
    </xs:complexContent>
  </xs:complexType>

```



```

        <xs:attribute ref="xml:lang" />
      </xs:restriction>
    </xs:complexContent>
  </xs:complexType>

  <!-- Geopriv transformations -->
  <xs:element name="set-retransmission-allowed"
    type="xs:boolean" default="false"/>
  <xs:element name="set-retention-expiry"
    type="xs:integer" default="0"/>
  <xs:element name="set-note-well"
    type="gp:notewellType"/>
  <xs:element name="keep-rule-reference"
    type="xs:boolean" default="false"/>

  <xs:element name="provide-location"
    type="gp:providelocationType"/>

  <xs:complexType name="notewellType">
    <xs:simpleContent>
      <xs:extension base="xs:string">
        <xs:attribute ref="xml:lang" />
      </xs:extension>
    </xs:simpleContent>
  </xs:complexType>

  <xs:complexType name="providelocationType">
    <xs:complexContent>
      <xs:restriction base="xs:anyType">
        <xs:choice minOccurs="0" maxOccurs="unbounded">
          <xs:any namespace="##other" processContents="lax"
            minOccurs="0" maxOccurs="unbounded"/>
        </xs:choice>
        <xs:attribute name="profile" type="xs:string" />
      </xs:restriction>
    </xs:complexContent>
  </xs:complexType>

</xs:schema>

```

10. XCAP Usage

This section defines the details necessary for clients to manipulate geolocation privacy documents from a server using XCAP. If used as part of a presence system, it uses the same Application Unique ID (AUID) as those rules. See [RFC5025] for a description of the XCAP usage in context with presence authorization rules.

10.1. Application Unique ID

XCAP requires application usages to define a unique Application Unique ID (AUID) in either the IETF tree or a vendor tree. This specification defines the "geolocation-policy" AUID within the IETF tree, via the IANA registration in Section 11.

10.2. XML Schema

XCAP requires application usages to define a schema for their documents. The schema for geolocation authorization documents is described in Section 9.

10.3. Default Namespace

XCAP requires application usages to define the default namespace for their documents. The default namespace is `urn:ietf:params:xml:ns:geolocation-policy`.

10.4. MIME Media Type

XCAP requires application usages to define the MIME media type for documents they carry. Geolocation privacy authorization documents inherit the MIME type of Common Policy documents, `application/auth-policy+xml`.

10.5. Validation Constraints

This specification does not define additional constraints.

10.6. Data Semantics

This document discusses the semantics of a geolocation privacy authorization.

10.7. Naming Conventions

When a Location Server receives a request to access location information of some user foo, it will look for all documents within `http://[xcaproot]/geolocation-policy/users/foo` and use all documents found beneath that point to guide authorization policy.

10.8. Resource Interdependencies

This application usage does not define additional resource interdependencies.

10.9. Authorization Policies

This application usage does not modify the default XCAP authorization policy, which is that only a user can read, write, or modify his/her own documents. A server can allow privileged users to modify documents that they do not own, but the establishment and indication of such policies is outside the scope of this document.

11. IANA Considerations

There are several IANA considerations associated with this specification.

11.1. Geolocation Policy XML Schema Registration

This section registers an XML schema in the IETF XML Registry as per the guidelines in [RFC3688].

URI: urn:ietf:params:xml:schema:geolocation-policy

Registrant Contact: IETF Geopriv Working Group (geopriv@ietf.org),
Hannes Tschofenig (hannes.tschofenig@nsn.com).

XML: The XML schema to be registered is contained in Section 9. Its first line is

```
<?xml version="1.0" encoding="UTF-8"?>
```

and its last line is

```
</xs:schema>
```

11.2. Geolocation Policy Namespace Registration

This section registers a new XML namespace in the IETF XML Registry as per the guidelines in [RFC3688].

URI: urn:ietf:params:xml:ns:geolocation-policy

Registrant Contact: IETF Geopriv Working Group (geopriv@ietf.org),
Hannes Tschofenig (hannes.tschofenig@nsn.com).

XML:

```
BEGIN
<?xml version="1.0"?>
<!DOCTYPE html PUBLIC "-//W3C//DTD XHTML Basic 1.0//EN"
  "http://www.w3.org/TR/xhtml-basic/xhtml-basic10.dtd">
<html xmlns="http://www.w3.org/1999/xhtml">
<head>
  <meta http-equiv="content-type"
    content="text/html; charset=iso-8859-1"/>
  <title>Geolocation Policy Namespace</title>
</head>
<body>
  <h1>Namespace for Geolocation Authorization Policies</h1>
  <h2>urn:ietf:params:xml:schema:geolocation-policy</h2>
<p>See <a href="http://www.rfc-editor.org/rfc/rfc6772.txt">
  RFC 6772</a>.</p>
</body>
</html>
END
```

11.3. Geolocation Policy Location Profile Registry

This document creates a registry of location profile names for the Geolocation Policy framework. Profile names are XML tokens. This registry will operate in accordance with RFC 5226 [RFC5226], Specification Required.

This document defines the following profile names:

geodetic-condition: Defined in Section 4.1.
civic-condition: Defined in Section 4.2.
geodetic-transformation: Defined in Section 6.5.2.
civic-transformation: Defined in Section 6.5.1.

11.4. Basic Location Profile XML Schema Registration

This section registers an XML schema in the IETF XML Registry as per the guidelines in [RFC3688].

URI: urn:ietf:params:xml:schema:basic-location-profiles

Registrant Contact: IETF Geopriv Working Group (geopriv@ietf.org),
Hannes Tschofenig (hannes.tschofenig@nsn.com).

XML: The XML schema to be registered is contained in Section 8. Its first line is

```
<?xml version="1.0" encoding="UTF-8"?>
```

and its last line is

```
</xs:schema>
```

11.5. Basic Location Profile Namespace Registration

This section registers a new XML namespace in the IETF XML Registry as per the guidelines in [RFC3688].

URI: urn:ietf:params:xml:ns:basic-location-profiles

Registrant Contact: IETF Geopriv Working Group (geopriv@ietf.org),
Hannes Tschofenig (hannes.tschofenig@nsn.com).

XML:

BEGIN

```
<?xml version="1.0"?>
<!DOCTYPE html PUBLIC "-//W3C//DTD XHTML Basic 1.0//EN"
  "http://www.w3.org/TR/xhtml-basic/xhtml-basic10.dtd">
<html xmlns="http://www.w3.org/1999/xhtml">
<head>
  <meta http-equiv="content-type"
    content="text/html; charset=iso-8859-1"/>
  <title>Basic Location Profile Namespace</title>
</head>
<body>
  <h1>Namespace for Basic Location Profile</h1>
  <h2>urn:ietf:params:xml:ns:basic-location-profiles</h2>
<p>See <a href="http://www.rfc-editor.org/rfc/rfc6772.txt">
  RFC 6772</a>.</p>
</body>
</html>
END
```

11.6. XCAP Application Usage ID

This section registers an XCAP Application Unique ID (AUID) in the "XML-XCAP Application Unique IDs" registry according to the IANA procedures defined in [RFC4825].

Name of the AUID: geolocation-policy

Description: Geolocation privacy rules are documents that describe the permissions that a Target has granted to Location Recipients that access information about his/her geographic location.

12. Internationalization Considerations

The policies described in this document are mostly meant for machine-to-machine communications; as such, many of its elements are tokens not meant for direct human consumption. If these tokens are presented to the end user, some localization may need to occur. The policies are, however, supposed to be created with the help of humans, and some of the elements and attributes are subject to internationalization considerations. The content of the <label> element is meant to be provided by a human (the Rule Maker) and also displayed to a human. Furthermore, the location condition element (<location-condition>, using the civic location profile, see Section 4.2) and the <set-note-well> element (see Section 6.3) may contain non-US-ASCII letters.

The geolocation policies utilize XML, and all XML processors are required to understand UTF-8 and UTF-16 encodings. Therefore, all entities processing these policies MUST understand UTF-8- and UTF-16-encoded XML. Additionally, geolocation policy-aware entities MUST NOT encode XML with encodings other than UTF-8 or UTF-16.

13. Security Considerations

13.1. Introduction

This document aims to allow users to prevent unauthorized access to location information and to restrict access to information dependent on the location of the Target, using location-based conditions. This is accomplished using authorization policies. This work builds on a series of other documents: security requirements are described in [RFC6280] and a discussion of generic security threats is available with [RFC3694]. Aspects of combining permissions in cases of multiple occurrence are addressed in [RFC4745].

In addition to the authorization policies, mechanisms for obfuscating location information are described. A theoretical treatment of location obfuscation is provided in [DUCKHAM05] and in [IFIP07]. [DUCKHAM05] provides the foundation, and [IFIP07] illustrates three different types of location obfuscation by enlarging the radius, by shifting the center, and by reducing the radius. The algorithm in Section 6.5.2 for geodetic location information obfuscation uses these techniques.

The requirements for protecting privacy-sensitive location information vary. The two obfuscation algorithms in this document provide a basis for protecting against unauthorized disclosure of location information, but they have limitations. Application and user requirements vary widely; therefore, an extension mechanism is support for defining and using different algorithms.

13.2. Obfuscation

Whenever location information is returned to a Location Recipient, it contains the location of the Target. This is also true when location is obfuscated, i.e., the Location Server does not lie about the Target's location but instead hides it within a larger location shape. Even without the Target's movement, there is a danger that information will be revealed over time. While the Target's location is not revealed within a particular region of the grid, the size of that returned region matters as well as the precise location of the Target within that region. Returning location shapes that are randomly computed will over time reveal more and more information about the Target.

Consider Figure 1, which shows three ellipses, a dotted area in the middle, and the Target's true location marked as 'x'. The ellipses illustrate the location shapes as received by a potential Location Recipient over time for requests of a Target's location information. Collecting information about the returned location information over time allows the Location Recipient to narrow the potential location of the Target down to the dotted area in the center of the graph.

For this purpose, the algorithm described in Section 6.5.2 uses a grid that ensures the same location information is reported while the Target remains in the same geographical area.




Figure 1: Obfuscation: A Static Target

An obscuring method that returns different results for consecutive requests can be exploited by recipients wishing to use this property. Rate limiting the generation of new obscured locations or providing the same obscured location to recipients for the same location might limit the information that can be obtained. Note, however, that providing a new obscured location based on a change in location provides some information to recipients when they observe a change in location.

When the Target is moving, then the location transformations reveal information when switching from one privacy region to another one. For example, when a transformation indicates that civic location is provided at a 'building' level of granularity, floor levels, room numbers, and other details normally internal to a building would be hidden. However, when the Target moves from one building to the next one, then the movement would still be recognizable as the disclosed location information would be reflected by the new civic location information indicating the new building. With additional knowledge about building entrances and floor plans, it would be possible to learn additional information.

13.3. Algorithm Limitations

The algorithm presented in Section 6.5.2 has some issues where information is leaked: when moving, when switching from one privacy region to another one, and also when the user regularly visits the same location.

The first issue arises if the algorithm provides different location information (privacy region) only when the previous one becomes inapplicable. The algorithm discloses new information the moment that the Target is on the border of the old privacy region.

Another issue arises if the algorithm produces the different values for the same location that is repeatedly visited. Suppose a user goes home every night. If the reported obfuscated locations are all randomly chosen, an analysis can reveal the home location with high precision.

In addition to these concerns, the combination of an obscured location with public geographic information (highways, lakes, mountains, cities, etc.) may yield much more precise location information than is desired. But even without it, just observing movements, once or multiple times, any obscuring algorithm can leak information about velocities or positions. Suppose a user wants to disclose location information with a radius of r . The privacy region, a circle with that radius, has an area of $A = \pi * r^2$. An adversary, observing the movements, will deduce that the target is

visiting, was visiting, or regularly visits, a region of size A_1 , smaller than A . The ratio A_1/A should be, even in the worst case, larger than a fixed known number, in order that the user can predict the worst-case information leakage. The choices of Section 6.5.2 are such that this maximum leakage can be established: by any statistical procedures, without using external information (highways, etc., as discussed above), the quotient A_1/A is larger than 0.13 (= $1/(5*1.5)$). Thus, for instance, when choosing a provided location of size 1000 km^2 , he will be leaking, in worst case, the location within a region of size 130 km^2 .

13.4. Usability

There is the risk that end users are specifying their location-based policies in such a way that very small changes in location yields a significantly different level of information disclosure. For example, a user might want to set authorization policies differently when they are in a specific geographical area (e.g., at home, in the office). Location might be the only factor in the policy that triggers a very different action and transformation to be executed. The accuracy of location information is not always sufficient to unequivocally determine whether a location is within a specific boundary [GEOPRIV-UNCERTAINTY]. In some situations, uncertainty in location information could produce unexpected results for end users. Providing adequate user feedback about potential errors arising from these limitation can help prevent unintentional information leakage.

Users might create policies that are nonsensical. To avoid such cases, the software used to create the authorization policies should perform consistency checks, and when authorization policies are uploaded to the policy servers, then further checks can be performed. When XCAP is used to upload authorization policies, then built-in features of XCAP can be utilized to convey error messages back to the user about an error condition. Section 8.2.5 of [RFC4825] indicates that some degree of application-specific checking is provided when authorization policies are added, modified, or deleted. The XCAP protocol may return a 409 response with a response that may contain a detailed conflict report containing the <constraint-failure> element. A human-readable description of the problem can be indicated in the 'phrase' attribute of that element.

13.5. Limitations of Obscuring Locations

Location-obscuring attempts to remove information about the location of a Target. The effectiveness of location obscuring is determined by how much uncertainty a Location Recipient has about the location of the Target. A location-obscuring algorithm is effective if the

Location Recipient cannot recover a location with better uncertainty than the obscuring algorithm was instructed to add.

Effective location obscuring is difficult. The amount of information that can be recovered by a determined and resourceful Location Recipient can be considerably more than is immediately apparent. A concise summary of the challenges is included in [DUCKHAM10].

A Location Recipient in possession of external information about the Target or geographical area that is reported can make assumptions or guesses aided by that information to recover more accurate location information. This is true even when a single location is reported, but it is especially true when multiple locations are reported for the same Target over time.

Furthermore, a Location Recipient that attempts to recover past locations for a Target can use later-reported locations to further refine any recovered location. A location-obscuring algorithm typically does not have any information about the future location of the Target.

The degree to which location information can be effectively degraded by an obscuring algorithm depends on the information that is used by the obscuring algorithm. If the information available to the obscuring algorithm is both more extensive and more effectively employed than the information available to the Location Recipient, then location obscuring might be effective.

Obscured locations can still serve a purpose where a Location Recipient is willing to respect privacy. A privacy-respecting Location Recipient can choose to interpret the existence of uncertainty as a request from a Rule Maker to not recover location.

Location obscuring is unlikely to be effective against a more determined or resourceful adversary. Withholding location information entirely is perhaps the most effective method of ensuring that it is not recovered.

As a final caution, we note that omitted data also conveys some information. Selective withholding of information reveals that there is something worth hiding. That information might be used to reveal something of the information that is being withheld. For example, if location is only obscured around a user's home and office, then the lack of location for that user and the current time will likely mean that the user is at home at night and in the office during the day, defeating the purpose of the controls.

14. References

14.1. Normative References

- [GML] OpenGIS, "OpenGIS Geography Markup Language (GML) Implementation Specification, Version 3.1.1, OGC 03-105r1", July 2004, <http://portal.opengeospatial.org/files/?artifact_id=4700>.
- [NIMA.TR8350.2-3e] "Department of Defense (DoD) World Geodetic System 1984 (WGS 84), Third Edition", NIMA TR8350.2, January 2000.
- [OGC-06-103r4] OpenGIS, "OpenGIS Implementation Specification for Geographic information - Simple feature access - Part 1: Common architecture", May 2011, <<http://www.opengeospatial.org/standards/sfa?>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3688] Mealling, M., "The IETF XML Registry", BCP 81, RFC 3688, January 2004.
- [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.
- [RFC5139] Thomson, M. and J. Winterbottom, "Revised Civic Location Format for Presence Information Data Format Location Object (PIDF-LO)", RFC 5139, February 2008.
- [RFC5491] Winterbottom, J., Thomson, M., and H. Tschofenig, "GEOPRIV Presence Information Data Format Location Object (PIDF-LO) Usage Clarification, Considerations, and Recommendations", RFC 5491, March 2009.

14.2. Informative References

- [DUCKHAM05] Duckham, M. and L. Kulik, "A Formal Model of Obfuscation and Negotiation for Location Privacy", In Proc. of the 3rd International Conference PERVASIVE 2005, Munich, Germany, May 2005.

- [DUCKHAM10] Duckham, M., "Moving Forward: Location Privacy and Location Awareness", In Proc. 3rd ACM SIGSPATIAL Workshop on Security and Privacy in GIS and LBS (SPRINGL), ACM, November 2010.
- [GEO-SHAPE] Thomson, M., "Geodetic Shapes for the Representation of Uncertainty in PIDF-LO", Work in Progress, December 2006.
- [GEOPRIV-UNCERTAINTY] Thomson, M. and J. Winterbottom, "Representation of Uncertainty and Confidence in PIDF-LO", Work in Progress, March 2012.
- [IFIP07] Ardagna, C., Cremonini, M., Damiani, E., De Capitani di Vimercati, S., and P. Samarati, "Location Privacy Protection through Obfuscation-Based Techniques", Proceedings of the 21st Annual IFIP WG 11.3 Working Conference on Data and Applications Security, Redondo Beach, CA, USA, July 2007.
- [RFC2392] Levinson, E., "Content-ID and Message-ID Uniform Resource Locators", RFC 2392, August 1998.
- [RFC2778] Day, M., Rosenberg, J., and H. Sugano, "A Model for Presence and Instant Messaging", RFC 2778, February 2000.
- [RFC3694] Danley, M., Mulligan, D., Morris, J., and J. Peterson, "Threat Analysis of the Geopriv Protocol", RFC 3694, February 2004.
- [RFC4079] Peterson, J., "A Presence Architecture for the Distribution of GEOPRIV Location Objects", RFC 4079, July 2005.
- [RFC4119] Peterson, J., "A Presence-based GEOPRIV Location Object Format", RFC 4119, December 2005.
- [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.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.

[RFC6280] Barnes, R., Lepinski, M., Cooper, A., Morris, J.,
Tschofenig, H., and H. Schulzrinne, "An Architecture for
Location and Location Privacy in Internet Applications",
BCP 160, RFC 6280, July 2011.

Appendix A. Acknowledgments

This document is informed by the discussions within the IETF GEOPRIV working group, including discussions at the GEOPRIV interim meeting in Washington, D.C., in 2003.

We particularly want to thank Allison Mankin <mankin@psg.com>, Randall Gellens <rg+iETF@qualcomm.com>, Andrew Newton <anewton@ecotroph.net>, Ted Hardie <hardie@qualcomm.com>, and Jon Peterson <jon.peterson@neustar.biz> for their help in improving the quality of this document.

We would like to thank Christian Guenther for his help with an earlier version of this document. Furthermore, we would like to thank Johnny Vrancken for his document reviews in September 2006, December 2006 and January 2007. James Winterbottom provided a detailed review in November 2006. Richard Barnes gave a detailed review in February 2008.

This document uses text from "Geodetic Shapes for the Representation of Uncertainty in PIDF-LO" [GEO-SHAPE], authored by Martin Thomson.

We would like to thank Matt Lepinski and Richard Barnes for their comments regarding the geodetic location transformation procedure. Richard provided us with a detailed text proposal.

Robert Sparks, and Warren Kumari deserve thanks for their input on the location obfuscation discussion. Robert implemented various versions of the algorithm in the graphical language "Processing" and thereby helped us tremendously to understand problems with the previously illustrated algorithm.

We would like to thank Dan Romascanu, Yoshiko Chong, and Jari Urpalainen for their last call comments.

Finally, we would like to thank the following individuals for their feedback as part of the IESG, GenArt, and SecDir review: Jari Arkko, Lisa Dusseault, Eric Gray, Sam Hartman, Russ Housley, Cullen Jennings, Chris Newman, Jon Peterson, Tim Polk, Carl Reed, and Brian Rosen.

Although John Morris is currently employed by the U.S. Government, he participated in the development of this document in his personal capacity, and the views expressed in the document may not reflect those of his employer.

Appendix B. Pseudocode

This section provides an informal description for the algorithm described in 6.5.2 and 7.5 as pseudocode. In addition to the algorithm, it randomly chooses among equidistant landmarks based on the previous location.

Constants

```
P = sqrt(3)/6 // approx 0.2887
q = 1 - p     // approx 0.7113
```

Parameters

```
prob: real // prob is a parameter in the range
      // 0.5 <= prob <=1
      // recommended is a value for prob between 0.7 and 0.9
      // the default of prob is 0.8
```

Inputs

```
M = (m,n) : real * real
      // M is a pair of reals: m and n
      // m is the longitude and n the latitude,
      // respectively, of the measured location
      // The values are given as real numbers, in the
      // range: -180 < m <= 180; -90 < n < 90
      // minus values for longitude m correspond to "West"
      // minus values for latitude n correspond to "South"

radius : integer // the 'radius' or uncertainty,
              // measured in meters

prev-M = (prev-m1, prev-n1): real * real
      // the *previously* provided location, if available
      // prev-m1 is the longitude and
      // prev-n1 the latitude, respectively

o : real

// this is the reference latitude for the geodesic projection
// The value of 'o' is chosen according to the table below.
// The area you want to project MUST be included in
// between a minimal latitude and a maximal latitude
// given by the two first columns of the table.
// (Otherwise the transformation is not available).
```

min lat	max lat	Examples	o
-45	45	Tropics and subtropics Africa Australia	0
25	50	Continental US Mediterranean most of China	25
35	55	Southern and Central Europe	35
45	60	Central and Northern Europe	45
55	65	most of Scandinavia	55
60	70		60
-50	-25	most of Chile and Argentina New Zealand	-50
-35	-55		-35
-45	-60		-45
-55	-65		-55
-60	-70		-60

Outputs

```
M1 = (m1,n1) : real * real // longitude and latitude,
              // respectively, of the provided location
```

Local Variables

```
d, d1, d2, l, r, b, t, x, y: real
SW, SE, NW, NE: real * real
  // pairs of real numbers, interpreted as coordinates
  // longitude and latitude, respectively
```

```
temp : Integer[1..8]
```

Function

```
choose(Ma, Mb: real * real): real * real;
  // This function chooses either Ma or Mb
  // depending on the parameter 'prob'
  // and on prev-M1, the previous value of M1:
  // If prev-M1 == Ma choose Ma with probability 'prob'
  // If prev-M1 == Mb choose Mb with probability 'prob'
  // Else choose Ma or Mb with probability 1/2
```

Begin

```
rand:= Random[0,1];
  // a real random number between 0 and 1
If    prev-M1 == Ma Then
      If rand < prob Then choose := Ma;
                          Else choose := Mb;  EndIf
Elseif prev-M1 == Mb Then
      If rand < prob Then choose := Mb;
                          Else choose := Ma;  EndIf
Else
      If rand < 0.5 Then choose := Ma;
                          Else choose := Mb;  EndIf
End // Function choose
```

Main // main procedure

```
Begin
d := radius/1000; // uncertainty, measured in km

d1:= (d * 180) / (pi*M*cos(o));

d2:= d / 110.6;

l := d1*floor(m/d1)
  // "floor" returns the largest integer
  // smaller or equal to a floating point number
r := l+d1;
```

```
b := o+d2*floor(n-o/d2);
t := b+d2;

x := (m-1)/(r-1);
y := (n-b)/(t-b);

SW := (1,b);
SE := (r,b);
NW := (1,t);
NE := (r,t);

If      x < p and y < p      Then M1 := SW;
Elseif  x < p and q <= y    Then M1 := NW;
Elseif  q <= x and y < p    Then M1 := SE;
Elseif  q <= x and q <= y   Then M1 := NE;
Elseif  p <= x and x < q and y < x and y < 1-x
        Then M1 := choose(SW,SE);
Elseif  p <= y and y < q and x <= y and y < 1-x
        Then M1 := choose(SW,NW);
Elseif  p <= y and y < q and y < x and 1-x <= y
        Then M1 := choose(SE,NE);
Elseif  p <= x and x < q and x <= y and 1-x <= y
        Then M1 := choose(NW,NE);
Endif

End // Main
```

Authors' Addresses

Henning Schulzrinne (editor)
Columbia University
Department of Computer Science
450 Computer Science Building
New York, NY 10027
USA

Phone: +1 212-939-7042
EMail: schulzrinne@cs.columbia.edu
URI: <http://www.cs.columbia.edu/~hgs>

Hannes Tschofenig (editor)
Nokia Siemens Networks
Linnoitustie 6
Espoo 02600
Finland

Phone: +358 (50) 4871445
EMail: Hannes.Tschofenig@gmx.net
URI: <http://www.tschofenig.priv.at>

Jorge R. Cuellar
Siemens
Otto-Hahn-Ring 6
Munich, Bavaria 81739
Germany

EMail: Jorge.Cuellar@siemens.com

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

Phone: +1 817-271-3552
EMail: jmpolk@cisco.com

John B. Morris, Jr.

EMail: ietf@jmorris.org

Martin Thomson
Microsoft
3210 Porter Drive
Palo Alto, CA 94304
USA

Phone: +1 650-353-1925
EMail: martin.thomson@gmail.com