

GEOPRIV
Internet-Draft
Intended status: Informational
Expires: May 13, 2011

R. Barnes
BBN Technologies
M. Thomson
J. Winterbottom
Andrew Corporation
H. Tschofenig
Nokia Siemens Networks
November 9, 2010

Location Configuration Extensions for Policy Management
draft-barnes-geopriv-policy-uri-02

Abstract

Current location configuration protocols are capable of provisioning an Internet host with a location URI that refers to the host's location. These protocols lack a mechanism for the target host to inspect or set the privacy rules that are applied to the URIs they distribute. This document extends the current location configuration protocols to provide hosts with a reference to the rules that are applied to a URI, so that the host can view or set these rules.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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

This Internet-Draft will expire on May 13, 2011.

Copyright Notice

Copyright (c) 2010 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	3
2. Definitions	4
3. Policy URIs	4
3.1. Policy URI Usage	4
3.2. Policy URI Allocation	5
4. Location Configuration Extensions	6
4.1. HELD	6
4.2. DHCP	7
5. Examples	8
5.1. HELD	8
5.2. DHCP	8
5.3. Basic access control policy	9
6. Acknowledgements	11
7. IANA Considerations	12
7.1. URN Sub-Namespace Registration for urn:ietf:params:xml:ns:geopriv:held:policy	12
7.2. XML Schema Registration	12
7.3. DHCP LuriType Registration	13
8. Operational Considerations	13
9. Security Considerations	14
9.1. Integrity and Confidentiality for Authorization Policy Data	14
9.2. Access Control for Authorization Policy	14
9.3. Location URI Allocation	15
10. References	16
10.1. Normative References	16
10.2. Informative References	17
Authors' Addresses	18

1. Introduction

A critical step in enabling Internet hosts to access location-based services is to provision those hosts with information about their own location. This is accomplished via a Location Configuration Protocol (LCP) [RFC5687], which allows a location provider (e.g., a local access network) to inform a host about its location.

There are two basic patterns for location configuration, namely configuration "by value" and "by reference" [RFC5808]. Configuration by value provisions a host directly with its location, by providing it location information that is directly usable (e.g., coordinates or a civic address). Configuration by reference provides a host with a URI that references the host's location, i.e., one that can be dereferenced to obtain the location (by value) of the host.

In some cases, location by reference offers a few benefits over location by value. From a privacy perspective, the required dereference transaction provides a policy enforcement point, so that the opaque URI itself can be safely conveyed over untrusted media (e.g., SIP through untrusted proxies [RFC5606]). If the target host is mobile, an application provider can use a single reference to obtain the location of the host multiple times, saving bandwidth to the host. For some configuration protocols, the location object referenced by a location URI provides a much more expressive syntax for location values than the configuration protocol itself (e.g., DHCP geodetic location [I-D.ietf-geopriv-rfc3825bis] versus GML in a PIDF-LO [RFC4119]).

From a privacy perspective, however, current LCPs are limited in their flexibility, in that they do not provide the Device (the client in an LCP) with a way to inform the Location Server with policy for how his location information should be handled. This document addresses this gap by defining a simple mechanism for referring to and manipulating policy, and by extending current LCPs to carry policy references. Using the mechanisms defined in this document, an LCP server (acting for the Location Server) can inform a client as to which policy document controls a given location resource, and the LCP client (in its Rule Maker role) can inspect this document and modify it as necessary.

The remainder of this document is structured as follows: After introducing a few relevant terms, we define policy URIs as a channel for referencing, inspecting, and updating policy documents. We then define extensions to the HELD protocol and the DHCP option for location by reference to allow these protocols to carry policy URIs. Examples are given that demonstrate how policy URIs are carried in these protocols and how they can be used by clients.

2. Definitions

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].

3. Policy URIs

A policy URI is an HTTP [RFC2616] URI that identifies a policy resource that contains the authorization policy for a linked location resource. Access to the location resource is governed by the contents of the authorization policy.

A policy URI identifies an HTTP resource that a Rule Maker can use to inspect and install policy documents that tell a Location Server how it should protect the associated location resource. A policy URI always identifies a resource that can be represented as a common-policy document [RFC4745] (possibly including some extensions; e.g., for geolocation policy [I-D.ietf-geopriv-policy]).

Note: RFC 3693 [RFC3693] identified the Rule Holder role as the one that stores policy information. In this document, the Location Server is also a Rule Holder.

3.1. Policy URI Usage

A Location Server that is the authority for policy URIs MUST support GET, PUT, and DELETE requests to these URIs, in order to allow clients to inspect, replace, and delete policy documents. Clients support the three request methods as they desire to perform these operations.

Knowledge of the policy URI can be considered adequate evidence of authorization. A Location Server SHOULD allow all requests, but it MAY deny certain requests based on local policy. For instance, a Location Server might allow clients to inspect policy (GET), but not to update it (PUT).

A GET request to a policy URI is a request for the referenced policy information. If the request is authorized, then the Location Server sends an HTTP 200 response containing the complete policy identified by the URI.

A PUT request to a policy URI is a request to replace the current policy. The entity-body of a PUT request includes a complete policy document. When a Location Server receives a PUT request, it MUST validate the policy document included in the body of the request. If

the request is valid and authorized, then the Location Server replaces the current policy with the policy provided in the request.

A DELETE request to a policy URI is a request to delete the referenced policy document and terminate access to the protected resource. If the request is authorized, then the Location Server deletes the policy referenced by the URI and disallows any further access to the location resource it governs.

The Location Server MUST support policy documents in the common-policy format [RFC4745], as identified by the MIME media type of "application/auth-policy+xml". The common-policy format MUST be provided as the default format in response to GET requests that do not include specific "Accept" headers, but content negotiation MAY be used to allow for other formats.

This usage of HTTP is generally compatible with the use of XCAP [RFC4825] or WebDAV [RFC4918] to manage policy documents, but this document does not define or require the use of these protocols.

3.2. Policy URI Allocation

A Location Server creates a policy URI for a specific location resource at the time that the location resource is created; that is, a policy URI is created at the same time as the location URI that it controls. The URI of the policy resource MUST be different to the location URI.

A policy URI is provided to a target device as part of the location configuration process. A policy URI MUST NOT be provided to an entity that is not authorized to view or set policy. A location server that provides a location configuration in addition to other location services (e.g., answering dereferencing requests [I-D.ietf-geopriv-deref-protocol] or requests from third parties [I-D.ietf-geopriv-held-identity-extensions]) MUST only include policy URIs in response to location configuration requests.

Each location URI has either one policy URI or no policy URI. A location server MUST NOT allocate multiple policy URIs controlling the same location URI. The initial policy that is referenced by a policy URI MUST be identical to the policy that would be applied in the absence of a policy URI. A client that does not support policy URIs can continue to use the location URI as they would have if no policy URI were provided.

Without a policy URI, clients have no way to know what this default policy is. The safest assumption for clients is that the default policy grants any request to dereference a location URI,

regardless of the requester's identity. With a policy URI, a client can ask the server to describe the default policy (with a GET request), or update the policy with a PUT request, prior to distributing the location URI.

A Location Server chooses whether or not to provide a policy URI based on local policy. A HELD-specific extension also allows a requester to specifically ask for a policy URI.

A policy URI is a shared secret between Location Server and its clients. Knowledge of a policy URI is all that is required to perform any operations allowed on the policy. Thus, a policy URI is constructed so that it is hard to predict (see Section 9).

4. Location Configuration Extensions

Location configuration protocols can provision hosts with location URIs that refer to the host's location. If the target host is to control policy on these URIs, it needs a way to access the policy that the Location Server uses to guide how it serves location URIs. This section defines extensions to LCPs to carry policy URIs that the target can use to control access to location resources.

4.1. HELD

The HELD protocol [I-D.ietf-geopriv-http-location-delivery] defines a "locationUriSet" element, which contain a set of one or more location URIs that reference the same resource and share a common access control policy. The schema in Figure 1 defines two extension elements for HELD: an empty "requestPolicyUri" element that is added to a location request to indicate that a Device desires that a policy URI be allocated; and a "policyUri" element that is included as a sub-element of the HELD "locationResponse" element.

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

  <xs:element name="requestPolicyUri">
    <xs:complexType name="empty"/>
  </xs:element>

  <xs:element name="policyUri" type="xs:anyURI"/>

</xs:schema>
```

Figure 1

The URI carried in a "policyUri" element refers to the common access control policy for requests for the target's location, including dereference requests for location URIs in the location response as well as third-party requests. The URI MUST be a policy URI as described in Section 3. A policy URI MUST use the "http:" or "https:" scheme, and the Location Server MUST support the specified operations on the URI.

A HELD request MAY contain an explicit request for a policy URI. The presence of the "requestPolicyUri" element in a location request indicates that a policy URI is desired. A location server may provide a policy URI regardless of the presence of this element.

4.2. DHCP

The DHCP location by reference option [I-D.ietf-geopriv-dhcp-lbyr-uri-option] provides location URIs in sub-options called LuriElements. This document defines a new LuriElement type for policy URIs.

LuriType=TBD Policy-URI - This is a policy URI that refers to the access control policy for the location URIs.

[NOTE TO IANA/RFC-EDITOR: Please replace TBD above with the assigned LuriType value and remove this note]

A Policy-URI LuriElement uses a UTF-8 character encoding.

A Policy-URI LuriElement identifies the policy resource for all location URIs included in the location URI option. The URI MUST be a policy URI as described in Section 3: It MUST use either the "http:"

or "https:" scheme, and the Location Server MUST support the specified operations on the URI.

5. Examples

In this section, we provide some brief illustrations of how policy URIs are delivered to target hosts and used by those hosts to manage policy.

5.1. HELD

A HELD request that explicitly requests the creation of a policy URI has the following form:

```
<locationRequest xmlns="urn:ietf:params:xml:ns:geopriv:held">
  <locationType exact="true">locationURI</locationType>
  <requestPolicyUri
    xmlns="urn:ietf:params:xml:ns:geopriv:held:policy"/>
</locationRequest>
```

A HELD response providing a single "locationUriSet", containing two URIs under a common policy, would have the following form:

```
<locationResponse xmlns="urn:ietf:params:xml:ns:geopriv:held">
  <locationUriSet expires="2011-01-01T13:00:00.0Z">
    <locationURI>
      https://ls.example.com:9768/357yc6s64ceyoiuy5ax3o
    </locationURI>
    <locationURI>
      sip:9769+357yc6s64ceyoiuy5ax3o@ls.example.com:
    </locationURI>
  </locationUriSet>
  <policyUri xmlns="urn:ietf:params:xml:ns:geopriv:held:policy">
    https://ls.example.com:9768/policy/357lp6f64prlbvhl5nk3b
  </policyUri>
</locationResponse>
```

5.2. DHCP

A DHCP option providing one of the location URIs and the corresponding policy URI from the previous example would have the following form:

0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9
option-code																				110																			
1					0					1					49					'h'																			
't'					't'					'p'					's'																								
':'					'/'					'/'					'l'																								
's'					'.'					...																													
TBD					56					'h'					't'																								
't'					'p'					's'					':'																								
'/'					'/'					...																													

[NOTE TO IANA/RFC-EDITOR: Please replace TBD above with the assigned LuriType value and remove this note]

5.3. Basic access control policy

Consider a user that gets the policy URI <https://ls.example.com:9768/policy/357lp6f64prlbvhl5nk3b>, as in the above LCP example. The first thing this allows the user to do is inspect the default policy that the LS has assigned to this URI:

```
GET /policy/357lp6f64prlbvhl5nk3b HTTP/1.1
Host: ls.example.com:9768
```

```
HTTP/1.1 200 OK
Content-type: application/auth-policy+xml
Content-length: 388
```

```
<?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>
      <validity>
        <until>2011-01-01T13:00:00.0Z</until>
      </validity>
    </conditions>
    <actions/>
    <transformations>
      <gp:provide-location/>
      <gp:set-retransmission-allowed>
        false
      </gp:set-retransmission-allowed>
      <gp:set-retention-expiry>0</gp:set-retention-expiry>
    </transformations>
  </rule>
</ruleset>
```

This policy allows any requester to obtain location information, as long as they know the location URI. If the user disagrees with this policy, and prefers for example, to only provide location to one friend, at a city level of granularity, then he can install this policy on the Location Server:

```
PUT /policy/357lp6f64prlbvhl5nk3b HTTP/1.1
Host: ls.example.com:9768
Content-type: application/auth-policy+xml
Content-length: 462
```

```
<?xml version="1.0" encoding="UTF-8"?>
<ruleset xmlns="urn:ietf:params:xml:ns:common-policy">
  <rule id="f3g44r1">
    <conditions>
      <identity>
        <one id="sip:friend@example.com"/>
      </identity>
      <validity>
        <until>2011-01-01T13:00:00.0Z</until>
      </validity>
    </conditions>
    <actions/>
    <transformations>
      <gp:provide-location
        profile="civic-transformation">
        <lp:provide-civic>city</lp:provide-civic>
      </gp:provide-location>
    </transformations>
  </rule>
</ruleset>
```

```
HTTP/1.1 200 OK
```

Finally, after using the URI for a period, the user wishes to permanently invalidate the URI.

```
DELETE /policy/357lp6f64prlbvhl5nk3b HTTP/1.1
Host: ls.example.com:9768
```

```
HTTP/1.1 200 OK
```

6. Acknowledgements

Thanks to Mary Barnes, Alissa Cooper, and Hannes Tschofenig for providing critical commentary and input on the ideas described in this document.

7. IANA Considerations

This document requires several IANA registrations, detailed below.

7.1. URN Sub-Namespace Registration for urn:ietf:params:xml:ns:geopriv:held:policy

This section registers a new XML namespace, "urn:ietf:params:xml:ns:geopriv:held:policy", per the guidelines in [RFC3688].

URI: urn:ietf:params:xml:ns:grip

Registrant Contact: IETF, GEOPRIV working group,
(geopriv@ietf.org), Richard Barnes (rbarnes@bbn.com).

XML:

```
BEGIN
<?xml version="1.0"?>
<!DOCTYPE html PUBLIC "-//W3C//DTD XHTML 1.0 Strict//EN"
  "http://www.w3.org/TR/xhtml1/DTD/xhtml1-strict.dtd">
<html xmlns="http://www.w3.org/1999/xhtml" xml:lang="en">
  <head>
    <title>HELD Policy URI Extension</title>
  </head>
  <body>
    <h1>Namespace for HELD Policy URI Extension</h1>
    <h2>urn:ietf:params:xml:ns:geopriv:held:policy</h2>
    [NOTE TO IANA/RFC-EDITOR: Please replace XXXX
with the RFC number for this specification.]
    <p>See RFCXXXX</p>
  </body>
</html>
END
```

7.2. XML Schema Registration

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

URI: urn:ietf:params:xml:ns:geopriv:held:policy

Registrant Contact: IETF, GEOPRIV working group (geopriv@ietf.org),
Richard Barnes (rbarnes@bbn.com)

Schema: The XML for this schema can be found in Section Section 4.1.

7.3. DHCP LuriType Registration

IANA is requested to add a value to the LuriTypes registry, as follows:

LuriType	Name	Reference
TBD*	Policy-URI	RFC XXXX**

* TBD is to be replaced with the assigned value

** RFC XXXX is to be replaced with this document's RFC number.

8. Operational Considerations

Associating a user's privacy preferences with a location URI can have a performance impact on the location configuration process, both in terms of protocol execution time and the state that a location server is required to store. There are additional protocol interactions (as described above), and the location server must store the user's privacy policies in addition to purely location-related state.

The mechanism that this document defines for installing policy conducts policy management actions through a separate set of interactions from the main location configuration transaction, rather than carrying policy-management messages in existing location configuration messages. This design decision imposes the cost of at least one an additional HTTP transaction on endpoints that wish to configure privacy policies. At the same time, however, it minimizes the changes that need to be made to a location configuration protocol, so that both HELD and DHCP can support policy management in basically the same fashion.

A server that supports this extension must store additional state for a location URI. By default, a location server only needs to keep location-related state for a location URI, so that it can compute location values to return in response to dereference requests. A server supporting this extension also has to store policy information. Such a server can mitigate the impact of this requirement by not storing policy information explicitly for each location URI. Until a user supplies his own policies, the server will apply a default policy, which doesn't need to be described separately for each location URI. So the amount of policy state that a server has to maintain scales as the number of users that actually

supply their own policy information. If policy URIs are constructed so that they can be associated with their corresponding location URIs algorithmically, then the server doesn't even need to maintain a table to store these associations.

Finally, a server that does not wish to be subject to any of these costs can opt not to support this extension at all. Such a server would simply never provide a "policyUri" element in a response, silently ignoring any "requestPolicyUri" element it might receive in a request.

9. Security Considerations

There are two main classes of risks associated with access control policy management: The risk of unauthorized disclosure of the protected resource via manipulation of the policy management process, and the risk of disclosure of policy information itself.

Protecting the policy management process from manipulation entails two primary requirements: First, the policy URI has to be faithfully and confidentially transmitted to the client, and second, the policy document has to be faithfully and confidentially transmitted to the Location Server. The mechanism also needs to ensure that only authorized entities are able to acquire or alter policy.

9.1. Integrity and Confidentiality for Authorization Policy Data

Each LCP ensures integrity and confidentiality through different means (see [I-D.ietf-geopriv-http-location-delivery] and [I-D.ietf-geopriv-dhcp-lbyr-uri-option]). These measures ensure that a policy URI is conveyed to the client without modification or interception.

To protect the integrity and confidentiality of policy data during management, the Location Server SHOULD provide policy URIs with the "https:" scheme and require the use of HTTP over TLS [RFC2818]. The cipher suites required by TLS [RFC5246] provide both integrity protection and confidentiality. If other means of protection are available, an "http:" URI MAY be used.

9.2. Access Control for Authorization Policy

Access control for the policy resource is based on knowledge of its URI. The URI of a policy resource operates under the same constraints as a possession model location URI [RFC5808] and is subject to the same constraints:

- o Knowledge of a policy URI MUST be restricted to authorized Rule Makers. Confidentiality is required for its conveyance in the location configuration protocol, and in the requests that are used to inspect, change or delete the policy resource.
- o The Location Server MUST ensure that the URI cannot be easily predicted. The policy URI MUST NOT be derived solely from information that might be public, including the Target identity or any location URI. The addition of random entropy increases the difficulty of guessing a policy URI.

Additional requestor authentication MAY be used for policy resources. For instance, in the particular case where the Device is identified to the Location Server by its IP address, the Location Server could use IP return routability as an additional authentication mechanism.

9.3. Location URI Allocation

A policy URI enables the authorization by access control lists model [RFC5808] for associated location URIs. Under this model, it might be possible to more widely distribute a location URI, relying on the authorization policy to constrain access to location information.

To allow for wider distribution, authorization by access control lists places additional constraints on the construction of location URIs.

If multiple Targets share a location URI, an unauthorized location recipient that acquires location URIs for the Targets can determine that the Targets are at the same location by comparing location URIs. With shared policy URIs, Targets are able to see and modify authorization policy for other Targets.

To allow for the creation of Target-specific authorization policies that are adequately privacy-protected, every location URI and policy URI that is issued to a different Target MUST be different. That is, no two clients can receive the same location URI or policy URI.

In some deployments it is not always apparent to a LCP server that two clients are different. In particular, where a middlebox [RFC3234] exists two or more clients might appear as a single client. An example of a deployment scenario of this nature is described in [RFC5687]. An LCP server MUST create a different location URI and policy URI for every request, unless the requests can be reliably identified as being from the same client.

Conversely, if a location server chooses to provide the same location URI and policy URI to multiple endpoints, then it MUST use a

restricted profile of the above protocol for policy management. (A server might do this to mitigate problems with link-layer confidentiality, e.g., for multiple clients on a shared medium.) Such a server MAY allow GET requests to allow clients to know the default policy, but it MUST NOT allow PUT or DELETE requests to control policy unless it has an out-of-band mechanism to distinguish and separately authorize clients.

10. References

10.1. Normative References

- [I-D.ietf-geopriv-dhcp-lbyr-uri-option]
Polk, J., "Dynamic Host Configuration Protocol (DHCP) IPv4 and IPv6 Option for a Location Uniform Resource Identifier (URI)", draft-ietf-geopriv-dhcp-lbyr-uri-option-09 (work in progress), October 2010.
- [I-D.ietf-geopriv-http-location-delivery]
Barnes, M., Winterbottom, J., Thomson, M., and B. Stark, "HTTP Enabled Location Delivery (HELD)", draft-ietf-geopriv-http-location-delivery-16 (work in progress), August 2009.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2616] Fielding, R., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., and T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1", RFC 2616, June 1999.
- [RFC2818] Rescorla, E., "HTTP Over TLS", RFC 2818, May 2000.
- [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.
- [RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, August 2008.

10.2. Informative References

- [I-D.ietf-geopriv-deref-protocol]
Winterbottom, J., Tschofenig, H., Schulzrinne, H., Thomson, M., and M. Dawson, "A Location Dereferencing Protocol Using HELD", draft-ietf-geopriv-deref-protocol-01 (work in progress), September 2010.
- [I-D.ietf-geopriv-held-identity-extensions]
Winterbottom, J., Thomson, M., Tschofenig, H., and R. Barnes, "Use of Device Identity in HTTP-Enabled Location Delivery (HELD)", draft-ietf-geopriv-held-identity-extensions-05 (work in progress), October 2010.
- [I-D.ietf-geopriv-policy]
Schulzrinne, H., Tschofenig, H., Morris, J., Cuellar, J., and J. Polk, "Geolocation Policy: A Document Format for Expressing Privacy Preferences for Location Information", draft-ietf-geopriv-policy-22 (work in progress), October 2010.
- [I-D.ietf-geopriv-rfc3825bis]
Polk, J., Schnizlein, J., Linsner, M., Thomson, M., and B. Aboba, "Dynamic Host Configuration Protocol Options for Coordinate-based Location Configuration Information", draft-ietf-geopriv-rfc3825bis-13 (work in progress), November 2010.
- [RFC3234] Carpenter, B. and S. Brim, "Middleboxes: Taxonomy and Issues", RFC 3234, February 2002.
- [RFC3693] Cuellar, J., Morris, J., Mulligan, D., Peterson, J., and J. Polk, "Geopriv Requirements", RFC 3693, February 2004.
- [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.
- [RFC4918] Dusseault, L., "HTTP Extensions for Web Distributed Authoring and Versioning (WebDAV)", RFC 4918, June 2007.
- [RFC5606] Peterson, J., Hardie, T., and J. Morris, "Implications of 'retransmission-allowed' for SIP Location Conveyance", RFC 5606, August 2009.

[RFC5687] Tschofenig, H. and H. Schulzrinne, "GEOPRIV Layer 7 Location Configuration Protocol: Problem Statement and Requirements", RFC 5687, March 2010.

[RFC5808] Marshall, R., "Requirements for a Location-by-Reference Mechanism", RFC 5808, May 2010.

Authors' Addresses

Richard Barnes
BBN Technologies
9861 Broken Land Parkway
Columbia, MD 21046
US

Phone: +1 410 290 6169
Email: rbarnes@bbn.com

Martin Thomson
Andrew Corporation
Andrew Building (39)
Wollongong University Campus
Northfields Avenue
Wollongong, NSW 2522
AU

Phone: +61 2 4221 2915
Email: martin.thomson@andrew.com

James Winterbottom
Andrew Corporation
Andrew Building (39)
Wollongong University Campus
Northfields Avenue
Wollongong, NSW 2522
AU

Phone: +61 242 212938
Email: james.winterbottom@andrew.com

Hannes Tschofenig
Nokia Siemens Networks
Linnoitustie 6
Espoo 02600
Finland

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

Network Working Group
Internet Draft
Intended status: Informational
Expires: September 15, 2011

S. Brim
M. Linsner
B. McLaughlin
K. Wierenga
Cisco
March 14, 2011

Mobility and Privacy
draft-brim-mobility-and-privacy-01.txt

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

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

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

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

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

This Internet-Draft will expire on September 15, 2011.

Copyright Notice

Copyright (c) 2011 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.

Abstract

Choices in Internet mobility architectures may have profound effects on privacy. This draft revisits this issue, stresses its increasing importance, and makes recommendations.

Table of Contents

1. Introduction.....	2
2. The risks of Being Traceable.....	3
3. Current Guidance on Privacy.....	4
4. Basic Mobility Requirements.....	6
5. Avoid Making a Mobile Node Traceable.....	7
6. Recommendations.....	9
7. Security Considerations.....	10
8. IANA Considerations.....	10
9. Acknowledgements.....	10
10. Normative References.....	10

1. Introduction

Significant steps are being taken right now to make the Internet's architecture more scalable and robust in routing, addressing, multihoming, mobility, including work on locator/identifier separation. However, since the Internet infrastructure is rapidly becoming an essential part of daily life for people around the world, our architectural changes need to take fundamental social issues and rights into account as a primary consideration. One of those is privacy, and in this case particularly privacy of end-user personal data. If we do not, we run the risk of colliding with established IETF principles (see for example [RFC3693]) as well as legal policy in many countries around the world.

When the Internet was designed, IP addresses were associated with timesharing machines and not with particular users. In the 1980s it began to be likely that a device and thus an IP address would be associated with a single user. Now a single IP address is very likely to be associated with a specific human being. Meanwhile, at the top of the stack, there has been a convergence of life functions using single devices using specific addresses. A person now uses his or her personal device and associated IP address for any activities: work, shopping, talking, exchanging mail and files, reading, listening to music, etc.

It is this convergence at both the top and bottom of the stack - to a single person per device and to many applications on that device - that makes the social issues more and more significant in IETF work. People use the Internet for many, more personal, activities than before. The Internet needs to fulfill the obligations expected of a communications system essential to modern human society. Our lower layer protocol designs have privacy implications beyond their intended scope.

2. The risks of Being Traceable

Issues with revealing geographic location are well-established elsewhere. For example the RAND review of the European Data Directive [RAND-EDPD] points out that "the interpretation of location data (e.g. which locations are visited, suggesting which shops are frequented, and which products and services are bought), may in the future permit the identification of the health, social, sexual or religious characteristics of the data subject" (section 3.3.1). The less well-known problem that this document focuses on is tracing the movement of mobile devices. Because mobile devices are used for so many things, any possibility of tracing them has significant, probably unpredictable, social implications, perhaps more so than revealing a single location. If an association can be made between a mobile device and a person at any location, if that device can be traced to a different geographic location then the association with the person can be inferred, usually correctly, even if the person believes they are anonymous at the new location. Consider scenarios such as:

- You are looking for a job, interviewing at other companies over your lunch hour, but you don't want your current management to know.
- You are planning a surprise gift or party for your spouse and are visiting specialty stores.
- You are a journalist gathering information on a corrupt politician from sources who wish to hide that they are dealing with you.
- You are infiltrating an organized crime ring and don't want them to know when you sneak in the back door of police headquarters.
- You are a very famous person trying to avoid paparazzi and assassins who are able to find you sporadically.

Mobility mechanisms need to take this issue into account. Obviously a mobile node must be reachable somehow, but a mobile node must be able to hide its actual movement from public view if it wishes.

3. Current Guidance on Privacy

In an attempt to define what privacy means to an end-user and the Internet, we have to start narrowing down the broad definition of "the state or condition of being free from being observed or disturbed by other people."

In this section we will examine a sampling of policies in various geographies to gain a sense of regulatory guidance around privacy. The data extracted from these policies will offer guidance in evaluating solution architectures and what pieces of data might be deemed a privacy risk.

The Internet exists within the remit of telecommunications legislation. It beholds the Internet community to be aware of and be able to adapt to the requirements of the legislative ecosystem to which our protocols and Architectures are to be deployed.

Here we will outline The European Union position as an example as it has existed for many years and has been well debated and understood globally. To be clear this is not an endorsement of specific legislation but is used merely as an example of the requirements our combined work will need operate within.

In October 1995 the EU introduced Directive 95/46/EC for the protection of individuals with regard to the processing of personal data. Included in Objective 1 of this directive is "fundamental rights and freedoms of natural persons, and in particular their right to privacy with respect to the processing of 'personal data'.

Personal data was defined as: 'personal data' shall mean any information relating to an identified or identifiable natural person ('data subject'); an identifiable person is one who can be identified, directly or indirectly, in particular by reference to an identification number or to one or more factors specific to his physical, physiological, mental, economic, cultural or social identity;

Directive 2002/58/EC included the following explicit mention of the Internet: The Internet is overturning traditional market structures by providing a common, global infrastructure for the delivery of a wide range of electronic communications services. Publicly available electronic communications services over the Internet open new

possibilities for users but also new risks for their personal data and privacy.

Also explicitly mentioned is requirement for consent for a valued added service beyond the contracted communications service.

"(30) Systems for the provision of electronic communications networks and services should be designed to limit the amount of personal data necessary to a strict minimum. Any activities related to the provision of the electronic communications service that go beyond the transmission of a communication and the billing thereof should be based on aggregated, traffic data that cannot be related to subscribers or users. Where such activities cannot be based on aggregated data, they should be considered as value added services for which the consent of the subscriber is required."

DIRECTIVE 2009/136/EC includes in section 56 explicit mention of location

"To achieve this aim, it is necessary to ensure that all fundamental rights of individuals, including the right to privacy and data protection, are safeguarded. When such devices are connected to publicly available electronic communications networks or make use of electronic communications services as a basic infrastructure, the relevant provisions of Directive 2002/58/EC (Directive on privacy and electronic communications), including those on security, traffic and location data and on confidentiality, should apply."

It should be noted that the legislative framework is evolving just as society and technology is evolving. A new principle is now proposed that rather than retrofitting privacy systems should be designed with privacy in mind. In 2009 a consultative document for the EU was published which discussed the technological requirements for Privacy by Design

"Technological standards should be developed and taken into consideration in the phase of system analysis by hardware and software engineers, so that difficulties in defining and specifying requirements deriving from the principle of 'privacy by design' are

minimized. Such standards may be general or specific with regard to various processing purposes and technologies.

4. Basic Mobility Requirements

A mobile node may need to be reachable by others, or it may act purely as a client of Internet-based services. Even if it is purely a client, it still needs at least two things:

- An authentication and authorization identifier that it can use with each access network it connects to. (Not required for open access networks.)
- A Layer 3 way for its correspondents to get packets back to it. This may no longer be simple due to potential innovations in routing architecture.

In addition, if the mobile node wants to be reachable as a peer or to offer services, it needs a few more things:

- An identifier (or identifiers) by which the node may be found by others, and a mechanism by which this identifier can be mapped to IP addresses/locators. Examples are domain names, SIP URIs, and the corresponding services.
- An IP address/locator for initially contacting the mobile node. This does not have to be associated with the mobile node's actual topological location. It can instead be associated with a rendezvous point or agent.
- A mechanism for "route optimization", whereby such an agent can be eliminated from a data path between the mobile node and a correspondent.
- An identifier or identifiers by which the mobile node can authenticate itself to its correspondents during initial contact, route optimization, and/or change of topological location. These identifiers can be at any layer, from 2 to 7. They can be associated with the mobile device's whole IP stack, individual transport sessions, or individual application instances.
- Identifiers by which the mobile node can be referred to by third parties.

If all mobile nodes are reduced to being clients only -- if they are

willing to register with servers in order to use the Internet and have others be able to reach them -- then there are fewer requirements. However, over the evolution of the Internet we have seen several times that it is not good to give up the symmetry of Internet communication and "permission-free" networking, i.e. the ability for anyone anywhere to communicate as a peer with other nodes on the Internet. For the rest of this document we assume that the IETF still wants to retain this model.

Every identifier listed above has a scope in which it needs to be known, but it is only required to be known in that scope. For example, an access authentication identifier only needs to be known to the mobile node, the access network, and a trusted third party (a mobile node's home network administration, or a bank, etc.). A session identifier only needs to be known among the parties using it, but not by the access network.

5. Avoid Making a Mobile Node Traceable

As a mobile node moves, if L3 or higher layer mobility mechanisms are used it will change its IP addresses/locators. The Internet already has sophisticated publicly available services for determining where a node is based on IP address alone. These mechanisms are not always precise or accurate, but they are in very many cases and even imprecise information is information. Protocol designers must assume that whatever IP address or locator a node has, it is likely that there is a service to turn that into a geographic location.

The tracing problem occurs when it is possible for a third party to correlate IP addresses/locators and something unique about the mobile node. Data can be gathered either through monitoring traffic or by accessing public information. It does not have to be done continuously -- periodic snapshots can make the mobile node just as vulnerable. Once the data is gathered, the third party can search for correlations.

Using identifiers for multiple purposes makes leakage of tracing information more likely. Different entities in different scopes may know different things about a mobile node or a person. Using overlapping identifiers mixes scopes and may make new, perhaps unexpected, correlations easier. For example if an access identifier such as a mobile phone's IMEI (hard-coded and not changeable, primarily used for access authentication) is also used for session continuity, or is registered in an Internet database service that is publicly accessible, changes in that device's IP addresses (and thus geographic location) can be traced.

Long-lasting identifiers make correlation easier as a device moves. They should not be used in scopes where they are not necessary.

The biggest concern is if information that makes a mobile node traceable is required to be publicly available in order for the Internet to function. If it is, it can be accessed not only without the mobile node's consent but even without its knowledge, perhaps without any audit trail of who is accessing the information that could be looked at after the fact. Some architecture for mobility and/or routing and addressing described in [I-D.irtf-rrg-recommendation] assume the use of DNS or other public mapping systems. In these, the mobile node is required to publish a mapping between its identifier and its current IP addresses/locators in order to be reachable, even if a mobile node is acting purely as a client (because otherwise packets would not get back to it). This architectural assumption removes all of the mobile node's freedom of choice about how much confidentiality to preserve -- either it exposes all of its movement to all of the world or it is simply not reachable. Public information systems like DNS are not designed to support confidentiality.

MIPv6's "home agent" [I-D.ietf-mext-rfc3775bis] is an example of how to avoid this problem: Contact with a mobile node is initially through a home agent, a rendezvous point for both data and control traffic. The home agent acts on behalf of the mobile node and encapsulates traffic to it. After an exchange of packets, the mobile node may decide, on its own, if it wants to reveal its topological location, and thus probably its geographic location, to the correspondent node. It controls its own location information. The decision to reveal it can be based on anything, including local policy.

The principle of hiding information that can expose geographic location in both data and control planes, and deferring revealing more until the mobile node or its agent decides what it wants to do, is essential. This can be included in any mobility architecture that is designed to allow it and does not insist on exposing location to a wide audience in order to gain efficiency. The obvious way to do it is an indirection mechanism such as a home agent, but this is just one way to do it. Any way will do.

Monitoring is a more subtle issue than exposure in public services, but still real, even if the mobile node is client-only. If packets contain an identifier that uniquely identifies the mobile node for some period of time, someone able to gather data on packet traffic can easily trace the mobile node's movements as the IP address/locator changes. It is not necessary for the watcher to be

able to gather this information in real time if it can access logs gathered by others. Here, approaches to the problem are more difficult to define because there is a conflict between three goals: to avoid overhead, to preserve session continuity with low delay, and to keep control over location information. Some designs such already try to find their balance. All protocol work should consider the tradeoffs with privacy and explicitly find a balance point.

6. Recommendations

Members of the Internet community who are creating or reviewing proposed architectural changes, particularly in mobility but also in other areas that impinge on mobility such as routing and addressing, should consider the following points:

- Architectural changes MUST avoid requiring the exposure of a mapping between any of a node's identifiers and IP addresses/locators to unknown observers. If they require exposure, they will experience a head-on collision with basic principles of the IETF and with privacy policies around the world. It will simply not be acceptable to require the loss of this much individual privacy.
- An architectural proposal MAY make it possible to use public information systems to optimize traffic flow, but ideally it should do so without sacrificing privacy. If it cannot do so without sacrificing privacy, the default case built into the architecture SHOULD be to preserve privacy instead of optimizing. The reason is that most users will not change defaults, and the default be one of privacy, only moving away from it by customer choice.
- If possible, information about who is gathering data about a user SHOULD be available to that user. Everyone deserves to know who is watching them.
- Proposals SHOULD address the issue of loss of geographic location privacy due to monitoring of packets crossing the Internet, and find an explicit balance between conflicting goals.
- Protocols SHOULD avoid using identifiers for multiple purposes. Different identifier scopes do not need to overlap. Confidentiality boundaries can be established by clearly defining limited interfaces.
- Protocols SHOULD avoid using long-lasting identifiers in scopes

where they are not necessary.

7. Security Considerations

In a sense this entire document is about security.

8. IANA Considerations

This document makes no request of IANA

Note to RFC Editor: this section may be removed on publication as an RFC.

9. Acknowledgements

Thanks to many with whom we have discussed this issue in recent months.

This document was prepared using 2-Word-v2.0.template.dot.

10. Normative References

[I-D.ietf-mext-rfc3775bis]

Perkins, C., Johnson, D., and J. Arkko, "Mobility Support in IPv6", draft-ietf-mext-rfc3775bis-08 (work in progress), October 2010.

[I-D.irtf-rrg-recommendation]

Li, T., "Recommendation for a Routing Architecture", draft-irtf-rrg-recommendation-14 (work in progress), September 2010.

[RAND-EDPD]

Robinson, N., Graux, H., Botterman, M., and L. Valeri, "Review of the European Data Protection Directive", May 2009.

[RFC2119]

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

[RFC3693]

Cuellar, J., Morris, J., Mulligan, D., Peterson, J., and J. Polk,

"Geopriv Requirements", RFC 3693, February 2004.

Authors' Addresses

Scott Brim
Cisco

Email: scott.brim@gmail.com

Marc Linsner
Cisco

Email: mlinsner@cisco.com

Bryan McLaughlin
Cisco

Email: brmclaug@cisco.com

KlaasWierenga
Cisco

Email: kwiereng@cisco.com

GEOPRIV
Internet Draft
Intended status: Informational

Expires: April 2011

C. Hoene
A. Krebs
C. Behle
M. Schmidt
Universitaet Tuebingen
October 14, 2010

Bayesian Location Identifier
draft-hoene-geopriv-bli-00.txt

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

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

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

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

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

This Internet-Draft will expire on April 14, 2011.

Copyright Notice

Copyright (c) 2010 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.

Abstract

Location Generators cannot always provide exact measures of particular locations. Instead, they estimate the location of objects. More precisely, they use filters to aggregate noisy sensor data and to calculate probability density distributions of estimated positions. In location tracking applications, typically Kalman-type, Gaussian-Sum, and Particle Filters are used.

We believe that it is reasonable to use the outputs of those filters to describe a location estimate and its uncertainty, because they are the natural result of location tracking algorithms. In addition, the results of those filters can be feed into sensor fusion and decision making engines easily.

Geometric representations such as polygons or ellipses might be demanded by an application. The output of filters can be converted to those application demanded shapes. However, these conversions come at the loss of precision and are not well understood scientifically. Thus, we think that transmitting filter results is a solution that is easier to implement.

In this draft, we present a transmission format for PIDF-LO, which is based on the output of Kalman-type, Gaussian-Sum, and Particle Filters.

Table of Contents

1. Introduction.....	3
2. Overview on Filters.....	3
2.1. Kalman Filters.....	4
2.2. Particle Filter.....	4
2.3. Gaussian Sum Particle Filters.....	5
3. Coordinate System and Datum.....	5
4. Examples.....	5
4.1. Kalman Filter Results.....	5
4.2. Particle Filter Results.....	6
4.3. Gaussian Sum Filter.....	7
4.4. Well-know Reference Frame.....	7
4.5. Relative Datum.....	8
5. Security Considerations.....	9
6. IANA Considerations.....	9
7. Conclusions.....	9
8. References.....	9
8.1. Informative References.....	9

1. Introduction

The location of an object cannot be measured precisely always. Especially, in an indoor environment numerous sources of measurement errors lead to measurement results, which are uncertain in a high degree. To represent this uncertainty, a previous draft [1] defines how uncertainty and its associated datum and confidence is expressed and interpreted. To simplify the representation, the draft limits the description of uncertainty to a number of well defined shapes, e.g. one point, a centroid, an arc-band centroid, or polygons. However, state of the art multimodal location tracking algorithms do not provide any location results of the above mentioned shapes. Instead, they typically applied some sort of Kalman or Particle Filters, which assume a Gaussian or an arbitrary error distribution. Thus, in this draft we propose to transmit as location information the results of Gaussian distributions or Particle Filters to present uncertainty.

This has the following advantages:

1. Any kind of uncertainty can be transmitted. More precisely, any form of probability distribution, which represents uncertainty, can be represented.
2. A conversion from particles to shapes is not required, which maintains precision and eases implementations.
3. Kalman and Particle filters are well understood statistical tools based on research of control theory, signal processing, Monte-Carlo simulations and Bayesian statistics. Numerous location tracking algorithms have been developed that work with those filters. Thus, the transmission format defined in this draft is based on a profound scientific basis.

2. Overview on Filters

Location tracking is based on physical measurements, which estimate time of flight, signal strength, angle of arrives, motions, objects in images, and many other forms of sensor input. All these sensor measurements are subjected to measurement noise. Because of that, filters are used to estimate the real value of the measurement despite the fact of measurement results that are subjected to noise.

Many filter types have been developed. However, in location tracking typically only a few are applied. These include different types of Kalman-Filters, filters that work with one or multiple Gaussian

Normal distributions, and Particle Filters. These filters are briefly described in the following.

2.1. Kalman Filters

A Kalman Filter uses a system model to estimate the probability of changes. This data is combined with a model of measurement data and control input, if any, to estimate the true value of the parameters under study. It only allows linear relations between filter variables and assumes Gaussian noise distributions. Despite that, it is very robust in many applications.

The result of a Kalman filter is a posteriori state vector (for example a location) and a posteriori estimated covariance. The state is given by a vector $\hat{x}_{k|k}$ and the covariance by $P_{k|k}$. Both estimates are given for the time index k [2]. If the vector has a dimension of d (for example 3 for xyz), then the covariance matrix has a size of $d*d$.

We suggest that, as PIDF-LO object, all these three variables shall be transmitted to indicate a position estimate, its distribution, and the time of measurement.

Also, this format should work well for non-linear filters such as the Extended or Unscented Kalman Filter.

2.2. Particle Filter

Particle Filter, also called sequential Monte Carlo methods (SMC), have the advantage that arbitrary distributions can be approximated [3]. As such, they approximate Bayesian models, which consist of probability distribution functions, which define the degree of "believe" to which a particular value is true.

Particle filters approximate probability distribution function with a number of particles. More particles are placed at positions that are more likely. Each particle has Dirac shape.

The a posteriori state of a particle filter is approximated by M particles called $x^{(M)}_i$, which are weighted with $w^{(m)}_i$. The PIDF-LO transmission object should contain the particles, their weights and again the time index k .

2.3. Gaussian Sum Particle Filters

Here we assume that the probability distributions are described by the sum of normal distributions [4]. As such, it can be seen as the combination of two previously mentioned filtering approaches.

The a posteriori state of a Gaussian Sum Filter is approximated by M Gaussian distributions called $\hat{x}^{(M)}_k$, which are weighted with $w^{(m)}_k$ and have distributions described by covariance matrices $P^{(M)}_{k|k}$.

Again, all those parameters shall be transmitted.

3. Coordinate System and Datum

Any location is relative to a frame of reference. The frame of reference defines the position, orientation and other properties of a coordinate system, in which an object is located. A number of geodetic reference frames have been defined such as WGS84, ETRS89, or ITRF2005. Typically, they define the reference point and the orientation of the coordinate system.

In robotics, reference frames are used, too. They are referencing to a zero point, have an orientation, and may be scaled, mirrored, rotated. For example, a so called transformation matrix can be applied to the location vector to transform coordinate systems.

Commonly in navigation, besides Cartesian also Polar coordinate systems are used. In addition, a polar coordinate system has the benefit that - for example - circular bands can be described easily if the rotating angles have a high uncertainty or are not defined.

In summary, to describe the location of an object, the reference frame has to be named or defined, and the type of coordinate system must be given.

4. Examples

This section shows examples on how to transmit the location ID described above.

4.1. Kalman Filter Results

Assuming, a Kalman filter estimates a position and assigns an uncertainty to this estimate. Then

```
<gp:geopriv>
  <gp:location-info>
    <gs:Kalman srsName="urn:ogc:def:crs:EPSG:??????">
      <gml:pos>-34.407242 150.882518 34</gml:pos>
      <gml:covariance>
        1 0 0
        0 4 0
        0 0 16
      </gml:covariance>
      <gms:timestamp>
        102000
      </gms:timestamp>
    </gs:Kalman>
  </gp:location-info>
</gp:geopriv>
```

defines a point at [-34.407242, 150.882518, 34] that has a Gaussian distribution with the standard deviation of 1, 4 and 16 for the X, Y, and Z-axes.

In addition, it states that the location has been estimated at a time stamp defined by the time index 102000.

4.2. Particle Filter Results

Next, we assume that a Particle Filter has calculated three particles. Then

```
<gp:geopriv>
  <gp:location-info>
    <gs:Particle srsName="urn:ogc:def:crs:EPSG:??????">
      <gml:particleList>
        -34.407242 150.882518 34 0.5
        -34.500000 150.000000 34 0.25
        -34.400000 151.000000 34 0.25
      </gml:particleList>
    </gs:Particle>
  </gp:location-info>
</gp:geopriv>
```

defines a point at [-34.407242, 150.882518, 34] with a weight of 0.5, a point [-34.5,150,34] with a weight of 0.25, and a point at [-34.4,151,34] with a weight of 0.25.

4.3. Gaussian Sum Filter

Next, we have a Gaussian Sum Filter with two Gaussian distributions. Then

```
<gp:geopriv>
  <gp:location-info>
    <gs:Particle srsName="urn:ogc:def:crs:EPSG:??????">
      <gml:particleList>
        -34.407242 150.882518 34 0.25
        -34.500000 150.000000 34 0.25
      </gml:particleList>
      <gml:covarianceList>
        1 0 0
        0 4 0
        0 0 16
        1 0.5 0.5
        0.5 1 0.5
        0.5 0.5 1
      </gml:covariance>
      <gml:particleList>
    </gs:Particle>
  </gp:location-info>
</gp:geopriv>
```

defines a multi-vector Gaussian distribution with the center at $[-34.407242, 150.882518, 34]$, a weight of 0.25, and standard deviations of 1, 4 and 16. In addition, a second distribution is at $[-34.5, 150, 34]$ with a weight of 0.25 and a covariance matrix of $[[1 \ 0.5 \ 0.5] [0.5 \ 1 \ 0.5] [0.5 \ 0.5 \ 1]]$.

In addition, because the sum of weights is lower than 1, we assume that the position estimate has only a belief of correctness (according to Bayesian network theory) of 0.5.

4.4. Well-know Reference Frame

Now, we extend the example given in Section 4.1. to well define a reference frame that is based on the UTM coordinate system.

```
<gp:geopriv>
  <gp:location-info name="Reference Point ID=0123456789">
    <gs:ReferenceFrame>
      UTM zone=32U
      EVRS2000
    </gs:ReferenceFrame>
    <gs:Kalman srsName="urn:ogc:def:crs:EPSG:??????">
```

```

    <gml:x>-34.407242</gml:x>
    <gml:y>150.882518</gml:y>
    <gml:z>34</gml:z>
    <gml:covariance>
      1 0 0
      0 4 0
      0 0 16
    </gml:covariance>
    <gms:timestamp>
      102000
    </gms:timestamp>
  </gs:Kalman>
</gp:location-info>
</gp:geopriv>

```

Here, the reference frame for the X and Y axes is given by an UTM map in the zone 32U (western part of Germany) and the height (Z axis) is given by the European Vertical Reference System (EVRS) based on the Normaal Amsterdams Peil (NAP).

In addition, we name this location "Reference Point ID=0123456789".

4.5. Relative Datum

Next, we assume that a reference frame is defined. This example is based on the particle filter given in Section 4.2.

```

<gp:geopriv>
  <gp:location-info name="Reference Point ID=0123456789">
    <gs:ReferenceFrameDefinition>
      "Reference Point ID=0123456789"
      1 0 0 0
      0 0 -1 0
      0 1 0 0
      0 0 0 1
      cartesian
    </gs:ReferenceFrameDefinition>
    <gs:Particle srsName="urn:ogc:def:crs:EPSG:??????">
      <gml:particleList>
        -34.407242 150.882518 34 0.5
        -34.500000 150.000000 34 0.25
        -34.400000 151.000000 34 0.25
      </gml:particleList>
    </gs:Particle>
  </gp:location-info>
</gp:geopriv>

```

The new coordinates are based on the reference point given in the previous section but it is rotated by 90 degree about their common normal, from old Z axis to new Z axis.

5. Security Considerations

Security issues have not yet been identified.

6. IANA Considerations

Well-known reference systems must be named or numbered. Thus might require a registration at IANA.

7. Conclusions

This initial draft presented a transmission format for uncertain, relative and transformed location estimates.

It is aim as a basis of discussion, because we believe that uncertainty shall be directly presented by the results of algorithms that determine uncertainty.

Questions that need to be address are:

- o Are the most common filter types covered?
- o Is the transmission format efficient? Especially, if many particle needs to be transmitted an XML description might cause too much overhead. Instead, a generic XML compression such as Binary XML can be used or an application-specific compression algorithm can be defined.
- o Does the representation of relative and transformation reference systems fit into the GEOPRIV framework?
- o Shall the time be given as an index or as a fourth dimension?

Any comments to enhance this draft are highly welcomed.

8. References

8.1. Informative References

- [1] M. Thomson, J. Winterbottom, "Representation of Uncertainty and Confidence in PIDF-LO", draft-thomson-geopriv-uncertainty-05, work in progress, June 1, 2010.

- [2] Wikipedia contributors, "Kalman filter", Publisher: Wikipedia, The Free Encyclopedia, Date of last revision: 7 October 2010 15:52 UTC, http://en.wikipedia.org/w/index.php?title=Kalman_filter&oldid=389338611

- [3] Wikipedia contributors, "Particle filter", Publisher: Wikipedia, The Free Encyclopedia, Date of last revision: 24 September 2010 22:18 UTC, http://en.wikipedia.org/w/index.php?title=Particle_filter&oldid=386830379

- [4] J.H. Kotecha, P.M. Djuric, "Gaussian sum particle filtering", IEEE Transactions on Signal Processing, Volume: 51, Issue:10, pages 2602 - 2612, Oct. 2003.

Authors' Addresses

Christian Hoene
Universitaet Tuebingen
WSI-ICS
Sand 13
72076 Tuebingen
Germany

Phone: +49 7071 2970532
Email: hoene@uni-tuebingen.de

Andreas Krebs
Universitaet Tuebingen
WSI
Sand 13
72076 Tuebingen
Germany

Phone: +49 7071 2977476
Email: krebs@informatik.uni-tuebingen.de

Christoph Behle
Universitaet Tuebingen
WSI
Sand 13
72076 Tuebingen
Germany

Phone: +49 7071 2977565
Email: behlec@informatik.uni-tuebingen.de

Mark Schmidt
Universitaet Tuebingen
WSI-ICS
Sand 13
72076 Tuebingen
Germany

Phone: +49 7071 2970510
Email: mark.schmidt@wsii.uni-tuebingen.de

GEOPRIV
Internet-Draft
Intended status: Standards Track
Expires: January 15, 2013

J. Winterbottom
Commscope
H. Tschofenig
Nokia Siemens Networks
H. Schulzrinne
Columbia University
M. Thomson
Microsoft
July 14, 2012

A Location Dereferencing Protocol Using HELD
draft-ietf-geopriv-deref-protocol-07

Abstract

This document describes how to use the Hypertext Transfer Protocol (HTTP) over Transport Layer Security (TLS) as a dereferencing protocol to resolve a reference to a Presence Information Data Format Location Object (PIDF-LO). The document assumes that a Location Recipient possesses a URI that can be used in conjunction with the HTTP-Enabled Location Delivery (HELD) protocol to request the location of the Target.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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

This Internet-Draft will expire on January 15, 2013.

Copyright Notice

Copyright (c) 2012 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	3
2. Terminology	3
3. HELD Dereference Protocol	4
3.1. HELD Usage Profile	4
3.2. HTTP GET Behavior	5
4. Authorization Models	6
4.1. Authorization by Possession	7
4.2. Authorization via Access Control	8
4.3. Access Control with HELD Deference	8
5. Examples	9
6. Security Considerations	12
7. IANA Considerations	13
8. Acknowledgements	14
9. References	14
9.1. Normative References	14
9.2. Informative references	15
Appendix A. GEOPRIV Using Protocol Compliance	16
Appendix B. Compliance to Location Reference Requirements	19
B.1. Requirements for a Location Configuration Protocol	20
B.2. Requirements for a Location Dereference Protocol	21
Authors' Addresses	22

1. Introduction

A location URI [RFC5808] identifies a resource that contains the location of an entity. This document specifies how a holder of an "http:" or "https:" location URI uses that URI to retrieve location information.

A location URI can be acquired using a location configuration protocol, such as HTTP-Enabled Location Delivery (HELD) [RFC5985] or the Dynamic Host Configuration Protocol (DHCP) location URI option [I-D.ietf-geopriv-dhcp-lbyr-uri-option].

A Location Recipient that dereferences a location URI acquires location information in the of a Presence Information Data Format - Location Object (PIDF-LO) document [RFC4119]. HELD parameters allow for specifying the type of location information, though some constraints are placed on allowable parameters.

Location URIs compatible with HELD dereferencing use the "https:" or "http:" scheme. HELD can be used by Location Recipients that are aware of the fact that the URI is a location URI. Mandatory support for an HTTP GET request ensures that the URI can be used even if it is not recognized as a location URI.

2. Terminology

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

This document uses key terminology from several sources:

- o terms for the GEOPRIV reference model defined in [RFC6280];
- o the term Location Information Server (LIS), from [RFC5687], is a node in the access network that provides location information to an end point; a LIS provides location URIs;
- o the term Location Server (LS), from [RFC6280], is used to identify the role that responds to a location dereference request; this might be the same entity as the LIS, but the model in [RFC5808] allows for the existence of separate - but related - entities; and
- o the term location URI is coined in [RFC5808].

3. HELD Dereference Protocol

This section describes how HELD can be used to dereference a location URI. This process can be applied when a Location Recipient is in possession of a location URI with a "https:" or "http:" URI scheme.

This document does not describe a specific authentication mechanism. This means that authorization policies are unable to specifically identify authorized Location Recipients.

A Location Recipient that wishes to dereference an "https:" or "http:" URI performs a HELD request on HTTP to the identified resource.

Note: In many cases, an "http:" URI does not provide sufficient security for location URIs. The absence of the security mechanisms provided by TLS means that the Rule Maker has no control over who receives location information and the Location Recipient has no assurance that the information is correct.

The Location Recipient establishes a connection to the LS, as described in [RFC2818].

The scheme of a location URI determines whether or not TLS is used on a given dereference transaction. Location Servers MUST be configured to issue only HTTPS URIs and respond to only to HTTPS dereference requests, unless confidentiality and integrity protection are provided by some other mechanism. For example, the server might only accept requests from clients within a trusted network, or via an IPsec-protected channel. When TLS is used, the TLS ciphersuite TLS_NULL_WITH_NULL_NULL MUST NOT be used and the LS MUST be authenticated [RFC6125] to ensure that the correct server is contacted.

A Location Server MAY reject a request and request that a Location Recipient provide authentication credentials if authorization is dependent on the Location Recipient identity. Future specifications could define an authentication mechanism and a means by which Location Recipients are identified in authorization policies. This document provides definitions for neither item.

3.1. HELD Usage Profile

Use of HELD as a location dereference protocol is largely the same as its use as a location configuration protocol. Aside from the restrictions noted in this document, HELD semantics do not differ from those established in [RFC5985].

The HELD "locationRequest" is the only request permitted by this specification. Similarly, request parameters other than the following MUST NOT be accepted by the LS: "responseTime", "locationType" (including the associated "exact" attribute).

Parameters and requests that do not have known behaviour for dereference requests MUST NOT be used. The LS MUST ignore any parameters that it does not understand unless it knows the parameters to be invalid. If parameters are understood by the LS and known to be invalid, the LS MAY generate a HELD error response. For instance, those defined in [RFC6155] are always invalid and can be rejected.

The LS MUST NOT generate location URIs or provide a "locationUriSet" in response to a dereference request. If the location request contains a "locationType" element that includes "locationURI", this parameter is either ignored or rejected as appropriate, based on the associated "exact" attribute.

3.2. HTTP GET Behavior

GET is the method assumed by generic HTTP user agents, therefore unless context identifies an "https:" URI as a HELD URI, such a user agent might simply send an HTTP GET. Rather than providing an HTTP 405 (Method Not Allowed) response indicating that POST is the only permitted method, a LIS MUST provide a HELD location response if it receives an HTTP GET request.

An HTTP GET request to a HELD URI produces a HELD response as if the following HELD request had been sent using HTTP POST:

```
<locationRequest xmlns="urn:ietf:params:xml:ns:geopriv:held">
  <locationType exact="false">
    geodetic civic
  </locationType>
</locationRequest>
```

Figure 1: GET Request Equivalent Location Request

HTTP GET requests MUST be safe and idempotent [RFC2616] - that is, there are no side-effects of making the request and a repeated request has no more effect than a single request. Repeating a HELD request might result in a different location, but only as a result of a change in the state of the resource: the location of the Target.

Only the creation of a location URI as a result of receiving a request causes a HELD request to have side-effects. A request to a location URI can be both safe and idempotent, since a location URI cannot be produced in response to a request to a location URI.

A Location Recipient MAY infer from a response containing the HELD content type, "application/held+xml", that a URI references a resource that supports HELD.

Content negotiation MAY be supported to produce a presence document in place of a HELD location response. Where the presence document would otherwise be included in a "locationResponse" document, it can be included in the body of the HTTP response directly by including an "Accept" header that includes "application/pidf+xml".

4. Authorization Models

This section discusses two extreme types of authorization models for dereferencing with HELD URIs, namely "Authorization by Possession" and "Authorization by Access Control". In the subsequent subsections we discuss the properties of these two models. Figure 2, from [RFC5808], shows the model applicable to location configuration, conveyance and dereference.

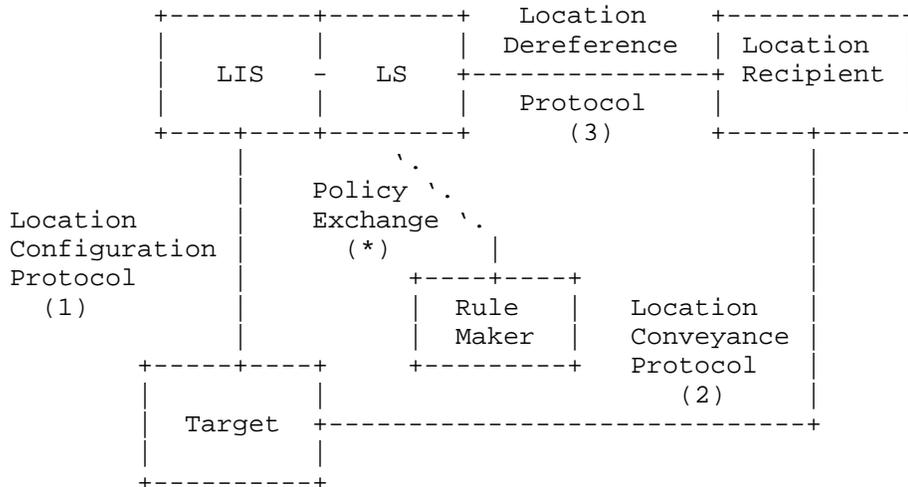


Figure 2: Communication Model

It is important to note that this document does not mandate a specific authorization model. It is possible to combine aspects of both models. However, no authentication framework is provided, which limits the policy options available when the "Authorization by Access Control" model is used.

For either authorization model, the overall process is similar. The following steps are followed, with minor alterations:

1. The Target acquires a location URI from the LIS. This uses a location configuration protocol (LCP), such as HELD or DHCP.
2. The Target then conveys the location URI to a third party, the Location Recipient (for example using SIP as described in [RFC6442]). This step is shown in (2) of Figure 2.
3. The Location Recipient then needs to dereference the location URI in order to obtain the Location Object (3). An "https:" or "http:" URI is dereferenced as described in this document; other URI schemes might be dereferenced using another method.

In this final step, the Location Server (LS) or LIS makes an authorization decision. How this decision is reached depends on the authorization model.

4.1. Authorization by Possession

In this model, possession - or knowledge - of the location URI is used to control access to location information. A location URI might be constructed such that it is hard to guess (see C8 of [RFC5808]) and the set of entities that it is disclosed to can be limited. The only authentication this would require by the LS is evidence of possession of the URI. The LS could immediately authorize any request that indicates this URI.

Authorization by possession does not require direct interaction with a Rule Maker; it is assumed that the Rule Maker is able to exert control over the distribution of the location URI. Therefore, the LIS can operate with limited policy input from a Rule Maker.

Limited disclosure is an important aspect of this authorization model. The location URI is a secret; therefore, ensuring that adversaries are not able to acquire this information is paramount. Encryption, such as might be offered by TLS [RFC5246] or S/MIME [RFC5751], protects the information from eavesdroppers.

Use of authorization by possession location URIs in a hop-by-hop protocol such as SIP [RFC3261] adds the possibility of on-path adversaries. Depending on the usage of the location URI for certain location based applications (e.g., emergency services, location based routing) specific treatment is important, as discussed in [RFC6442].

Using possession as a basis for authorization means that, once granted, authorization cannot be easily revoked. Cancellation of a location URI ensures that legitimate users are also affected; application of additional policy is theoretically possible, but could be technically infeasible. Expiration of location URIs limits the

usable time for a location URI, requiring that an attacker continue to learn new location URIs to retain access to current location information.

A very simple policy might be established at the time that a location URI is created. This policy specifies that the location URI expires after a certain time, which limits any inadvertent exposure of location information to adversaries. The expiration time of the location URI might be negotiated at the time of its creation, or it might be unilaterally set by the LIS.

4.2. Authorization via Access Control

Use of explicit access control provides a Rule Maker greater control over the behaviour of an LS. In contrast to authorization by possession, possession of this form of location URI does not imply authorization. Since an explicit policy is used to authorize access to location information, the location URI can be distributed to many potential Location Recipients.

Either before creation or dissemination of the location URI, the Rule Maker establishes an authorization policy with the LS. In reference to Figure 2, authorization policies might be established at creation (Step 1), and need to be established before the location URI is published (Step 2) to ensure that the policy grants access to the desired Location Recipients. Depending on the mechanism used, it might also be possible to change authorization policies at any time.

A possible format for these authorization policies is available with GEOPRIV Common Policy [RFC4745] and Geolocation Policy [I-D.ietf-geopriv-policy]. Additional constraints might be established by other means.

The LS enforces the authorization policy when a Location Recipient dereferences the URI. Explicit authorization policies allow a Rule Maker to specify how location information is provided to Location Recipients.

4.3. Access Control with HELD Deference

This document does not describe a specific authentication mechanism; therefore, the authorization by access control model is not an option. Instead, this document assumes the authorization by possession model.

Other policy mechanisms, such as those described in [I-D.ietf-geopriv-policy], can be applied for different Location Recipients if each recipient is given a different location URIs.

Each location URI can be assigned different authorization policy. Selective disclosure used in this fashion can be used in place of identity-based authorization.

How policy is associated with a location URI is not defined by this document. [I-D.ietf-geopriv-policy-uri] describes one possible mechanism.

Use of identity-based authorization policy is not precluded. A Location Server MAY support an authentication mechanism that enables identity-based authorization policies to be used. Future specifications might define means of identifying recipients.

Note: Policy frameworks like [RFC4745] degrade in a way that protects privacy if features are not supported. If a policy specifies a rule that is conditional on the identity of a recipient and the protocol does not (or cannot) provide an assertion identity of the recipient, the rule has no effect and the policy defaults to providing less information.

5. Examples

An example scenario envisioned by this document is shown in Figure 3. This diagram shows how a location dereference protocol fits with location configuration and conveyance. [RFC5808] contains more information on this scenario and others like it.

Again, this is identical to the response in Section 10.1 of [RFC5985] - unless policy specifies otherwise, the Location Recipient receives the same information as the Device.

```
HTTP/1.1 200 OK
Server: Example LIS
Date: Mon, 10 Jan 2011 03:42:29 GMT
Expires: Tue, 11 Jan 2011 03:42:29 GMT
Cache-control: private
Content-Type: application/held+xml
Content-Length: 676
```

```
<?xml version="1.0"?>
<locationResponse xmlns="urn:ietf:params:xml:ns:geopriv:held">
<presence xmlns="urn:ietf:params:xml:ns:pidf"
  entity="pres:3650n87934c@ls.example.com">
  <tuple id="b650sf789nd">
    <status>
      <geopriv xmlns="urn:ietf:params:xml:ns:pidf:geopriv10"
        xmlns:gbp="urn:ietf:params:xml:ns:pidf:geopriv10:basic-policy">
        <location-info>
          <Point xmlns="http://www.opengis.net/gml"
            srsName="urn:ogc:def:crs:EPSG::4326">
            <pos>-34.407 150.88001</pos>
          </Point>
        </location-info>
        <usage-rules>
          <gbp:retransmission-allowed>
            false</gbp:retransmission-allowed>
          <gbp:retention-expiry>
            2011-01-11T03:42:29+00:00</gbp:retention-expiry>
          </usage-rules>
          <method>Wiremap</method>
        </geopriv>
      </status>
      <timestamp>2006-01-10T03:42:28+00:00</timestamp>
    </tuple>
  </presence>
</locationResponse>
```

Figure 5: Response with Location Information

The following GET request is treated in an equivalent fashion. The LS treats this request as though it were a location request of the form shown in Figure 1. The same response might be provided.

```
GET /uri/w3g61nf5n66p0 HTTP/1.1
Host: ls.example.com:49152
Accept: application/held+xml
```

Figure 6: GET Request

The following GET request uses content negotiation to indicate a preference for a presence document.

```
GET /uri/w3g61nf5n66p0 HTTP/1.1
Host: ls.example.com:49152
Accept: application/pidf+xml,application/held+xml;q=0.5
```

Figure 7: GET Request with Content Negotiation

The response only differs from a normal HELD location response to a POST request in that the "locationResponse" element is omitted and the "Content-Type" header reflects the changed content.

```
HTTP/1.1 200 OK
Server: Example LIS
Date: Mon, 10 Jan 2011 03:42:29 GMT
Expires: Tue, 11 Jan 2011 03:42:29 GMT
Cache-control: private
Content-Type: application/pidf+xml
Content-Length: 591
```

```
<?xml version="1.0"?>
<presence xmlns="urn:ietf:params:xml:ns:pidf"
  entity="pres:3650n87934c@ls.example.com">
  <!-- PIDF contents are identical to the previous example -->
</presence>
```

Figure 8: GET Response with PIDF-LO

6. Security Considerations

Privacy of location information is the most important security consideration for this document. Two measures in particular are used to protect privacy: TLS and authorization policies. TLS provides a means of ensuring confidentiality of location information through

encryption and mutual authentication. An authorization policy allows a Rule Maker to explicitly control how location information is provided to Location Recipients.

The process by which a Rule Maker establishes an authorization policy is not covered by this document; several methods are possible, for instance: [I-D.ietf-geopriv-policy-uri], [RFC4825].

TLS MUST be used for dereferencing location URIs unless confidentiality and integrity are provided by some other mechanism, as discussed in Section 3. Location Recipients MUST authenticate the host identity using the domain name included in the location URI, using the procedure described in Section 3.1 of [RFC2818]. Local policy determines what a Location Recipient does if authentication fails or cannot be attempted.

The authorization by possession model (Section 4.1) further relies on TLS when transmitting the location URI to protect the secrecy of the URI. Possession of such a URI implies the same privacy considerations as possession of the PIDF-LO document that the URI references.

Location URIs MUST only be disclosed to authorized Location Recipients. The GEOPRIV architecture [RFC6280] identifies the Rule Maker role as being the entity that authorizes disclosure of this nature.

Protection of the location URI is necessary, since the policy attached to such a location URI permits any who have the URI to view it. This aspect of security is covered in more detail in the specification of location conveyance protocols, such as [RFC6442].

The LS MUST NOT provide any information about the Target except its location, unless policy from a Rule Maker allows otherwise. In particular, the requirements in [RFC5808] mandate this measure to protect the identity of the Target. To this end, an unlinked pseudonym MUST be provided in the "entity" attribute of the PIDF-LO document.

Further security considerations and requirements relating to the use of location URIs are described in [RFC5808].

7. IANA Considerations

This document makes no request of IANA.

[[IANA/RFC-EDITOR: Please remove this section before publication.]]

8. Acknowledgements

Thanks to Barbara Stark and Guy Caron for providing early comments. Thanks to Rohan Mahy for constructive comments on the scope and format of the document. Thanks to Ted Hardie for his strawman proposal that provided assistance with the security section of this document. Richard Barnes made helpful observations on the application of authorization policy. Bernard Aboba and Julian Reschke contributed constructive reviews.

The participants of the GEOPRIV interim meeting 2008 provided significant feedback on this document.

James Polk provided input on security in June 2008.

Martin Dawson was an original author of this document. Sadly, he passed away prior to its publication.

9. References

9.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2616] Fielding, R., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., and T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1", RFC 2616, June 1999.
- [RFC2818] Rescorla, E., "HTTP Over TLS", RFC 2818, May 2000.
- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, January 2005.
- [RFC4119] Peterson, J., "A Presence-based GEOPRIV Location Object Format", RFC 4119, December 2005.
- [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.
- [RFC5985] Barnes, M., "HTTP-Enabled Location Delivery (HELD)", RFC 5985, September 2010.
- [RFC6125] Saint-Andre, P. and J. Hodges, "Representation and

Verification of Domain-Based Application Service Identity within Internet Public Key Infrastructure Using X.509 (PKIX) Certificates in the Context of Transport Layer Security (TLS)", RFC 6125, March 2011.

9.2. Informative references

- [I-D.ietf-geopriv-dhcp-lbyr-uri-option]
Polk, J., "Dynamic Host Configuration Protocol (DHCP) IPv4 and IPv6 Option for a Location Uniform Resource Identifier (URI)", draft-ietf-geopriv-dhcp-lbyr-uri-option-15 (work in progress), May 2012.
- [I-D.ietf-geopriv-policy]
Schulzrinne, H., Tschofenig, H., Cuellar, J., Polk, J., Morris, J., and M. Thomson, "Geolocation Policy: A Document Format for Expressing Privacy Preferences for Location Information", draft-ietf-geopriv-policy-26 (work in progress), June 2012.
- [I-D.ietf-geopriv-policy-uri]
Thomson, M., Winterbottom, J., Barnes, R., and H. Tschofenig, "Location Configuration Extensions for Policy Management", draft-ietf-geopriv-policy-uri-04 (work in progress), November 2011.
- [RFC3261] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., and E. Schooler, "SIP: Session Initiation Protocol", RFC 3261, June 2002.
- [RFC3693] Cuellar, J., Morris, J., Mulligan, D., Peterson, J., and J. Polk, "Geopriv Requirements", RFC 3693, February 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.
- [RFC4825] Rosenberg, J., "The Extensible Markup Language (XML) Configuration Access Protocol (XCAP)", RFC 4825, May 2007.
- [RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, August 2008.
- [RFC5687] Tschofenig, H. and H. Schulzrinne, "GEOPRIV Layer 7 Location Configuration Protocol: Problem Statement and Requirements", RFC 5687, March 2010.

- [RFC5751] Ramsdell, B. and S. Turner, "Secure/Multipurpose Internet Mail Extensions (S/MIME) Version 3.2 Message Specification", RFC 5751, January 2010.
- [RFC5808] Marshall, R., "Requirements for a Location-by-Reference Mechanism", RFC 5808, May 2010.
- [RFC6155] Winterbottom, J., Thomson, M., Tschofenig, H., and R. Barnes, "Use of Device Identity in HTTP-Enabled Location Delivery (HELD)", RFC 6155, March 2011.
- [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.
- [RFC6442] Polk, J., Rosen, B., and J. Peterson, "Location Conveyance for the Session Initiation Protocol", RFC 6442, December 2011.

Appendix A. GEOPRIV Using Protocol Compliance

This section describes how use of HELD as a location dereference protocol complies with the GEOPRIV requirements described in [RFC3693].

Req. 1. (Location Object generalities):

This section relates to the PIDF-LO [RFC4119] document, which is used by HELD. These requirements are addressed by [RFC4119] and [RFC5491].

Req. 2. (Location Object fields):

This section relates to the PIDF-LO [RFC4119] document, which is used by HELD. These requirements are addressed by [RFC4119] and [RFC5491].

Req. 3. (Location Data Types):

This section relates to the PIDF-LO [RFC4119] document, which is used by HELD. These requirements are addressed by [RFC4119] and [RFC5491].

Section 7.2 of [RFC3693] details the requirements of a "Using Protocol". These requirements are restated, followed by a statement of compliance:

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

Compliant: This specification describes the use of HTTP over TLS for carrying the PIDF-LO from the LS to the Location Recipient. The sending and receiving parties are expected to comply with the instructions carried inside the object.

Though discouraged, using unsecured http: URIs is permitted. Using unsecured HTTP is likely to result in non-compliance with this requirement.

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

Compliant: This document specifies that authentication of the LS uses the established public key infrastructure used by HTTP over TLS [RFC2818]. Authentication of Location Recipients is either based on distribution of a secret (the location URI) using a conveyance protocol (for instance, [RFC6442]), allowances are made for later work to define alternative methods.

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

Not Compliant: The XML encoding specified in [RFC4119] is not suited to single packet transfers. Use of compressed content encoding [RFC2616] might allow this condition to be met.

Section 7.3 of [RFC3693] details the requirements of a "Rule based Location Data Transfer". These requirements are restated where they are applicable to this document:

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

Compliant: This document describes two alternative methods by which a Rule Maker is able to control access to location information. Rule Maker policy is enforced by the LS when

a location URI is dereferenced. However, this document does not describe how a location URI is created, or how a Rule Maker associates policy with a location URI. These are covered by other specifications.

- Req. 8. (LG Rules) Not Applicable: This relationship between LS and the source of its information (be that Location Generator (LG) or LIS) is out of scope for this document.
- Req. 9. "(Viewer Rules) A Viewer does not need to be aware of the full Rules defined by the Rule Maker (because a Viewer SHOULD NOT retransmit Location Information), and thus a Viewer SHOULD receive only the subset of Privacy Rules necessary for the Viewer to handle the LO in compliance with the full Privacy Rules (such as, instruction on the time period for which the LO can be retained)."

Compliant: The Rule Maker might define (via mechanisms outside the scope of this document) which policy rules are disclosed to other entities. For instance, if [RFC4745] is used to convey authorization policies from Rule Maker to LS, this is possible using the parameters specified in [I-D.ietf-geopriv-policy].

In order to comply with these rules, a Location Recipient MUST NOT redistribute a location URI without express permission. Depending on the access control model, the location URI might be secret (see Section 3.3 of [RFC5808]).

- Req. 10. (Full Rule language) Not Applicable: Note however that Geopriv has defined a rule language capable of expressing a wide range of privacy rules (see [RFC4745] and [I-D.ietf-geopriv-policy].
- Req. 11. (Limited Rule language) Not Applicable: This requirement applies to (and is addressed by) PIDF-LO [RFC4119].

Section 7.4 of [RFC3693] details the requirements of "Location Object Privacy and Security". These requirements are restated where they are applicable to this document:

- Req. 12. (Identity Protection) Compliant: Identity protection of the Target is provided as long as both of the following conditions are true:

- (a) the location URI is not associated with the identity of the Target in any context, and
- (b) the PIDF-LO does not contain information about the identity of the Target.

For instance, this requirement is complied with if the protocol that conveys the location URI does not link the identity of the Target to the location URI and the LS doesn't include meaningful identification information in the PIDF-LO document. Section 6 recommends that an unlinked pseudonym is used by the LS.

- Req. 13. (Credential Requirements) Compliant: The primary security mechanism specified in this document is Transport Layer Security. TLS offers the ability to use different types of credentials, including symmetric, asymmetric credentials or a combination of them.
- Req. 14. (Security Features) Compliant: Geopriv defines a few security requirements for the protection of Location Objects such as mutual end-point authentication, data object integrity, data object confidentiality and replay protection. The ability to use Transport Layer security fulfills most of these requirements. Authentication of Location Recipients in this document relies on proof of a shared secret - the location URI. This does not preclude the addition of more robust authentication procedures.
- Req. 15. (Minimal Crypto) Compliant: The mandatory to implement ciphersuite is provided in the TLS layer security specification.

Appendix B. Compliance to Location Reference Requirements

This section describes how HELD complies to the location reference requirements stipulated in [RFC5808]. Compliance of [RFC5985] to the Location Configuration Protocol is included.

Note that use of HELD as a location dereference protocol does not necessarily imply that HELD is the corresponding LCP. This document is still applicable to HTTP location URIs that are acquired by other means.

B.1. Requirements for a Location Configuration Protocol

- C1. "Location URI support: The location configuration protocol MUST support a location reference in URI form."

Compliant: HELD only provides location references in URI form.

- C2. "Location URI expiration: When a location URI has a limited validity interval, its lifetime MUST be indicated."

Compliant: HELD indicates the expiry time of location URIs using the "expires" attribute. [I-D.ietf-geopriv-policy-uri] provides a way to control expiration of a location URI.

- C3. "Location URI cancellation: The location configuration protocol MUST support the ability to request a cancellation of a specific location URI."

Compliant with Extension: [I-D.ietf-geopriv-policy-uri] describes how a location URI can be cancelled through the application of policy. Without extensions, HELD does not provide a method for cancelling location URIs.

- C4. "Location Information Masking: The location URI MUST ensure, by default, through randomization and uniqueness, that the location URI does not contain location information specific components."

Compliant: The HELD specification explicitly references this requirement in providing guidance on the format of the location URI.

- C5. "Target Identity Protection: The location URI MUST NOT contain information that identifies the Target (e.g., user or device)."

Compliant: The HELD specification provides specific guidance on the anonymity of the Target with regards to the generation of location URIs. Section 6 expands on this guidance.

- C6. "Reuse indicator: There SHOULD be a way to allow a Target to control whether a location URI can be resolved once only, or multiple times."

Not Compliant: Specific extensions to the protocol or authorization policy formats is needed to alter the default behavior, which allows unlimited resolution of the location URI.

- C7. "Selective disclosure: The location configuration protocol MUST provide a mechanism that allows the Rule Maker to control what information is being disclosed about the Target."

Compliant with Extension: Use of policy mechanisms and [I-D.ietf-geopriv-policy-uri] enable this capability. Note that this document recommends that only location information be provided.

- C8. "Location URI Not guessable: As a default, the location configuration protocol MUST return location URIs that are random and unique throughout the indicated lifetime. A location URI with 128-bits of randomness is RECOMMENDED."

Compliant: HELD specifies that location URIs conform to this requirement. The amount of randomness is not specifically identified since it depends on a number of factors that change over time, such as the number of valid location URIs, the validity period of those URIs and the rate that guesses can be made.

- C9. "Location URI Options: In the case of user-provided authorization policies, where anonymous or non-guessable location URIs are not warranted, the location configuration protocol MAY support a variety of optional location URI conventions, as requested by a Target to a location configuration server, (e.g., embedded location information within the location URI)."

Not Compliant: HELD does not support Device-specified location URI forms.

B.2. Requirements for a Location Dereference Protocol

- D1. "Location URI support: The location dereference protocol MUST support a location reference in URI form."

Compliant: HELD only provides location references in URI form.

- D2. "Authentication: The location dereference protocol MUST include mechanisms to authenticate both the client and the server."

Partially Compliant: TLS provides means for mutual authentication. This document only specifies the required mechanism for server authentication. Client authentication is not precluded.

- D3. "Dereferenced Location Form: The value returned by the dereference protocol MUST contain a well-formed PIDF-LO document."

Compliant: HELD requires that location objects are in the form of a PIDF-LO that complies with [RFC5491].

- D4. "Location URI Repeated Use: The location dereference protocol MUST support the ability for the same location URI to be resolved more than once, based on dereference server configuration."

Compliant: A Location Recipient may access and use a location URI as many times as desired until URI expiration results in the URI being invalidated. Authorization policies might include rules that modify this behavior.

- D5. "The location dereference protocol MUST support confidentiality protection of messages sent between the Location Recipient and the location server."

Compliant: This document strongly recommends the use of TLS for confidentiality and HELD mandates its implementation. Unsecured HTTP is permitted: the associated risks are described in Section 3.

Authors' Addresses

James Winterbottom
Commscope
Andrew Building (39)
Wollongong University Campus
Northfields Avenue
Wollongong, NSW 2522
AU

Phone: +61 242 212938
Email: james.winterbottom@commscope.com

Hannes Tschofenig
Nokia Siemens Networks
Linnoitustie 6
Espoo 02600
Finland

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

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

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

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

Phone: +1 650-353-1925
Email: martin.thomson@skype.net

GEOPRIV
Internet-Draft
Intended status: Standards Track
Expires: March 10, 2014

M. Thomson
Microsoft
J. Winterbottom
Unaffiliated
September 06, 2013

Using Device-provided Location-Related Measurements in Location
Configuration Protocols
draft-ietf-geopriv-held-measurements-09

Abstract

This document describes a protocol for a Device to provide location-related measurement data to a Location Information Server (LIS) within a request for location information. Location-related measurement information are observations concerning properties related to the position of a Device, which could be data about network attachment or about the physical environment. A LIS is able to use the location-related measurement data to improve the accuracy of the location estimate it provides to the Device. A basic set of location-related measurements are defined, including common modes of network attachment as well as assisted Global Navigation Satellite System (GNSS) parameters.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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

This Internet-Draft will expire on March 10, 2014.

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.	Conventions used in this document	5
3.	Location-Related Measurements in LCPs	5
4.	Location-Related Measurement Data Types	7
4.1.	Measurement Container	7
4.1.1.	Time of Measurement	8
4.1.2.	Expiry Time on Location-Related Measurement Data	8
4.2.	RMS Error and Number of Samples	8
4.2.1.	Time RMS Error	9
4.3.	Measurement Request	10
4.4.	Identifying Location Provenance	11
5.	Location-Related Measurement Data Types	13
5.1.	LLDP Measurements	14
5.2.	DHCP Relay Agent Information Measurements	15
5.3.	802.11 WLAN Measurements	15
5.3.1.	Wifi Measurement Requests	19
5.4.	Cellular Measurements	19
5.4.1.	Cellular Measurement Requests	22
5.5.	GNSS Measurements	22
5.5.1.	GNSS System and Signal	24
5.5.2.	Time	24
5.5.3.	Per-Satellite Measurement Data	24
5.5.4.	GNSS Measurement Requests	25
5.6.	DSL Measurements	25
5.6.1.	L2TP Measurements	26
5.6.2.	RADIUS Measurements	26
5.6.3.	Ethernet VLAN Tag Measurements	27
5.6.4.	ATM Virtual Circuit Measurements	28
6.	Privacy Considerations	28
6.1.	Measurement Data Privacy Model	28
6.2.	LIS Privacy Requirements	29
6.3.	Measurement Data and Location URIs	29
6.4.	Third-Party-Provided Measurement Data	30
7.	Security Considerations	30
7.1.	Threat Model	30
7.1.1.	Acquiring Location Information Without Authorization	31

7.1.2.	Extracting Network Topology Data	32
7.1.3.	Exposing Network Topology Data	32
7.1.4.	Lying By Proxy	32
7.1.5.	Measurement Replay	33
7.1.6.	Environment Spoofing	34
7.2.	Mitigation	35
7.2.1.	Measurement Validation	36
7.2.1.1.	Effectiveness	36
7.2.1.2.	Limitations (Unique Observer)	37
7.2.2.	Location Validation	38
7.2.2.1.	Effectiveness	38
7.2.2.2.	Limitations	38
7.2.3.	Supporting Observations	39
7.2.3.1.	Effectiveness	39
7.2.3.2.	Limitations	40
7.2.4.	Attribution	40
7.2.5.	Stateful Correlation of Location Requests	41
7.3.	An Unauthorized or Compromised LIS	42
8.	Measurement Schemas	42
8.1.	Measurement Container Schema	42
8.2.	Measurement Source Schema	44
8.3.	Base Type Schema	45
8.4.	LLDP Measurement Schema	48
8.5.	DHCP Measurement Schema	49
8.6.	WiFi Measurement Schema	50
8.7.	Cellular Measurement Schema	54
8.8.	GNSS Measurement Schema	56
8.9.	DSL Measurement Schema	58
9.	IANA Considerations	60
9.1.	IANA Registry for GNSS Types	60
9.2.	URN Sub-Namespace Registration for urn:ietf:params:xml:ns:pidf:geopriv10:lmsrc	61
9.3.	URN Sub-Namespace Registration for urn:ietf:params:xml:ns:geopriv:lm	62
9.4.	URN Sub-Namespace Registration for urn:ietf:params:xml:ns:geopriv:lm:basetypes	63
9.5.	URN Sub-Namespace Registration for urn:ietf:params:xml:ns:geopriv:lm:lldp	63
9.6.	URN Sub-Namespace Registration for urn:ietf:params:xml:ns:geopriv:lm:dhcp	64
9.7.	URN Sub-Namespace Registration for urn:ietf:params:xml:ns:geopriv:lm:wifi	65
9.8.	URN Sub-Namespace Registration for urn:ietf:params:xml:ns:geopriv:lm:cell	65
9.9.	URN Sub-Namespace Registration for urn:ietf:params:xml:ns:geopriv:lm:gnss	66
9.10.	URN Sub-Namespace Registration for urn:ietf:params:xml:ns:geopriv:lm:dsl	67

9.11. XML Schema Registration for Measurement Source Schema . .	67
9.12. XML Schema Registration for Measurement Container Schema	68
9.13. XML Schema Registration for Base Types Schema	68
9.14. XML Schema Registration for LLDP Schema	68
9.15. XML Schema Registration for DHCP Schema	68
9.16. XML Schema Registration for WiFi Schema	69
9.17. XML Schema Registration for Cellular Schema	69
9.18. XML Schema Registration for GNSS Schema	69
9.19. XML Schema Registration for DSL Schema	69
10. Acknowledgements	70
11. References	70
11.1. Normative References	70
11.2. Informative References	72
Authors' Addresses	73

1. Introduction

A Location Configuration Protocol (LCP) provides a means for a Device to request information about its physical location from an access network. A location information server (LIS) is the server that provides location information that is available due to the knowledge it has about the network and physical environment.

As a part of the access network, the LIS is able to acquire measurement results related to Device location from network elements. The LIS also has access to information about the network topology that can be used to turn measurement data into location information. This information can be further enhanced with information acquired from the Device itself.

A Device is able to make observations about its network attachment, or its physical environment. The location-related measurement data might be unavailable to the LIS; alternatively, the LIS might be able to acquire the data, but at a higher cost, in time or an other metric. Providing measurement data gives the LIS more options in determining location, which could improve the quality of the service provided by the LIS. Improvements in accuracy are one potential gain, but improved response times and lower error rates are possible.

This document describes a means for a Device to report location-related measurement data to the LIS. Examples based on the HELD [RFC5985] location configuration protocol are provided.

2. Conventions used in this document

The terms LIS and Device are used in this document in a manner consistent with the usage in [RFC5985].

This document also uses the following definitions:

Location Measurement: An observation about the physical properties of a particular Device's position in time and space. The result of a location measurement - "location-related measurement data", or simply "measurement data" given sufficient context - can be used to determine the location of a Device. Location-related measurement data does not directly identify a Device, though it could do indirectly. Measurement data can change with time if the location of the Device also changes.

Location-related measurement data does not necessarily contain location information directly, but it can be used in combination with contextual knowledge and/or algorithms to derive location information. Examples of location-related measurement data are: radio signal strength or timing measurements, Ethernet switch and port identifiers.

Location-related measurement data can be considered sighting information, based on the definition in [RFC3693].

Location Estimate: A location estimate is an approximation of where the Device is located. Location estimates are derived from location measurements. Location estimates are subject to uncertainty, which arise from errors in measurement results.

GNSS: Global Navigation Satellite System. A satellite-based system that provides positioning and time information. For example, the US Global Positioning System (GPS) or the European Galileo system.

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

3. Location-Related Measurements in LCPs

This document defines a standard container for the conveyance of location-related measurement parameters in location configuration protocols. This is an XML container that identifies parameters by type and allows the Device to provide the results of any measurement it is able to perform. A set of measurement schemas are also defined that can be carried in the generic container.

A simple example of measurement data conveyance is illustrated by the example message in Figure 1. This shows a HELD location request message with an Ethernet switch and port measurement taken using LLDP [IEEE.8021AB].

```
<locationRequest xmlns="urn:ietf:params:xml:ns:geopriv:held">
  <locationType exact="true">civic</locationType>
  <measurements xmlns="urn:ietf:params:xml:ns:geopriv:lm"
    time="2008-04-29T14:33:58">
    <lldp xmlns="urn:ietf:params:xml:ns:geopriv:lm:lldp">
      <chassis type="4">0a01003c</chassis>
      <port type="6">c2</port>
    </lldp>
  </measurements>
</locationRequest>
```

Figure 1: HELD Location Request with Measurement Data

This LIS can ignore measurement data that it does not support or understand. The measurements defined in this document follow this rule: extensions that could result in backward incompatibility MUST be added as new measurement definitions rather than extensions to existing types.

Multiple sets of measurement data, either of the same type or from different sources, can be included in the "measurements" element. See Section 4.1.1 for details on repetition of this element.

A LIS can choose to use or ignore location-related measurement data in determining location, as long as rules regarding use and retention (Section 6) are respected. The "method" parameter in the Presence Information Data Format - Location Object (PIDF-LO) [RFC4119] SHOULD be adjusted to reflect the method used. A correct "method" can assist location recipients in assessing the quality (both accuracy and integrity) of location information, though there could be reasons to withhold information about the source of data.

Measurement data is typically only used to serve the request that it is included in. There may be exceptions, particularly with respect to location URIs. Section 6 provides more information on usage rules.

Location-related measurement data need not be provided exclusively by Devices. A third party location requester (for example, see [RFC6155]) can request location information using measurement data, if the requester is able to acquire measurement data and authorized to distribute it. There are specific privacy considerations relating to the use of measurements by third parties, which are discussed in Section 6.4.

Location-related measurement data and its use presents a number of privacy and security challenges. These are described in more detail in Section 6 and Section 7.

4. Location-Related Measurement Data Types

A common container is defined for the expression of location measurement data, as well as a simple means of identifying specific types of measurement data for the purposes of requesting them.

The following example shows a measurement container with measurement time and expiration time included. A WiFi measurement is enclosed.

```
<lm:measurements xmlns:lm="urn:ietf:params:xml:ns:geopriv:lm"
  time="2008-04-29T14:33:58"
  expires="2008-04-29T17:33:58">
  <wifi xmlns="urn:ietf:params:xml:ns:geopriv:lm:wifi">
    <ap serving="true">
      <bssid>00-12-F0-A0-80-EF</bssid>
      <ssid>wlan-home</ssid>
    </ap>
  </wifi>
</lm:measurements>
```

Figure 2: Measurement Example

4.1. Measurement Container

The "measurements" element is used to encapsulate measurement data that is collected at a certain point in time. It contains time-based attributes that are common to all forms of measurement data, and permits the inclusion of arbitrary measurement data. The elements that are included within the "measurements" element are generically referred to as "measurement elements".

This container can be added to a request for location information in any protocol capable of carrying XML, such as a HELD location request [RFC5985].

4.1.1. Time of Measurement

The "time" attribute records the time that the measurement or observation was made. This time can be different to the time that the measurement information was reported. Time information can be used to populate a timestamp on the location result, or to determine if the measurement information is used.

The "time" attribute SHOULD be provided whenever possible. This allows a LIS to avoid selecting an arbitrary timestamp. Exceptions to this, where omitting time might make sense, include relatively static types of measurement (for instance, the DSL measurements in Section 5.6) or for legacy Devices that don't record time information (such as the Home Location Register/Home Subscriber Server for cellular).

The "time" attribute is attached to the root "measurement" element. Multiple measurements can often be given the same timestamp, even when the measurements were not actually taken at the same time (consider a set of measurements taken sequentially, where the difference in time between observations is not significant). Measurements cannot be grouped if they have different types, or there is a need for independent time values on each measurement. In these instances, multiple measurement sets are necessary.

4.1.2. Expiry Time on Location-Related Measurement Data

A Device is able to indicate an expiry time in the location measurement using the "expires" attribute. Nominally, this attribute indicates how long information is expected to be valid, but it can also indicate a time limit on the retention and use of the measurement data. A Device can use this attribute to request that the LIS not retain measurement data beyond the indicated time.

Note: Movement of the Device might result in the measurement data being invalidated before the expiry time.

A Device is advised to set the "expires" attribute to earlier of: the time that measurements are likely to be unusable, and the time that it desires to have measurements discarded by the LIS. A Device that does not desire measurement data to be retained can omit the "expires" attribute. Section 6 describes more specific rules regarding measurement data retention.

4.2. RMS Error and Number of Samples

Often a measurement is taken more than once. Reporting the average of a number of measurement results mitigates the effects of random errors that occur in the measurement process.

Reporting each measurement individually can be the most effective method of reporting multiple measurements. This is achieved by providing multiple measurement elements for different times.

The alternative is to aggregate multiple measurements and report a mean value across the set of measurements. Additional information about the distribution of the results can be useful in determining location uncertainty.

Two attributes are provided for use on some measurement values:

rmsError: The root-mean-squared (RMS) error of the set of measurement values used in calculating the result. RMS error is expressed in the same units as the measurement, unless otherwise stated. If an accurate value for RMS error is not known, this value can be used to indicate an upper bound or estimate for the RMS error.

samples: The number of samples that were taken in determining the measurement value. If omitted, this value can be assumed to be large enough that the RMS error is an indication of the standard deviation of the sample set.

For some measurement techniques, measurement error is largely dependent on the measurement technique employed. In these cases, measurement error is largely a product of the measurement technique and not the specific circumstances, so RMS error does not need to be actively measured. A fixed value MAY be provided for RMS error where appropriate.

The "rmsError" and "samples" elements are added as attributes of specific measurement data types.

4.2.1. Time RMS Error

Measurement of time can be significant in certain circumstances. The GNSS measurements included in this document are one such case where a small error in time can result in a large error in location. Factors such as clock drift and errors in time synchronization can result in small, but significant, time errors. Including an indication of the quality of time measurements can be helpful.

A "timeError" attribute MAY be added to the "measurement" element to indicate the RMS error in time. "timeError" indicates an upper bound on the time RMS error in seconds.

The "timeError" attribute does not apply where multiple samples of a measurement are taken over time. If multiple samples are taken, each SHOULD be included in a different "measurement" element.

4.3. Measurement Request

A measurement request is used by a protocol peer to describe a set of measurement data that it desires. A "measurementRequest" element is defined that can be included in a protocol exchange.

For instance, a LIS can use a measurement request in HELD responses. If the LIS is unable to provide location information, but it believes that a particular measurement type would enable it to provide a location, it can include a measurement request in an error response.

The "measurement" element of the measurement request identifies the type of measurement that is requested. The "type" attribute of this element indicates the type of measurement, as identified by an XML qualified name. An "samples" attribute MAY be used to indicate how many samples of the identified measurement are requested.

The "measurement" element can be repeated to request multiple (or alternative) measurement types.

Additional XML content might be defined for a particular measurement type that is used to further refine a request. These elements either constrain what is requested or specify non-mandatory components of the measurement data that are needed. These are defined along with the specific measurement type.

In the HELD protocol, the inclusion of a measurement request in an error response with a code of "locationUnknown" indicates that providing measurements would increase the likelihood of a subsequent request being successful.

The following example shows a HELD error response that indicates that WiFi measurement data would be useful if a later request were made. Additional elements indicate that received signal strength for an 802.11n access point is requested.

```
<error xmlns="urn:ietf:params:xml:ns:geopriv:held"
  code="locationUnknown">
  <message xml:lang="en">Insufficient measurement data</message>
  <measurementRequest
    xmlns="urn:ietf:params:xml:ns:geopriv:lm"
    xmlns:wifi="urn:ietf:params:xml:ns:geopriv:lm:wifi">
    <measurement type="wifi:wifi">
    <wifi:type>n</wifi:type>
    <wifi:parameter context="ap">wifi:rcpi</wifi:parameter>
    </measurement>
  </measurementRequest>
</error>
```

Figure 3: HELD Error Requesting Measurement Data

A measurement request that is included in other HELD messages has undefined semantics and can be safely ignored. Other specifications might define semantics for measurement requests under other conditions.

4.4. Identifying Location Provenance

An extension is made to the PIDF-LO [RFC4119] that allows a location recipient to identify the source (or sources) of location information and the measurement data that was used to determine that location information.

The "source" element is added to the "geopriv" element of the PIDF-LO. This element does not identify specific entities. Instead, it identifies the type of source.

The following types of measurement source are identified:

lis: Location information is based on measurement data that the LIS or sources that it trusts have acquired. This label MAY be used if measurement data provided by the Device has been completely validated by the LIS.

device: A LIS MUST include this value if the location information is based (in whole or part) on measurement data provided by the Device and if the measurement data isn't completely validated.

other: Location information is based on measurement data that a third party has provided. This might be an authorized third party that uses identity parameters [RFC6155] or any other entity. The LIS MUST include this, unless the third party is trusted by the LIS to provide measurement data.

No assertion is made about the veracity of the measurement data from sources other than the LIS. A combination of tags MAY be included to indicate that measurement data from multiple types of sources was used.

For example, the first tuple of the following PIDF-LO indicates that measurement data from a LIS and a device was combined to produce the result, the second tuple was produced by the LIS alone.

```
<presence xmlns="urn:ietf:params:xml:ns:pidf"
  xmlns:gp="urn:ietf:params:xml:ns:pidf:geopriv10"
  xmlns:gml="http://www.opengis.net/gml"
  xmlns:gs="http://www.opengis.net/pidf/1.0"
  xmlns:lmsrc="urn:ietf:params:xml:ns:pidf:geopriv10:lmsrc"
  entity="pres:lm@example.com">
  <tuple id="deviceLoc">
    <status>
    <gp:geopriv>
      <gp:location-info>
        <gs:Circle srsName="urn:ogc:def:crs:EPSG::4326">
          <gml:pos>7.34324 134.47162</gml:pos>
          <gs:radius uom="urn:ogc:def:uom:EPSG::9001">
            850.24
          </gs:radius>
        </gs:Circle>
      </gp:location-info>
      <gp:usage-rules/>
      <gp:method>OTDOA</gp:method>
      <lmsrc:source>lis device</lmsrc:source>
    </gp:geopriv>
    </status>
  </tuple>
  <tuple id="lisLoc">
    <status>
    <gp:geopriv>
      <gp:location-info>
        <gs:Circle srsName="urn:ogc:def:crs:EPSG::4326">
          <gml:pos>7.34379 134.46484</gml:pos>
          <gs:radius uom="urn:ogc:def:uom:EPSG::9001">
            9000
          </gs:radius>
        </gs:Circle>
      </gp:location-info>
      <gp:usage-rules/>
      <gp:method>Cell</gp:method>
      <lmsrc:source>lis</lmsrc:source>
    </gp:geopriv>
    </status>
  </tuple>
</presence>
```

PIDF-LO document with source labels

5. Location-Related Measurement Data Types

This document defines location-related measurement data types for a range of common network types.

All included measurement data definitions allow for arbitrary extension in the corresponding schema. New parameters that are applicable to location determination are added as new XML elements in a unique namespace, not by adding elements to an existing namespace.

5.1. LLDP Measurements

Link-Layer Discovery Protocol (LLDP) [IEEE.8021AB] messages are sent between adjacent nodes in an IEEE 802 network (e.g. wired Ethernet, WiFi, 802.16). These messages all contain identification information for the sending node, which can be used to determine location information. A Device that receives LLDP messages can report this information as a location-related measurement to the LIS, which is then able to use the measurement data in determining the location of the Device.

Note: The LLDP extensions defined in LLDP Media Endpoint Discovery (LLDP-MED) [ANSI-TIA-1057] provide the ability to acquire location information directly from an LLDP endpoint. Where this information is available, it might be unnecessary to use any other form of location configuration.

Values are provided as hexadecimal sequences. The Device MUST report the values directly as they were provided by the adjacent node. Attempting to adjust or translate the type of identifier is likely to cause the measurement data to be useless.

Where a Device has received LLDP messages from multiple adjacent nodes, it should provide information extracted from those messages by repeating the "lldp" element.

An example of an LLDP measurement is shown in Figure 4. This shows an adjacent node (chassis) that is identified by the IP address 192.0.2.45 (hexadecimal c000022d) and the port on that node is numbered using an agent circuit ID [RFC3046] of 162 (hexadecimal a2).

```
<measurements xmlns="urn:ietf:params:xml:ns:geopriv:lm"
  time="2008-04-29T14:33:58">
  <lldp xmlns="urn:ietf:params:xml:ns:geopriv:lm:lldp">
    <chassis type="4">c000022d</chassis>
    <port type="6">a2</port>
  </lldp>
</measurements>
```

Figure 4: LLDP Measurement Example

IEEE 802 Devices that are able to obtain information about adjacent network switches and their attachment to them by other means MAY use this data type to convey this information.

5.2. DHCP Relay Agent Information Measurements

The DHCP Relay Agent Information option [RFC3046] provides measurement data about the network attachment of a Device. This measurement data can be included in the "dhcp-rai" element.

The elements in the DHCP relay agent information options are opaque data types assigned by the DHCP relay agent. The three items MAY be omitted if unknown: circuit identifier ("circuit", circuit [RFC3046], Interface-Id [RFC3315]), remote identifier ("remote", Remote ID [RFC3046], or remote-id [RFC4649]) and subscriber identifier ("subscriber", subscriber-id [RFC3993], Subscriber-ID [RFC4580]). The DHCPv6 remote-id has an associated enterprise number [IANA.enterprise] as an XML attribute.

```
<measurements xmlns="urn:ietf:params:xml:ns:geopriv:lm"
  time="2008-04-29T14:33:58">
  <dhcp-rai xmlns="urn:ietf:params:xml:ns:geopriv:lm:dhcp">
    <giaddr>192.0.2.158</giaddr>
    <circuit>108b</circuit>
  </dhcp-rai>
</measurements>
```

Figure 5: DHCP Relay Agent Information Measurement Example

The "giaddr" is specified as a dotted quad IPv4 address or an RFC 4291 [RFC4291] IPv6 address, using the forms defined in [RFC3986]; IPv6 addresses SHOULD use the form described in [RFC5952]. The enterprise number is specified as a decimal integer. All other information is included verbatim from the DHCP request in hexadecimal format.

The "subscriber" element could be considered sensitive. This information MUST NOT be provided to a LIS that is not authorized to receive information about the access network. See Section 7.1.3 for more details.

5.3. 802.11 WLAN Measurements

In WiFi, or 802.11 [IEEE.80211], networks a Device might be able to provide information about the access point (AP) that it is attached to, or other WiFi points it is able to see. This is provided using the "wifi" element, as shown in Figure 6, which shows a single complete measurement for a single access point.

```
<measurements xmlns="urn:ietf:params:xml:ns:geopriv:lm"
  time="2011-04-29T14:33:58">
  <wifi xmlns="urn:ietf:params:xml:ns:geopriv:lm:wifi">
    <nicType>Intel(r)PRO/Wireless 2200BG</nicType>
    <ap serving="true">
      <bssid>AB-CD-EF-AB-CD-EF</bssid>
      <ssid>example</ssid>
      <channel>5</channel>
      <location>
        <gml:Point xmlns:gml="http://opengis.net/gml">
          <gml:pos>-34.4 150.8</gml:pos>
        </gml:Point>
      </location>
      <type>a</type>
      <band>5</band>
      <regclass country="AU">2</regclass>
      <antenna>2</antenna>
      <flightTime rmsError="4e-9" samples="1">2.56e-9</flightTime>
      <apSignal>
        <transmit>23</transmit>
        <gain>5</gain>
        <rcpi dBm="true" rmsError="12" samples="1">-59</rcpi>
        <rsni rmsError="15" samples="1">23</rsni>
      </apSignal>
      <deviceSignal>
        <transmit>10</transmit>
        <gain>9</gain>
        <rcpi dBm="true" rmsError="9.5" samples="1">-98.5</rcpi>
        <rsni rmsError="6" samples="1">7.5</rsni>
      </deviceSignal>
    </ap>
  </wifi>
</measurements>
```

Figure 6: 802.11 WLAN Measurement Example

A wifi element is made up of one or more access points, and a "nicType" element, which MAY be omitted. Each access point is described using the "ap" element, which is comprised of the following fields:

bssid: The basic service set identifier. In an Infrastructure BSS network, the bssid is the 48 bit MAC address of the access point.

The "verified" attribute of this element describes whether the device has verified the MAC address or it authenticated the access point or the network operating the access point (for example, a captive portal accessed through the access point has been

authenticated). This attribute defaults to a value of "false" when omitted.

ssid: The service set identifier (SSID) for the wireless network served by the access point.

The SSID is a 32-octet identifier that is commonly represented as a ASCII [ASCII] or UTF-8 [RFC3629] encoded string. To represent octets that cannot be directly included in an XML element, escaping is used. Sequences of octets that do not represent a valid UTF-8 encoding can be escaped using a backslash ('\') followed by two case-insensitive hexadecimal digits representing the value of a single octet.

The canonical or value-space form of an SSID is a sequence of up to 32 octets that is produced from the concatenation of UTF-8 encoded sequences of unescaped characters and octets derived from escaped components.

channel: The channel number (frequency) that the access point operates on.

location: The location of the access point, as reported by the access point. This element contains any valid location, using the rules for a "location-info" element, as described in [RFC5491].

type: The network type for the network access. This element includes the alphabetic suffix of the 802.11 specification that introduced the radio interface, or PHY; e.g. "a", "b", "g", or "n".

band: The frequency band for the radio, in gigahertz (GHz). 802.11 [IEEE.80211] specifies PHY layers that use 2.4, 3.7 and 5 gigahertz frequency bands.

regclass: The operating class (regulatory domain and class in older versions in 802.11), see Annex E of [IEEE.80211]. The "country" attribute optionally includes the applicable two character country identifier (dot11CountryString), which can be followed by an 'O', 'I' or 'X'. The element text content includes the value of the regulatory class: an 8-bit integer in decimal form.

antenna: The antenna identifier for the antenna that the access point is using to transmit the measured signals.

flightTime: Flight time is the difference between the time of departure (TOD) of signal from a transmitting station and time of arrival (TOA) of signal at a receiving station, as defined in

[IEEE.80211]. Measurement of this value requires that stations synchronize their clocks. This value can be measured by access point or Device; because the flight time is assumed to be the same in either direction - aside from measurement errors - only a single element is provided. This element permits the use of the "rmsError" and "samples" attributes. RMS error might be derived from the reported RMS error in TOD and TOA.

apSignal: Measurement information for the signal transmitted by the access point, as observed by the Device. Some of these values are derived from 802.11v [IEEE.80211] messages exchanged between Device and access point. The contents of this element include:

transmit: The transmit power reported by the access point, in dBm.

gain: The gain of the access point antenna reported by the access point, in dB.

rcpi: The received channel power indicator for the access point signal, as measured by the Device. This value SHOULD be in units of dBm (with RMS error in dB). If power is measured in a different fashion, the "dBm" attribute MUST be set to "false". Signal strength reporting on current hardware uses a range of different mechanisms; therefore, the value of the "nicType" element SHOULD be included if the units are not known to be in dBm and the value reported by the hardware should be included without modification. This element permits the use of the "rmsError" and "samples" attributes.

rsni: The received signal to noise indicator in dB. This element permits the use of the "rmsError" and "samples" attributes.

deviceSignal: Measurement information for the signal transmitted by the device, as reported by the access point. This element contains the same child elements as the "ap" element, with the access point and Device roles reversed.

The only mandatory element in this structure is "bssid".

The "nicType" element is used to specify the make and model of the wireless network interface in the Device. Different 802.11 chipsets report measurements in different ways, so knowing the network interface type aids the LIS in determining how to use the provided measurement data. The content of this field is unconstrained and no mechanisms are specified to ensure uniqueness. This field is unlikely to be useful, except under tightly controlled circumstances.

5.3.1. Wifi Measurement Requests

Two elements are defined for requesting WiFi measurements in a measurement request:

type: The "type" element identifies the desired type (or types that are requested).

parameter: The "parameter" element identifies measurements that are requested for each measured access point. An element is identified by its qualified name. The "context" parameter can be used to specify if an element is included as a child of the "ap" or "device" elements; omission indicates that it applies to both.

Multiple types or parameters can be requested by repeating either element.

5.4. Cellular Measurements

Cellular Devices are common throughout the world and base station identifiers can provide a good source of coarse location information. Cellular measurements can be provided to a LIS run by the cellular operator, or may be provided to an alternative LIS operator that has access to one of several global cell-id to location mapping databases.

A number of advanced location determination methods have been developed for cellular networks. For these methods a range of measurement parameters can be collected by the network, Device, or both in cooperation. This document includes a basic identifier for the wireless transmitter only; future efforts might define additional parameters that enable more accurate methods of location determination.

The cellular measurement set allows a Device to report to a LIS any LTE (Figure 7), UMTS (Figure 8), GSM (Figure 9) or CDMA (Figure 10) cells that it is able to observe. Cells are reported using their global identifiers. All 3GPP cells are identified by public land mobile network (PLMN), which is formed of mobile country code (MCC) and mobile network code (MNC); specific fields are added for each network type.

Formats for 3GPP cell identifiers are described in [TS.3GPP.23.003]. Bit-level formats for CDMA cell identifiers are described in [TIA-2000.5]; decimal representations are used.

MCC and MNC are provided as decimal digit sequences; a leading zero in an MCC or MNC is significant. All other values are decimal integers.

```
<measurements xmlns="urn:ietf:params:xml:ns:geopriv:lm"
  time="2008-04-29T14:33:58">
  <cellular xmlns="urn:ietf:params:xml:ns:geopriv:lm:cell">
    <servingCell>
      <mcc>465</mcc><mnc>20</mnc><eucid>80936424</eucid>
    </servingCell>
    <observedCell>
      <mcc>465</mcc><mnc>06</mnc><eucid>10736789</eucid>
    </observedCell>
  </cellular>
</measurements>
```

Long term evolution (LTE) cells are identified by a 28-bit cell identifier (eucid).

Figure 7: Example LTE Cellular Measurement

```
<measurements xmlns="urn:ietf:params:xml:ns:geopriv:lm"
  time="2008-04-29T14:33:58">
  <cellular xmlns="urn:ietf:params:xml:ns:geopriv:lm:cell">
    <servingCell>
      <mcc>465</mcc><mnc>20</mnc>
      <rnc>2000</rnc><cid>65000</cid>
    </servingCell>
    <observedCell>
      <mcc>465</mcc><mnc>06</mnc>
      <lac>16383</lac><cid>32767</cid>
    </observedCell>
  </cellular>
</measurements>
```

Universal mobile telephony service (UMTS) cells are identified by 12- or 16-bit radio network controller (rnc) id and a 16-bit cell id (cid).

Figure 8: Example UMTS Cellular Measurement

```
<measurements xmlns="urn:ietf:params:xml:ns:geopriv:lm"
  time="2008-04-29T14:33:58">
  <cellular xmlns="urn:ietf:params:xml:ns:geopriv:lm:cell">
    <servingCell>
      <mcc>465</mcc><mnc>06</mnc>
      <lac>16383</lac><cid>32767</cid>
    </servingCell>
```

```

    </cellular>
  </measurements>

```

Global System for Mobile communication (GSM) cells are identified by a 16-bit location area code (lac) and 16-bit cell id (cid).

Figure 9: Example GSM Cellular Measurement

```

<measurements xmlns="urn:ietf:params:xml:ns:geopriv:lm"
  time="2008-04-29T14:33:58">
  <cellular xmlns="urn:ietf:params:xml:ns:geopriv:lm:cell">
    <servingCell>
      <sid>15892</sid><nid>4723</nid><baseid>12</baseid>
    </servingCell>
    <observedCell>
      <sid>15892</sid><nid>4723</nid><baseid>13</baseid>
    </observedCell>
  </cellular>
</measurements>

```

Code division multiple access (CDMA) cells are not identified by PLMN, instead these use a 15-bit system id (sid), a 16-bit network id (nid) and a 16-bit base station id (baseid).

Figure 10: Example CDMA Cellular Measurement

In general, a cellular Device will be attached to the cellular network and so the notion of a serving cell exists. Cellular network also provide overlap between neighbouring sites, so a mobile Device can hear more than one cell. The measurement schema supports sending both the serving cell and any other cells that the mobile might be able to hear. In some cases, the Device could simply be listening to cell information without actually attaching to the network, mobiles without a SIM are an example of this. In this case the Device could report cells it can hear without identifying any particular cell as serving cell. An example of this is shown in Figure 11.

```

<measurements xmlns="urn:ietf:params:xml:ns:geopriv:lm"
  time="2008-04-29T14:33:58">
  <cellular xmlns="urn:ietf:params:xml:ns:geopriv:lm:cell">
    <observedCell>
      <mcc>465</mcc><mnc>20</mnc>
      <rnc>2000</rnc><cid>65000</cid>
    </observedCell>
    <observedCell>
      <mcc>465</mcc><mnc>06</mnc>
      <lac>16383</lac><cid>32767</cid>
    </observedCell>
  </cellular>
</measurements>

```

```
</cellular>
</measurements>
```

Figure 11: Example Observed Cellular Measurement

5.4.1. Cellular Measurement Requests

Two elements can be used in measurement requests for cellular measurements:

type: A label indicating the type of identifier to provide: one of "gsm", "umts", "lte", or "cdma".

network: The network portion of the cell identifier. For 3GPP networks, this is the combination of MCC and MNC; for CDMA, this is the network identifier.

Multiple identifier types or networks can be identified by repeating either element.

5.5. GNSS Measurements

A Global Navigation Satellite System (GNSS) uses orbiting satellites to transmit signals. A Device with a GNSS receiver is able to take measurements from the satellite signals. The results of these measurements can be used to determine time and the location of the Device.

Determining location and time in autonomous GNSS receivers follows three steps:

Signal acquisition: During the signal acquisition stage, the receiver searches for the repeating code that is sent by each GNSS satellite. Successful operation typically requires measurement data for a minimum of 5 satellites. At this stage, measurement data is available to the Device.

Navigation message decode: Once the signal has been acquired, the receiver then receives information about the configuration of the satellite constellation. This information is broadcast by each satellite and is modulated with the base signal at a low rate; for instance, GPS sends this information at about 50 bits per second.

Calculation: The measurement data is combined with the data on the satellite constellation to determine the location of the receiver and the current time.

A Device that uses a GNSS receiver is able to report measurements after the first stage of this process. A LIS can use the results of these measurements to determine a location. In the case where there are fewer results available than the optimal minimum, the LIS might be able to use other sources of measurement information and combine these with the available measurement data to determine a position.

Note: The use of different sets of GNSS `_assistance data_` can reduce the amount of time required for the signal acquisition stage and obviate the need for the receiver to extract data on the satellite constellation. Provision of assistance data is outside the scope of this document.

Figure 12 shows an example of GNSS measurement data. The measurement shown is for the GPS system and includes measurement data for three satellites only.

```
<measurements xmlns="urn:ietf:params:xml:ns:geopriv:lm"
  time="2008-04-29T14:33:58" timeError="2e-5">
  <gnss xmlns="urn:ietf:params:xml:ns:geopriv:lm:gnss"
    system="gps" signal="L1">
    <sat num="19">
      <doppler>499.9395</doppler>
      <codephase rmsError="1.6e-9">0.87595747</codephase>
      <cn0>45</cn0>
    </sat>
    <sat num="27">
      <doppler>378.2657</doppler>
      <codephase rmsError="1.6e-9">0.56639479</codephase>
      <cn0>52</cn0>
    </sat>
    <sat num="20">
      <doppler>-633.0309</doppler>
      <codephase rmsError="1.6e-9">0.57016835</codephase>
      <cn0>48</cn0>
    </sat>
  </gnss>
</measurements>
```

Figure 12: Example GNSS Measurement

Each "gnss" element represents a single set of GNSS measurement data, taken at a single point in time. Measurements taken at different times can be included in different "gnss" elements to enable iterative refinement of results.

GNSS measurement parameters are described in more detail in the following sections.

5.5.1. GNSS System and Signal

The GNSS measurement structure is designed to be generic and to apply to different GNSS types. Different signals within those systems are also accounted for and can be measured separately.

The GNSS type determines the time system that is used. An indication of the type of system and signal can ensure that the LIS is able to correctly use measurements.

Measurements for multiple GNSS types and signals can be included by repeating the "gnss" element.

This document creates an IANA registry for GNSS types. Two satellite systems are registered by this document: GPS [GPS.ICD] and Galileo [Galileo.ICD]. Details for the registry are included in Section 9.1.

5.5.2. Time

Each set of GNSS measurements is taken at a specific point in time. The "time" attribute is used to indicate the time that the measurement was acquired, if the receiver knows how the time system used by the GNSS relates to UTC time.

Alternative to (or in addition to) the measurement time, the "gnssTime" element MAY be included. The "gnssTime" element includes a relative time in milliseconds using the time system native to the satellite system. For the GPS satellite system, the "gnssTime" element includes the time of week in milliseconds. For the Galileo system, the "gnssTime" element includes the time of day in milliseconds.

The accuracy of the time measurement provided is critical in determining the accuracy of the location information derived from GNSS measurements. The receiver SHOULD indicate an estimated time error for any time that is provided. An RMS error can be included for the "gnssTime" element, with a value in milliseconds.

5.5.3. Per-Satellite Measurement Data

Multiple satellites are included in each set of GNSS measurements using the "sat" element. Each satellite is identified by a number in the "num" attribute. The satellite number is consistent with the identifier used in the given GNSS.

Both the GPS and Galileo systems use satellite numbers between 1 and 64.

The GNSS receiver measures the following parameters for each satellite:

doppler: The observed Doppler shift of the satellite signal, measured in meters per second. This is converted from a value in Hertz by the receiver to allow the measurement to be used without knowledge of the carrier frequency of the satellite system. This value permits the use of RMS error attributes, also measured in meters per second.

codephase: The observed code phase for the satellite signal, measured in milliseconds. This is converted from the system-specific value of chips or wavelengths into a system independent value. Larger values indicate larger distances from satellite to receiver. This value permits the use of RMS error attributes, also measured in milliseconds.

cn0: The signal to noise ratio for the satellite signal, measured in decibel-Hertz (dB-Hz). The expected range is between 20 and 50 dB-Hz.

mp: An estimation of the amount of error that multipath signals contribute in meters. This parameter MAY be omitted.

cq: An indication of the carrier quality. Two attributes are included: "continuous" can be either "true" or "false"; direct can be either "direct" or "inverted". This parameter MAY be omitted.

adr: The accumulated Doppler range, measured in meters. This parameter MAY be omitted and is not useful unless multiple sets of GNSS measurements are provided or differential positioning is being performed.

All values are converted from measures native to the satellite system to generic measures to ensure consistency of interpretation. Unless necessary, the schema does not constrain these values.

5.5.4. GNSS Measurement Requests

Measurement requests can include a "gnss" element, which includes the "system" and "signal" attributes. Multiple elements can be included to indicate a requests for GNSS measurements from multiple systems or signals.

5.6. DSL Measurements

Digital Subscriber Line (DSL) networks rely on a range of network technologies. DSL deployments regularly require cooperation between

multiple organizations. These fall into two broad categories: infrastructure providers and Internet service providers (ISPs). For the same end user, an infrastructure and Internet service can be provided by different entities. Infrastructure providers manage the bulk of the physical infrastructure including cabling. End users obtain their service from an ISP, which manages all aspects visible to the end user including IP address allocation and operation of a LIS. See [DSL.TR025] and [DSL.TR101] for further information on DSL network deployments and the parameters that are available.

Exchange of measurement information between these organizations is necessary for location information to be correctly generated. The ISP LIS needs to acquire location information from the infrastructure provider. However, since the infrastructure provider could have no knowledge of Device identifiers, it can only identify a stream of data that is sent to the ISP. This is resolved by passing measurement data relating to the Device to a LIS operated by the infrastructure provider.

5.6.1. L2TP Measurements

Layer 2 Tunneling Protocol (L2TP) [RFC2661] is a common means of linking the infrastructure provider and the ISP. The infrastructure provider LIS requires measurement data that identifies a single L2TP tunnel, from which it can generate location information. Figure 13 shows an example L2TP measurement.

```
<measurements xmlns="urn:ietf:params:xml:ns:geopriv:lm"
  time="2008-04-29T14:33:58">
  <dsl xmlns="urn:ietf:params:xml:ns:geopriv:lm:dsl">
    <l2tp>
      <src>192.0.2.10</src>
      <dest>192.0.2.61</dest>
      <session>528</session>
    </l2tp>
  </dsl>
</measurements>
```

Figure 13: Example DSL L2TP Measurement

5.6.2. RADIUS Measurements

When authenticating network access, the infrastructure provider might employ a RADIUS [RFC2865] proxy at the DSL Access Module (DSLAM) or Access Node (AN). These messages provide the ISP RADIUS server with an identifier for the DSLAM or AN, plus the slot and port that the Device is attached to. These data can be provided as a measurement, which allows the infrastructure provider LIS to generate location information.

The format of the AN, slot and port identifiers are not defined in the RADIUS protocol. Slot and port together identify a circuit on the AN, analogous to the circuit identifier in [RFC3046]. These items are provided directly, as they were in the RADIUS message. An example is shown in Figure 14.

```
<measurements xmlns="urn:ietf:params:xml:ns:geopriv:lm"
  time="2008-04-29T14:33:58">
  <dsl xmlns="urn:ietf:params:xml:ns:geopriv:lm:dsl">
    <an>AN-7692</an>
    <slot>3</slot>
    <port>06</port>
  </dsl>
</measurements>
```

Figure 14: Example DSL RADIUS Measurement

5.6.3. Ethernet VLAN Tag Measurements

For Ethernet-based DSL access networks, the DSL Access Module (DSLAM) or Access Node (AN) provide two VLAN tags on packets. A C-TAG is used to identify the incoming residential circuit, while the S-TAG is used to identify the DSLAM or AN. The C-TAG and S-TAG together can be used to identify a single point of network attachment. An example is shown in Figure 15.

```
<measurements xmlns="urn:ietf:params:xml:ns:geopriv:lm"
  time="2008-04-29T14:33:58">
  <dsl xmlns="urn:ietf:params:xml:ns:geopriv:lm:dsl">
    <stag>613</stag>
    <ctag>1097</ctag>
  </dsl>
</measurements>
```

Figure 15: Example DSL VLAN Tag Measurement

Alternatively, the C-TAG can be replaced by data on the slot and port that the Device is attached to. This information might be included in RADIUS requests that are proxied from the infrastructure provider to the ISP RADIUS server.

5.6.4. ATM Virtual Circuit Measurements

An ATM virtual circuit can be employed between the ISP and infrastructure provider. Providing the virtual port ID (VPI) and virtual circuit ID (VCI) for the virtual circuit gives the infrastructure provider LIS the ability to identify a single data stream. A sample measurement is shown in Figure 16.

```
<measurements xmlns="urn:ietf:params:xml:ns:geopriv:lm"
  time="2008-04-29T14:33:58">
  <dsl xmlns="urn:ietf:params:xml:ns:geopriv:lm:dsl">
    <vpi>55</vpi>
    <vci>6323</vci>
  </dsl>
</measurements>
```

Figure 16: Example DSL ATM Measurement

6. Privacy Considerations

Location-related measurement data can be as privacy sensitive as location information [RFC6280].

Measurement data is effectively equivalent to location information if the contextual knowledge necessary to generate one from the other is readily accessible. Even where contextual knowledge is difficult to acquire, there can be no assurance that an authorized recipient of the contextual knowledge is also authorized to receive location information.

In order to protect the privacy of the subject of location-related measurement data, measurement data **MUST** be protected with the same degree of protection as location information. The confidentiality and authentication provided by TLS **MUST** be used in order to convey measurement data over HELD [RFC5985]. Other protocols **MUST** provide comparable guarantees.

6.1. Measurement Data Privacy Model

It is not necessary to distribute measurement data in the same fashion as location information. Measurement data is less useful to location recipients than location information. A simple distribution model is described in this document.

In this simple model, the Device is the only entity that is able to distribute measurement data. To use an analogy from the GEOPRIV architecture, the Device - as the Location Generator, or the Measurement Data Generator - is the sole entity that can act for the role of both Rule Maker and Location Server.

A Device that provides location-related measurement data, MUST only do so as explicitly authorized by a Rule Maker. This depends on having an interface that allows Rule Makers (for instance, users or administrators) to control where and how measurement data is provided.

No entity is permitted to redistribute measurement data. The Device directs other entities in how measurement data is used and retained.

The GEOPRIV model [RFC6280] protects the location of a Target using direction provided by a Rule Maker. For the purposes of measurement data distribution, this model relies on the assumptions made in Section 3 of HELD [RFC5985]. These assumptions effectively declare the Device to be a proxy for both Target and Rule Maker.

6.2. LIS Privacy Requirements

A LIS MUST NOT reveal location-related measurement data to any other entity. A LIS MUST NOT reveal location information based on measurement data to any other entity unless directed to do so by the Device.

By adding measurement data to a request for location information, the Device implicitly grants permission for the LIS to generate the requested location information using the measurement data. Permission to use this data for any other purpose is not implied.

As long as measurement data is only used in serving the request that contains it, rules regarding data retention are not necessary. A LIS MUST discard location-related measurement data after servicing a request, unless the Device grants permission to use that information for other purposes.

6.3. Measurement Data and Location URIs

A LIS MAY use measurement data provided by the Device to serve requests to location URIs, if the Device permits it. A Device permits this by including measurement data in a request that explicitly requests a location URI. By requesting a location URI, the Device grants permission for the LIS to use the measurement data in serving requests to that location URI. The LIS cannot provide location recipients with measurement data, as defined in Section 6.1.

Note: In HELD, the "any" type is not an explicit request for a location URI, though a location URI might be provided.

The usefulness of measurement data that is provided in this fashion is limited. The measurement data is only valid at the time that it was acquired by the Device. At the time that a request is made to a location URI, the Device might have moved, rendering the measurement data incorrect.

A Device is able to explicitly limit the time that a LIS retains measurement data by adding an expiry time to the measurement data. A LIS MUST NOT retain location-related measurement data in memory, storage or logs beyond the time indicated in the "expires" attribute (Section 4.1.2). A LIS MUST NOT retain measurement data if the "expires" attribute is absent.

6.4. Third-Party-Provided Measurement Data

An authorized third-party request for the location of a Device (see [RFC6155]) can include location-related measurement data. This is possible where the third-party is able to make observations about the Device.

A third-party that provides measurement data MUST be authorized to provide the specific measurement for the identified device. A third-party MUST either be trusted by the LIS for the purposes of providing measurement data of the provided type, or the measurement data MUST be validated (see Section 7.2.1) before being used.

How a third-party authenticates its identity or gains authorization to use measurement data is not covered by this document.

7. Security Considerations

Use of location-related measurement data has privacy considerations that are discussed in Section 6.

7.1. Threat Model

The threat model for location-related measurement data concentrates on the Device providing falsified, stolen or incorrect measurement data.

A Device that provides location-related measurement data might use data to:

- o acquire the location of another Device, without authorization;

- o extract information about network topology; or
- o coerce the LIS into providing falsified location information based on the measurement data.

Location-related measurement data describes the physical environment or network attachment of a Device. A third party adversary in the proximity of the Device might be able to alter the physical environment such that the Device provides measurement data that is controlled by the third party. This might be used to indirectly control the location information that is derived from measurement data.

7.1.1. Acquiring Location Information Without Authorization

Requiring authorization for location requests is an important part of privacy protections of a location protocol. A location configuration protocol usually operates under a restricted policy that allows a requester to obtain their own location. HELD identity extensions [RFC6155] allows other entities to be authorized, conditional on a Rule Maker providing sufficient authorization.

The intent of these protections is to ensure that a location recipient is authorized to acquire location information. Location-related measurement data could be used by an attacker to circumvent such authorization checks if the association between measurement data and Target Device is not validated by a LIS.

A LIS can be coerced into providing location information for a Device that a location recipient is not authorized to receive. A request identifies one Device (implicitly or explicitly), but measurement data is provided for another Device. If the LIS does not check that the measurement data is for the identified Device, it could incorrectly authorize the request.

By using unverified measurement data to generate a response, the LIS provides information about a Device without appropriate authorization.

The feasibility of this attack depends on the availability of information that links a Device with measurement data. In some cases, measurement data that is correlated with a target is readily available. For instance, LLDP measurements (Section 5.1) are broadcast to all nodes on the same network segment. An attacker on that network segment can easily gain measurement data that relates a Device with measurements.

For some types of measurement data, it's necessary for an attacker to know the location of the target in order to determine what measurements to use. This attack is meaningless for types of measurement data that require that the attacker first know the location of the target before measurement data can be acquired or fabricated. GNSS measurements (Section 5.5) share this trait with many wireless location determination methods.

7.1.2. Extracting Network Topology Data

Allowing requests with measurements might be used to collect information about network topology.

Network topology can be considered sensitive information by a network operator for commercial or security reasons. While it is impossible to completely prevent a Device from acquiring some knowledge of network topology if a location service is provided, a network operator might desire to limit how much of this information is made available.

Mapping a network topology does not require that an attacker be able to associate measurement data with a particular Device. If a requester is able to try a number of measurements, it is possible to acquire information about network topology.

It is not even necessary that the measurements are valid; random guesses are sufficient, provided that there is no penalty or cost associated with attempting to use the measurements.

7.1.3. Exposing Network Topology Data

A Device could reveal information about a network to entities outside of that network if it provides location measurement data to a LIS that is outside of that network. With the exception of GNSS measurements, the measurements in this document provide information about an access network that could reveal topology information to an unauthorized recipient.

A Device **MUST NOT** provide information about network topology without a clear signal that the recipient is authorized. A LIS that is discovered using DHCP as described in LIS discovery [RFC5986] can be considered to be authorized to receive information about the access network.

7.1.4. Lying By Proxy

Location information is a function of its inputs, which includes measurement data. Thus, falsified measurement data can be used to alter the location information that is provided by a LIS.

Some types of measurement data are relatively easy to falsify in a way that causes the resulting location information to be selected with little or no error. For instance, GNSS measurements are easy to use for this purpose because all the contextual information necessary to calculate a position using measurements is broadcast by the satellites [HARPER].

An attacker that falsifies measurement data gains little if they are the only recipients of the result. The attacker knows that the location information is bad. The attacker only gains if the information can somehow be attributed to the LIS by another location recipient. By coercing the LIS into providing falsified location information, any credibility that the LIS might have - that the attacker does not - is gained by the attacker.

A third-party that is reliant on the integrity of the location information might base an evaluation of the credibility of the information on the source of the information. If that third party is able to attribute location information to the LIS, then an attacker might gain.

Location information that is provided to the Device without any means to identify the LIS as its source is not subject to this attack. The Device is identified as the source of the data when it distributes the location information to location recipients.

Location information is attributed to the LIS either through the use of digital signatures or by having the location recipient directly interact with the LIS. A LIS that digitally signs location information becomes identifiable as the source of the data. Similarly, the LIS is identified as a source of data if a location recipient acquires information directly from a LIS using a location URI.

7.1.5. Measurement Replay

The value of some measured properties do not change over time for a single location. For properties of a network, time-invariance is often directly as a result of the practicalities of operating the network. Limiting the changes to a network ensures greater consistency of service. A largely static network also greatly simplifies the data management tasks involved with providing a location service. However, time invariant properties allow for simple replay attacks, where an attacker acquires measurements that can later be used without being detected as being invalid.

Measurement data is frequently an observation of an time-invariant property of the environment at the subject location. For measurements of this nature, nothing in the measurement itself is sufficient proof that the Device is present at the resulting location. Measurement data might have been previously acquired and reused.

For instance, the identity of a radio transmitter, if broadcast by that transmitter, can be collected and stored. An attacker that wishes it known that they exist at a particular location, can claim to observe this transmitter at any time. Nothing inherent in the claim reveals it to be false.

7.1.6. Environment Spoofing

Some types of measurement data can be altered or influenced by a third party so that a Device unwittingly provides falsified data. If it is possible for a third party to alter the measured phenomenon, then any location information that is derived from this data can be indirectly influenced.

Altering the environment in this fashion might not require involvement with either Device or LIS. Measurement that is passive - where the Device observes a signal or other phenomenon without direct interaction - are most susceptible to alteration by third parties.

Measurement of radio signal characteristics is especially vulnerable since an adversary need only be in the general vicinity of the Device and be able to transmit a signal. For instance, a GNSS spoofer is able to produce fake signals that claim to be transmitted by any satellite or set of satellites (see [GPS.SPOOF]).

Measurements that require direct interaction increases the complexity of the attack. For measurements relating to the communication medium, a third party cannot avoid direct interaction, they need only be on the communications path (that is, man in the middle).

Even if the entity that is interacted with is authenticated, this does not provide any assurance about the integrity of measurement data. For instance, the Device might authenticate the identity of a radio transmitter through the use of cryptographic means and obtain signal strength measurements for that transmitter. Radio signal strength is trivial for an attacker to increase simply by receiving and amplifying the raw signal; it is not necessary for the attacker to be able to understand the signal content.

Note: This particular "attack" is more often completely legitimate. Radio repeaters are commonplace mechanism used to increase radio coverage.

Attacks that rely on altering the observed environment of a Device require countermeasures that affect the measurement process. For radio signals, countermeasures could include the use of authenticated signals, or altered receiver design. In general, countermeasures are highly specific to the individual measurement process. An exhaustive discussion of these issues is left to the relevant literature for each measurement technology.

A Device that provides measurement data is assumed to be responsible for applying appropriate countermeasures against this type of attack.

Where a Device is the sole recipient of location information derived from measurement data, a LIS might choose to provide location information without any validation. The responsibility for ensuring the veracity of the measurement data lies with the Device.

Measurement data that is susceptible to this sort of influence SHOULD be treated as though it were produced by an untrusted Device for those cases where a location recipient might attribute the location information to the LIS. GNSS measurements and radio signal strength measurements can be affected relatively cheaply, though almost all other measurement types can be affected with varying costs to an attacker, with the largest cost often being a requirement for physical access. To the extent that it is feasible, measurement data SHOULD be subjected to the same validation as for other types of attacks that rely on measurement falsification.

Note: Altered measurement data might be provided by a Device that has no knowledge of the alteration. Thus, an otherwise trusted Device might still be an unreliable source of measurement data.

7.2. Mitigation

The following measures can be applied to limit or prevent attacks. The effectiveness of each depends on the type of measurement data and how that measurement data is acquired.

Two general approaches are identified for dealing with untrusted measurement data:

1. Require independent validation of measurement data or the location information that is produced.
2. Identify the types of sources that provided the measurement data that location information was derived from.

This section goes into more detail on the different forms of validation in Section 7.2.1, Section 7.2.2, and Section 7.2.3. The impact of attributing location information to sources is discussed in more detail in Section 7.2.4.

Any costs in validation are balanced against the degree of integrity desired from the resulting location information.

7.2.1. Measurement Validation

Detecting that measurement data has been falsified is difficult in the absence of integrity mechanisms.

Independent confirmation of the veracity of measurement data ensures that the measurement is accurate and that it applies to the correct Device. When it's possible to gather the same measurement data from a trusted and independent source without undue expense, the LIS can use the trusted data in place of what the untrusted Device has sent. In cases where that is impractical, the untrusted data can provide hints that allow corroboration of the data (see Section 7.2.1.1).

Measurement information might contain no inherent indication that it is falsified. On the contrary, it can be difficult to obtain information that would provide any degree of assurance that the measurement device is physically at any particular location. Measurements that are difficult to verify require other forms of assurance before they can be used.

7.2.1.1. Effectiveness

Measurement validation **MUST** be used if measurement data for a particular Device can be easily acquired by unauthorized location recipients, as described in Section 7.1.1. This prevents unauthorized access to location information using measurement data.

Validation of measurement data can be significantly more effective than independent acquisition of the same. For instance, a Device in a large Ethernet network could provide a measurement indicating its point of attachment using LLDP measurements. For a LIS, acquiring the same measurement data might require a request to all switches in that network. With the measurement data, validation can target the identified switch with a specific query.

Validation is effective in identifying falsified measurement data (Section 7.1.4), including attacks involving replay of measurement data (Section 7.1.5). Validation also limits the amount of network topology information (Section 7.1.2) made available to Devices to that portion of the network topology that they are directly attached.

Measurement validation has no effect if the underlying effect is being spoofed (Section 7.1.6).

7.2.1.2. Limitations (Unique Observer)

A Device is often in a unique position to make a measurement. It alone occupies the point in space-time that the location determination process seeks to determine. The Device becomes a unique observer for a particular property.

The ability of the Device to become a unique observer makes the Device invaluable to the location determination process. As a unique observer, it also makes the claims of a Device difficult to validate and easily to spoof.

As long as no other entity is capable of making the same measurements, there is also no other entity that can independently check that the measurements are correct and applicable to the Device. A LIS might be unable to validate all or part of the measurement data it receives from a unique observer. For instance, a signal strength measurement of the signal from a radio tower cannot be validated directly.

Some portion of the measurement data might still be independently verified, even if all information cannot. In the previous example, the radio tower might be able to provide verification that the Device is present if it is able to observe a radio signal sent by the Device.

If measurement data can only be partially validated, the extent to which it can be validated determines the effectiveness of validation against these attacks.

The advantage of having the Device as a unique observer is that it makes it difficult for an attacker to acquire measurements without the assistance of the Device. Attempts to use measurements to gain unauthorized access to measurement data (Section 7.1.1) are largely ineffectual against a unique observer.

7.2.2. Location Validation

Location information that is derived from location-related measurement data can also be verified against trusted location information. Rather than validating inputs to the location determination process, suspect locations are identified at the output of the process.

Trusted location information is acquired using sources of measurement data that are trusted. Untrusted location information is acquired using measurement data provided from untrusted sources, which might include the Device. These two locations are compared. If the untrusted location agrees with the trusted location, the untrusted location information is used.

Algorithms for the comparison of location information are not included in this document. However, a simple comparison for agreement might require that the untrusted location be entirely contained within the uncertainty region of the trusted location.

There is little point in using a less accurate, less trusted location. Untrusted location information that has worse accuracy than trusted information can be immediately discarded. There are multiple factors that affect accuracy, uncertainty and currency being the most important. How location information is compared for accuracy is not defined in this document.

7.2.2.1. Effectiveness

Location validation limits the extent to which falsified - or erroneous - measurement data can cause an incorrect location to be reported.

Location validation can be more efficient than validation of inputs, particularly for a unique observer (Section 7.2.1.2).

Validating location ensures that the Device is at or near the resulting location. Location validation can be used to limit or prevent all of the attacks identified in this document.

7.2.2.2. Limitations

The trusted location that is used for validation is always less accurate than the location that is being checked. The amount by which the untrusted location is more accurate, is the same amount that an attacker can exploit.

For example, a trusted location might indicate a five kilometer radius uncertainty region. An untrusted location that describes a 100 meter uncertainty within the larger region might be accepted as more accurate. An attacker might still falsify measurement data to select any location within the larger uncertainty region. While the 100 meter uncertainty that is reported seems more accurate, a falsified location could be anywhere in the five kilometer region.

Where measurement data might have been falsified, the actual uncertainty is effectively much higher. Local policy might allow differing degrees of trust to location information derived from untrusted measurement data. This might be a boolean operation with only two possible outcomes: untrusted location information might be used entirely or not at all. Alternatively, untrusted location could be combined with trusted location information using different weightings, based on a value set in local policy.

7.2.3. Supporting Observations

Replay attacks using previously acquired measurement data are particularly hard to detect without independent validation. Rather than validate the measurement data directly, supplementary data might be used to validate measurements or the location information derived from those measurements.

These supporting observations could be used to convey information that provides additional assurance that the Device was acquired at a specific time and place. In effect, the Device is requested to provide proof of its presence at the resulting location.

For instance, a Device that measures attributes of a radio signal could also be asked to provide a sample of the measured radio signal. If the LIS is able to observe the same signal, the two observations could be compared. Providing that the signal cannot be predicted in advance by the Device, this could be used to support the claim that the Device is able to receive the signal. Thus, the Device is likely to be within the range that the signal is transmitted. A LIS could use this to attribute a higher level of trust in the associated measurement data or resulting location.

7.2.3.1. Effectiveness

The use of supporting observations is limited by the ability of the LIS to acquire and validate these observations. The advantage of selecting observations independent of measurement data is that observations can be selected based on how readily available the data is for both LIS and Device. The amount and quality of the data can be selected based on the degree of assurance that is desired.

Use of supporting observations is similar to both measurement validation and location validation. All three methods rely on independent validation of one or more properties. Applicability of each method is similar.

Use of supporting observations can be used to limit or prevent all of the attacks identified in this document.

7.2.3.2. Limitations

The effectiveness of the validation method depends on the quality of the supporting observation: how hard it is to obtain at a different time or place, how difficult it is to guess, and what other costs might be involved in acquiring this data.

In the example of an observed radio signal, requesting a sample of the signal only provides an assurance that the Device is able to receive the signal transmitted by the measured radio transmitter. This only provides some assurance that the Device is within range of the transmitter.

As with location validation, a Device might still be able to provide falsified measurements that could alter the value of the location information as long as the result is within this region.

Requesting additional supporting observations can reduce the size of the region over which location information can be altered by an attacker, or increase trust in the result, but each additional measurement imposes an acquisition cost. Supporting observations contribute little or nothing toward the primary goal of determining the location of the Device.

7.2.4. Attribution

Lying by proxy (Section 7.1.4) relies on the location recipient being able to attribute location information to a LIS. The effectiveness of this attack is negated if location information is explicitly attributed to a particular source.

This requires an extension to the location object that explicitly identifies the source (or sources) of each item of location information.

Rather than relying on a process that seeks to ensure that location information is accurate, this approach instead provides a location recipient with the information necessary to reach their own conclusion about the trustworthiness of the location information.

Including an authenticated identity for all sources of measurement data presents a number of technical and operational challenges. It is possible that the LIS has a transient relationship with a Device. A Device is not expected to share authentication information with a LIS. There is no assurance that Device identification is usable by a potential location recipient. Privacy concerns might also prevent the sharing identification information, even if it were available and usable.

Identifying the type of measurement source allows a location recipient to make a decision about the trustworthiness of location information without depending on having authenticated identity information for each source. An element for this purpose is defined in Section 4.4.

When including location information that is based on measurement data from sources that might be untrusted, a LIS SHOULD include alternative location information that is derived from trusted sources of measurement data. Each item of location information can then be labelled with the source of that data.

A location recipient that is able to identify a specific source of measurement data (whether it be LIS or Device) can use this information to attribute location information to either or both entity. The location recipient is then better able to make decisions about trustworthiness based on the source of the data.

A location recipient that does not understand the "source" element is unable to make this distinction. When constructing a PIDF-LO document, trusted location information MUST be placed in the PIDF-LO so that it is given higher priority to any untrusted location information according to Rule #8 of [RFC5491].

Attribution of information does nothing to address attacks that alter the observed parameters that are used in location determination (Section 7.1.6).

7.2.5. Stateful Correlation of Location Requests

Stateful examination of requests can be used to prevent a Device from attempting to map network topology using requests for location information (Section 7.1.2).

Simply limiting the rate of requests from a single Device reduces the amount of data that a Device can acquire about network topology. A LIS could also make observations about the movements of a Device. A Device that is attempting to gather topology information is likely to be assigned a location that changes significantly between subsequent requests, possibly violating physical laws (or lower limits that might still be unlikely) with respect to speed and acceleration.

7.3. An Unauthorized or Compromised LIS

A compromised LIS, or a compromise in LIS discovery [RFC5986] could lead to an unauthorized entity obtaining measurement data. This information could then be used or redistributed. A Device MUST ensure that it authenticate a LIS, as described in Section 9 of [RFC5985].

An entity that is able to acquire measurement data can, in addition to using those measurements to learn the location of a Device, also use that information for other purposes. This information can be used to provide insight into network topology (Section 7.1.2).

Measurement data might also be exploited in other ways. For example, revealing the type of 802.11 transceiver that a Device uses could allow an attacker to use specific vulnerabilities to attack a Device. Similarly, revealing information about network elements could enable targeted attacks on that infrastructure.

8. Measurement Schemas

The schema are broken up into their respective functions. There is a base container schema into which all measurements are placed, plus definitions for a measurement request (Section 8.1). A PIDF-LO extension is defined in a separate schema (Section 8.2). There is a basic types schema, that contains various base type definitions for things such as the "rmsError" and "samples" attributes IPv4, IPv6 and MAC addresses (Section 8.3). Then each of the specific measurement types is defined in its own schema.

8.1. Measurement Container Schema

```
<?xml version="1.0"?>
<xs:schema
  xmlns:lm="urn:ietf:params:xml:ns:geopriv:lm"
  xmlns:bt="urn:ietf:params:xml:ns:geopriv:lm:basetypes"
```

```
xmlns:xs="http://www.w3.org/2001/XMLSchema"
targetNamespace="urn:ietf:params:xml:ns:geopriv:lm"
elementFormDefault="qualified"
attributeFormDefault="unqualified">

<xs:annotation>
  <xs:appinfo
    source="urn:ietf:params:xml:schema:geopriv:lm"
  </xs:appinfo>
  <xs:documentation source="http://www.ietf.org/rfc/rfcXXXX.txt">
    <!-- [[NOTE TO RFC-EDITOR: Please replace above URL with URL of
    published RFC and remove this note.]] -->
    This schema defines a framework for location measurements.
  </xs:documentation>
</xs:annotation>

<xs:import namespace="urn:ietf:params:xml:ns:geopriv:lm:basetypes"/>

<xs:element name="measurements">
  <xs:complexType>
    <xs:complexContent>
      <xs:restriction base="xs:anyType">
        <xs:sequence>
          <xs:any namespace="##other" processContents="lax"
            minOccurs="0" maxOccurs="unbounded"/>
        </xs:sequence>
        <xs:attribute name="time" type="xs:dateTime"/>
        <xs:attribute name="timeError" type="bt:positiveDouble"/>
        <xs:attribute name="expires" type="xs:dateTime"/>
        <xs:anyAttribute namespace="##any" processContents="lax"/>
      </xs:restriction>
    </xs:complexContent>
  </xs:complexType>
</xs:element>

<xs:element name="measurementRequest"
  type="lm:measurementRequestType"/>
<xs:complexType name="measurementRequestType">
  <xs:complexContent>
    <xs:restriction base="xs:anyType">
      <xs:sequence>
        <xs:element ref="lm:measurement"
          minOccurs="0" maxOccurs="unbounded"/>
        <xs:any namespace="##other" processContents="lax"
          minOccurs="0" maxOccurs="unbounded"/>
      </xs:sequence>
    </xs:restriction>
  </xs:complexContent>
</xs:complexType>
```

```

</xs:complexType>

<xs:element name="measurement" type="lm:measurementType"/>
<xs:complexType name="measurementType">
  <xs:complexContent>
    <xs:restriction base="xs:anyType">
      <xs:sequence>
        <xs:any namespace="##other" processContents="lax"
          minOccurs="0" maxOccurs="unbounded"/>
      </xs:sequence>
      <xs:attribute name="type" type="xs:QName" use="required"/>
      <xs:attribute name="samples" type="xs:positiveInteger"/>
    </xs:restriction>
  </xs:complexContent>
</xs:complexType>

<!-- PIDF-LO extension for source -->
<xs:element name="source" type="lm:sourceType"/>
<xs:simpleType name="sourceType">
  <xs:list>
    <xs:simpleType>
      <xs:restriction base="xs:token">
        <xs:enumeration value="lis"/>
        <xs:enumeration value="device"/>
        <xs:enumeration value="other"/>
      </xs:restriction>
    </xs:simpleType>
  </xs:list>
</xs:simpleType>
</xs:schema>

```

Measurement Container Schema

8.2. Measurement Source Schema

```

<?xml version="1.0"?>
<xs:schema
  xmlns:lmsrc="urn:ietf:params:xml:ns:pidf:geopriv10:lmsrc"
  xmlns:xs="http://www.w3.org/2001/XMLSchema"
  targetNamespace="urn:ietf:params:xml:ns:pidf:geopriv10:lmsrc"
  elementFormDefault="qualified"
  attributeFormDefault="unqualified">

  <xs:annotation>
    <xs:appinfo
      source="urn:ietf:params:xml:schema:pidf:geopriv10:lmsrc">
    </xs:appinfo>
    <xs:documentation source="http://www.ietf.org/rfc/rfcXXXX.txt">

```

```

    <!-- [[NOTE TO RFC-EDITOR: Please replace above URL with URL of
    published RFC and remove this note.]] -->
    This schema defines an extension to PIDF-LO that indicates the
    type of source that produced the measurement data used in
    generating the associated location information.
  </xs:documentation>
</xs:annotation>

<xs:element name="source" type="lmsrc:sourceType"/>
<xs:simpleType name="sourceType">
  <xs:list>
    <xs:simpleType>
      <xs:restriction base="xs:token">
        <xs:enumeration value="lis"/>
        <xs:enumeration value="device"/>
        <xs:enumeration value="other"/>
      </xs:restriction>
    </xs:simpleType>
  </xs:list>
</xs:simpleType>
</xs:schema>

```

Measurement Source PIDF-LO Extension Schema

8.3. Base Type Schema

Note that the pattern rules in the following schema wrap due to length constraints. None of the patterns contain whitespace.

```

<?xml version="1.0"?>
<xs:schema
  xmlns:bt="urn:ietf:params:xml:ns:geopriv:lm:basetypes"
  xmlns:xs="http://www.w3.org/2001/XMLSchema"
  targetNamespace="urn:ietf:params:xml:ns:geopriv:lm:basetypes"
  elementFormDefault="qualified"
  attributeFormDefault="unqualified">

  <xs:annotation>
    <xs:appinfo
      source="urn:ietf:params:xml:schema:geopriv:lm:basetypes">
    </xs:appinfo>
    <xs:documentation source="http://www.ietf.org/rfc/rfcXXXX.txt">
      <!-- [[NOTE TO RFC-EDITOR: Please replace above URL with URL of
      published RFC and remove this note.]] -->
      This schema defines a set of base type elements.
    </xs:documentation>
  </xs:annotation>

```

```
<xs:simpleType name="byteType">
  <xs:restriction base="xs:integer">
    <xs:minInclusive value="0"/>
    <xs:maxInclusive value="255"/>
  </xs:restriction>
</xs:simpleType>
<xs:simpleType name="twoByteType">
  <xs:restriction base="xs:integer">
    <xs:minInclusive value="0"/>
    <xs:maxInclusive value="65535"/>
  </xs:restriction>
</xs:simpleType>

<xs:simpleType name="nonNegativeDouble">
  <xs:restriction base="xs:double">
    <xs:minInclusive value="0.0"/>
  </xs:restriction>
</xs:simpleType>
<xs:simpleType name="positiveDouble">
  <xs:restriction base="bt:nonNegativeDouble">
    <xs:minExclusive value="0.0"/>
  </xs:restriction>
</xs:simpleType>

<xs:complexType name="doubleWithRMSError">
  <xs:simpleContent>
    <xs:extension base="xs:double">
      <xs:attribute name="rmsError" type="bt:positiveDouble"/>
      <xs:attribute name="samples" type="xs:positiveInteger"/>
    </xs:extension>
  </xs:simpleContent>
</xs:complexType>
<xs:complexType name="nnDoubleWithRMSError">
  <xs:simpleContent>
    <xs:restriction base="bt:doubleWithRMSError">
      <xs:minInclusive value="0"/>
    </xs:restriction>
  </xs:simpleContent>
</xs:complexType>

<xs:simpleType name="ipAddressType">
  <xs:union memberTypes="bt:IPv6AddressType bt:IPv4AddressType"/>
</xs:simpleType>

<!-- IPv6 format definition -->
<xs:simpleType name="IPv6AddressType">
  <xs:annotation>
    <xs:documentation>
```

```

An IP version 6 address, based on RFC 4291.
</xs:documentation>
</xs:annotation>
<xs:restriction base="xs:token">
  <!-- Fully specified address -->
  <xs:pattern value="[0-9A-Fa-f]{1,4}(:[0-9A-Fa-f]{1,4}){7}"/>
  <!-- Double colon start -->
  <xs:pattern value="(:[0-9A-Fa-f]{1,4}){1,7}"/>
  <!-- Double colon middle -->
  <xs:pattern value="([0-9A-Fa-f]{1,4}:){1,6}
    (:[0-9A-Fa-f]{1,4}){1}"/>
  <xs:pattern value="([0-9A-Fa-f]{1,4}:){1,5}
    (:[0-9A-Fa-f]{1,4}){1,2}"/>
  <xs:pattern value="([0-9A-Fa-f]{1,4}:){1,4}
    (:[0-9A-Fa-f]{1,4}){1,3}"/>
  <xs:pattern value="([0-9A-Fa-f]{1,4}:){1,3}
    (:[0-9A-Fa-f]{1,4}){1,4}"/>
  <xs:pattern value="([0-9A-Fa-f]{1,4}:){1,2}
    (:[0-9A-Fa-f]{1,4}){1,5}"/>
  <xs:pattern value="([0-9A-Fa-f]{1,4}:){1}
    (:[0-9A-Fa-f]{1,4}){1,6}"/>
  <!-- Double colon end -->
  <xs:pattern value="([0-9A-Fa-f]{1,4}:){1,7}:"/>
  <!-- IPv4-Compatible and IPv4-Mapped Addresses -->
  <xs:pattern value="((:0{1,4}){0,3}:[fF]{4})|(0{1,4}:
    (:0{1,4}){0,2}:[fF]{4})|((0{1,4}:){2}
    (:0{1,4})?:[fF]{4})|((0{1,4}:){3}:[fF]{4})
    |((0{1,4}:){4}[fF]{4})|(25[0-5]|2[0-4][0-9]|
    [0-1]?[0-9]?[0-9])\.(25[0-5]|2[0-4][0-9]|
    [0-1]?[0-9]?[0-9])\.(25[0-5]|2[0-4][0-9]|
    [0-1]?[0-9]?[0-9])\.(25[0-5]|2[0-4][0-9]|
    [0-1]?[0-9]?[0-9])"/>
  <!-- The unspecified address -->
  <xs:pattern value=":"/>
</xs:restriction>
</xs:simpleType>

<!-- IPv4 format definition -->
<xs:simpleType name="IPv4AddressType">
  <xs:restriction base="xs:token">
    <xs:pattern value="(25[0-5]|2[0-4][0-9]|[0-1]?[0-9]?[0-9])\.(
      25[0-5]|2[0-4][0-9]|[0-1]?[0-9]?[0-9])\.(
      25[0-5]|2[0-4][0-9]|[0-1]?[0-9]?[0-9])\.(
      25[0-5]|2[0-4][0-9]|[0-1]?[0-9]?[0-9])"/>
  </xs:restriction>
</xs:simpleType>

<!-- MAC address (EUI-48) or EUI-64 address -->

```

```

<xs:simpleType name="macAddressType">
  <xs:restriction base="xs:token">
    <xs:pattern
value="[\da-fA-F]{2}(-[\da-fA-F]{2}){5}((-[\da-fA-F]{2}){2})?" />
    </xs:restriction>
  </xs:simpleType>

</xs:schema>

```

Base Type Schema

8.4. LLDP Measurement Schema

```

<?xml version="1.0"?>
<xs:schema
  xmlns:lldp="urn:ietf:params:xml:ns:geopriv:lm:lldp"
  xmlns:bt="urn:ietf:params:xml:ns:geopriv:lm:basetypes"
  xmlns:xs="http://www.w3.org/2001/XMLSchema"
  targetNamespace="urn:ietf:params:xml:ns:geopriv:lm:lldp"
  elementFormDefault="qualified"
  attributeFormDefault="unqualified">

  <xs:annotation>
    <xs:appinfo
      source="urn:ietf:params:xml:schema:geopriv:lm:lldp">
    </xs:appinfo>
    <xs:documentation source="http://www.ietf.org/rfc/rfcXXXX.txt">
      <!-- [[NOTE TO RFC-EDITOR: Please replace above URL with URL of
        published RFC and remove this note.]] -->
      This schema defines a set of LLDP location measurements.
    </xs:documentation>
  </xs:annotation>

  <xs:import namespace="urn:ietf:params:xml:ns:geopriv:lm:basetypes"/>

  <xs:element name="lldp" type="lldp:lldpMeasurementType"/>
  <xs:complexType name="lldpMeasurementType">
    <xs:complexContent>
      <xs:restriction base="xs:anyType">
        <xs:sequence>
          <xs:element name="chassis" type="lldp:lldpDataTypes"/>
          <xs:element name="port" type="lldp:lldpDataTypes"/>
          <xs:any namespace="##other" processContents="lax"
            minOccurs="0" maxOccurs="unbounded"/>
        </xs:sequence>
        <xs:anyAttribute namespace="##any" processContents="lax"/>
      </xs:restriction>
    </xs:complexContent>

```

```

</xs:complexType>

<xs:complexType name="lldpDataType">
  <xs:simpleContent>
    <xs:extension base="lldp:lldpOctetStringType">
      <xs:attribute name="type" type="bt:byteType"
        use="required" />
    </xs:extension>
  </xs:simpleContent>
</xs:complexType>

<xs:simpleType name="lldpOctetStringType">
  <xs:restriction base="xs:hexBinary">
    <xs:minLength value="1" />
    <xs:maxLength value="255" />
  </xs:restriction>
</xs:simpleType>

</xs:schema>

```

LLDP measurement schema

8.5. DHCP Measurement Schema

```

<?xml version="1.0"?>
<xs:schema
  xmlns:dhcp="urn:ietf:params:xml:ns:geopriv:lm:dhcp"
  xmlns:xs="http://www.w3.org/2001/XMLSchema"
  xmlns:bt="urn:ietf:params:xml:ns:geopriv:lm:basetypes"
  targetNamespace="urn:ietf:params:xml:ns:geopriv:lm:dhcp"
  elementFormDefault="qualified"
  attributeFormDefault="unqualified">

  <xs:annotation>
    <xs:appinfo
      source="urn:ietf:params:xml:schema:geopriv:lm:dhcp">
    </xs:appinfo>
    <xs:documentation source="http://www.ietf.org/rfc/rfcXXXX.txt">
      <!-- [[NOTE TO RFC-EDITOR: Please replace above URL with URL of
        published RFC and remove this note.]] -->
      This schema defines a set of DHCP location measurements.
    </xs:documentation>
  </xs:annotation>

  <xs:import namespace="urn:ietf:params:xml:ns:geopriv:lm:basetypes" />

  <!-- DHCP Relay Agent Information Option -->
  <xs:element name="dhcp-rai" type="dhcp:dhcpType" />

```

```

<xs:complexType name="dhcpType">
  <xs:complexContent>
    <xs:restriction base="xs:anyType">
      <xs:sequence>
        <xs:element name="giaddr" type="bt:ipAddressType"/>
        <xs:element name="circuit"
          type="xs:hexBinary" minOccurs="0"/>
        <xs:element name="remote"
          type="dhcp:dhcpRemoteType" minOccurs="0"/>
        <xs:element name="subscriber"
          type="xs:hexBinary" minOccurs="0"/>
        <xs:any namespace="##other" processContents="lax"
          minOccurs="0" maxOccurs="unbounded"/>
      </xs:sequence>
      <xs:anyAttribute namespace="##any" processContents="lax"/>
    </xs:restriction>
  </xs:complexContent>
</xs:complexType>

<xs:complexType name="dhcpRemoteType">
  <xs:simpleContent>
    <xs:extension base="xs:hexBinary">
      <xs:attribute name="enterprise" type="xs:positiveInteger"
        use="optional"/>
    </xs:extension>
  </xs:simpleContent>
</xs:complexType>

</xs:schema>

```

DHCP measurement schema

8.6. WiFi Measurement Schema

```

<?xml version="1.0"?>
<xs:schema
  xmlns:wifi="urn:ietf:params:xml:ns:geopriv:lm:wifi"
  xmlns:bt="urn:ietf:params:xml:ns:geopriv:lm:basetypes"
  xmlns:gml="http://www.opengis.net/gml"
  xmlns:xs="http://www.w3.org/2001/XMLSchema"
  targetNamespace="urn:ietf:params:xml:ns:geopriv:lm:wifi"
  elementFormDefault="qualified"
  attributeFormDefault="unqualified">

  <xs:annotation>
    <xs:appinfo
      source="urn:ietf:params:xml:schema:geopriv:lm:wifi">
      802.11 location measurements
    </xs:appinfo>
  </xs:annotation>

```

```
</xs:appinfo>
<xs:documentation source="http://www.ietf.org/rfc/rfcXXXX.txt">
  <!-- [[NOTE TO RFC-EDITOR: Please replace above URL with URL of
    published RFC and remove this note.]] -->
  This schema defines a basic set of 802.11 location measurements.
</xs:documentation>
</xs:annotation>

<xs:import namespace="urn:ietf:params:xml:ns:geopriv:lm:basetypes"/>
<xs:import namespace="http://www.opengis.net/gml"/>

<xs:element name="wifi" type="wifi:wifiNetworkType"/>

<xs:complexType name="wifiNetworkType">
  <xs:complexContent>
    <xs:restriction base="xs:anyType">
      <xs:sequence>
        <xs:element name="nicType" type="xs:token"
          minOccurs="0"/>
        <xs:element name="ap" type="wifi:wifiType"
          maxOccurs="unbounded"/>
      </xs:sequence>
      <xs:anyAttribute namespace="##any" processContents="lax"/>
    </xs:restriction>
  </xs:complexContent>
</xs:complexType>

<xs:complexType name="wifiType">
  <xs:complexContent>
    <xs:restriction base="xs:anyType">
      <xs:sequence>
        <xs:element name="bssid" type="wifi:bssidType"/>
        <xs:element name="ssid" type="wifi:ssidType"
          minOccurs="0"/>
        <xs:element name="channel" type="xs:nonNegativeInteger"
          minOccurs="0"/>
        <xs:element name="location" minOccurs="0"
          type="xs:anyType"/>
        <xs:element name="type" type="wifi:networkType"
          minOccurs="0"/>
        <xs:element name="regclass" type="wifi:regclassType"
          minOccurs="0"/>
        <xs:element name="antenna" type="wifi:octetType"
          minOccurs="0"/>
        <xs:element name="flightTime" minOccurs="0"
          type="bt:nnDoubleWithRMSError"/>
        <xs:element name="apSignal" type="wifi:signalType"
          minOccurs="0"/>
      </xs:sequence>
    </xs:restriction>
  </xs:complexContent>
</xs:complexType>
```

```
<xs:element name="deviceSignal" type="wifi:signalType"
  minOccurs="0"/>
<xs:any namespace="##other" processContents="lax"
  minOccurs="0" maxOccurs="unbounded"/>
</xs:sequence>
<xs:attribute name="serving" type="xs:boolean"
  default="false"/>
<xs:anyAttribute namespace="##any" processContents="lax"/>
</xs:restriction>
</xs:complexContent>
</xs:complexType>

<xs:complexType name="bssidType">
  <xs:simpleContent>
    <xs:extension base="bt:macAddressType">
      <xs:attribute name="verified" type="xs:boolean"
        default="false"/>
    </xs:extension>
  </xs:simpleContent>
</xs:complexType>

<!-- Note that this pattern does not prevent multibyte UTF-8
  sequences that result in a SSID longer than 32 octets. -->
<xs:simpleType name="ssidType">
  <xs:restriction base="xs:token">
    <xs:pattern value="(\\[\\da-fA-F]{2}|[^\])}{0,32}"/>
  </xs:restriction>
</xs:simpleType>

<xs:simpleType name="networkType">
  <xs:restriction base="xs:token">
    <xs:pattern value="[a-zA-Z]+"/>
  </xs:restriction>
</xs:simpleType>

<xs:complexType name="regclassType">
  <xs:simpleContent>
    <xs:extension base="wifi:octetType">
      <xs:attribute name="country">
        <xs:simpleType>
          <xs:restriction base="xs:token">
            <xs:pattern value="[A-Z]{2}[OIX]?"/>
          </xs:restriction>
        </xs:simpleType>
      </xs:attribute>
    </xs:extension>
  </xs:simpleContent>
</xs:complexType>
```

```
<xs:simpleType name="octetType">
  <xs:restriction base="xs:nonNegativeInteger">
    <xs:maxInclusive value="255"/>
  </xs:restriction>
</xs:simpleType>

<xs:complexType name="signalType">
  <xs:complexContent>
    <xs:restriction base="xs:anyType">
      <xs:sequence>
        <xs:element name="transmit" type="xs:double"
          minOccurs="0"/>
        <xs:element name="gain" type="xs:double" minOccurs="0"/>
        <xs:element name="rcpi" type="wifi:rssiType"
          minOccurs="0"/>
        <xs:element name="rsni" type="bt:doubleWithRMSError"
          minOccurs="0"/>
        <xs:any namespace="##other" processContents="lax"
          minOccurs="0" maxOccurs="unbounded"/>
      </xs:sequence>
    </xs:restriction>
  </xs:complexContent>
</xs:complexType>

<xs:complexType name="rssiType">
  <xs:simpleContent>
    <xs:extension base="bt:doubleWithRMSError">
      <xs:attribute name="dBm" type="xs:boolean" default="true"/>
    </xs:extension>
  </xs:simpleContent>
</xs:complexType>

<!-- Measurement Request elements -->
<xs:element name="type" type="wifi:networkType"/>
<xs:element name="parameter" type="wifi:parameterType"/>

<xs:complexType name="parameterType">
  <xs:simpleContent>
    <xs:extension base="xs:QName">
      <xs:attribute name="context" use="optional">
        <xs:simpleType>
          <xs:restriction base="xs:token">
            <xs:enumeration value="ap"/>
            <xs:enumeration value="device"/>
          </xs:restriction>
        </xs:simpleType>
      </xs:attribute>
    </xs:extension>
  </xs:simpleContent>
</xs:complexType>
```

```
    </xs:simpleContent>
  </xs:complexType>

</xs:schema>
```

WiFi measurement schema

8.7. Cellular Measurement Schema

```
<?xml version="1.0"?>
<xs:schema
  xmlns:cell="urn:ietf:params:xml:ns:geopriv:lm:cell"
  xmlns:xs="http://www.w3.org/2001/XMLSchema"
  targetNamespace="urn:ietf:params:xml:ns:geopriv:lm:cell"
  elementFormDefault="qualified"
  attributeFormDefault="unqualified">

  <xs:annotation>
    <xs:appinfo
      source="urn:ietf:params:xml:schema:geopriv:lm:cell">
    </xs:appinfo>
    <xs:documentation source="http://www.ietf.org/rfc/rfcXXXX.txt">
      <!-- [[NOTE TO RFC-EDITOR: Please replace above URL with URL of
        published RFC and remove this note.]] -->
      This schema defines a set of cellular location measurements.
    </xs:documentation>
  </xs:annotation>

  <xs:element name="cellular" type="cell:cellularType"/>

  <xs:complexType name="cellularType">
    <xs:complexContent>
      <xs:restriction base="xs:anyType">
        <xs:sequence>
          <xs:choice>
            <xs:element name="servingCell" type="cell:cellType"/>
            <xs:element name="observedCell" type="cell:cellType"/>
          </xs:choice>
          <xs:element name="observedCell" type="cell:cellType"
            minOccurs="0" maxOccurs="unbounded"/>
        </xs:sequence>
        <xs:anyAttribute namespace="##any" processContents="lax"/>
      </xs:restriction>
    </xs:complexContent>
  </xs:complexType>

  <xs:complexType name="cellType">
    <xs:complexContent>
```

```

    <xs:restriction base="xs:anyType">
    <xs:choice>
    <xs:sequence>
    <xs:element name="mcc" type="cell:mccType"/>
    <xs:element name="mnc" type="cell:mncType"/>
    <xs:choice>
    <xs:sequence>
    <xs:choice>
    <xs:element name="rnc" type="cell:cellIdType"/>
    <xs:element name="lac" type="cell:cellIdType"/>
    </xs:choice>
    <xs:element name="cid" type="cell:cellIdType"/>
    </xs:sequence>
    <xs:element name="eucid" type="cell:cellIdType"/>
    </xs:choice>
    <xs:any namespace="##other" processContents="lax"
    minOccurs="0" maxOccurs="unbounded"/>
    </xs:sequence>
    <xs:sequence>
    <xs:element name="sid" type="cell:cellIdType"/>
    <xs:element name="nid" type="cell:cellIdType"/>
    <xs:element name="baseid" type="cell:cellIdType"/>
    <xs:any namespace="##other" processContents="lax"
    minOccurs="0" maxOccurs="unbounded"/>
    </xs:sequence>
    <xs:any namespace="##other" processContents="lax"
    minOccurs="0" maxOccurs="unbounded"/>
    </xs:choice>
    </xs:restriction>
  </xs:complexContent>
</xs:complexType>

<xs:simpleType name="mccType">
  <xs:restriction base="xs:token">
    <xs:pattern value="[0-9]{3}"/>
  </xs:restriction>
</xs:simpleType>

<xs:simpleType name="mncType">
  <xs:restriction base="xs:token">
    <xs:pattern value="[0-9]{2,3}"/>
  </xs:restriction>
</xs:simpleType>

<xs:simpleType name="cellIdType">
  <xs:restriction base="xs:nonNegativeInteger">
    <xs:maxInclusive value="268435455"/> <!-- 2^28 (eucid) -->
  </xs:restriction>

```

```

</xs:simpleType>

<!-- Measurement Request elements -->

<xs:element name="type" type="cell:typeType"/>
<xs:simpleType name="typeType">
  <xs:restriction base="xs:token">
    <xs:enumeration value="gsm"/>
    <xs:enumeration value="umts"/>
    <xs:enumeration value="lte"/>
    <xs:enumeration value="cdma"/>
  </xs:restriction>
</xs:simpleType>

<xs:element name="network" type="cell:networkType"/>
<xs:complexType name="networkType">
  <xs:complexContent>
    <xs:restriction base="xs:anyType">
      <xs:choice>
        <xs:sequence>
          <xs:element name="mcc" type="cell:mccType"/>
          <xs:element name="mnc" type="cell:mncType"/>
        </xs:sequence>
        <xs:element name="nid" type="cell:cellIdType"/>
      </xs:choice>
    </xs:restriction>
  </xs:complexContent>
</xs:complexType>

</xs:schema>

```

Cellular measurement schema

8.8. GNSS Measurement Schema

```

<?xml version="1.0"?>
<xs:schema
  xmlns:gnss="urn:ietf:params:xml:ns:geopriv:lm:gnss"
  xmlns:bt="urn:ietf:params:xml:ns:geopriv:lm:basetypes"
  xmlns:xs="http://www.w3.org/2001/XMLSchema"
  targetNamespace="urn:ietf:params:xml:ns:geopriv:lm:gnss"
  elementFormDefault="qualified"
  attributeFormDefault="unqualified">

  <xs:annotation>
    <xs:appinfo
      source="urn:ietf:params:xml:schema:geopriv:lm:gnss">
    </xs:appinfo>

```

```
<xs:documentation source="http://www.ietf.org/rfc/rfcXXXX.txt">
  <!-- [[NOTE TO RFC-EDITOR: Please replace above URL with URL of
    published RFC and remove this note.]] -->
  This schema defines a set of GNSS location measurements
</xs:documentation>
</xs:annotation>

<xs:import namespace="urn:ietf:params:xml:ns:geopriv:lm:basetypes"/>

<!-- GNSS -->
<xs:element name="gnss" type="gnss:gnssMeasurementType">
  <xs:unique name="gnssSatellite">
    <xs:selector xpath="sat"/>
    <xs:field xpath="@num"/>
  </xs:unique>
</xs:element>

<xs:complexType name="gnssMeasurementType">
  <xs:complexContent>
    <xs:restriction base="xs:anyType">
      <xs:sequence>
        <xs:element name="gnssTime" type="bt:nnDoubleWithRMSError"
          minOccurs="0"/>
        <xs:element name="sat" type="gnss:gnssSatelliteType"
          minOccurs="1" maxOccurs="64"/>
        <xs:any namespace="##other" processContents="lax"
          minOccurs="0" maxOccurs="unbounded"/>
      </xs:sequence>
      <xs:attribute name="system" type="xs:token" use="required"/>
      <xs:attribute name="signal" type="xs:token"/>
      <xs:anyAttribute namespace="##any" processContents="lax"/>
    </xs:restriction>
  </xs:complexContent>
</xs:complexType>

<xs:complexType name="gnssSatelliteType">
  <xs:complexContent>
    <xs:restriction base="xs:anyType">
      <xs:sequence>
        <xs:element name="doppler" type="bt:doubleWithRMSError"/>
        <xs:element name="codephase"
          type="bt:nnDoubleWithRMSError"/>
        <xs:element name="cn0" type="bt:nonNegativeDouble"/>
        <xs:element name="mp" type="bt:positiveDouble"
          minOccurs="0"/>
        <xs:element name="cq" type="gnss:codePhaseQualityType"
          minOccurs="0"/>
        <xs:element name="adr" type="xs:double" minOccurs="0"/>
      </xs:sequence>
    </xs:restriction>
  </xs:complexContent>
</xs:complexType>
```

```

    </xs:sequence>
    <xs:attribute name="num" type="xs:positiveInteger"
        use="required"/>
    </xs:restriction>
</xs:complexContent>
</xs:complexType>

<xs:complexType name="codePhaseQualityType">
  <xs:complexContent>
    <xs:restriction base="xs:anyType">
      <xs:attribute name="continuous" type="xs:boolean"
        default="true"/>
      <xs:attribute name="direct" use="required">
        <xs:simpleType>
          <xs:restriction base="xs:token">
            <xs:enumeration value="direct"/>
            <xs:enumeration value="inverted"/>
          </xs:restriction>
        </xs:simpleType>
      </xs:attribute>
    </xs:restriction>
  </xs:complexContent>
</xs:complexType>
</xs:schema>

```

GNSS measurement Schema

8.9. DSL Measurement Schema

```

<?xml version="1.0"?>
<xs:schema
  xmlns:dsl="urn:ietf:params:xml:ns:geopriv:lm:dsl"
  xmlns:bt="urn:ietf:params:xml:ns:geopriv:lm:basetypes"
  xmlns:xs="http://www.w3.org/2001/XMLSchema"
  targetNamespace="urn:ietf:params:xml:ns:geopriv:lm:dsl"
  elementFormDefault="qualified"
  attributeFormDefault="unqualified">

  <xs:annotation>
    <xs:appinfo
      source="urn:ietf:params:xml:schema:geopriv:lm:dsl">
      DSL measurement definitions
    </xs:appinfo>
    <xs:documentation source="http://www.ietf.org/rfc/rfcXXXX.txt">
      <!-- [[NOTE TO RFC-EDITOR: Please replace above URL with URL of
        published RFC and remove this note.]] -->
      This schema defines a basic set of DSL location measurements.
    </xs:documentation>
  </xs:annotation>

```

```
</xs:annotation>

<xs:import namespace="urn:ietf:params:xml:ns:geopriv:lm:basetypes"/>

<xs:element name="dsl" type="dsl:dslVlanType"/>
<xs:complexType name="dslVlanType">
  <xs:complexContent>
    <xs:restriction base="xs:anyType">
      <xs:choice>
        <xs:element name="l2tp">
          <xs:complexType>
            <xs:complexContent>
              <xs:restriction base="xs:anyType">
                <xs:sequence>
                  <xs:element name="src" type="bt:ipAddressType"/>
                  <xs:element name="dest" type="bt:ipAddressType"/>
                  <xs:element name="session"
                    type="xs:nonNegativeInteger"/>
                </xs:sequence>
              </xs:restriction>
            </xs:complexContent>
          </xs:complexType>
        </xs:element>
        <xs:sequence>
          <xs:element name="an" type="xs:token"/>
          <xs:group ref="dsl:dslSlotPort"/>
        </xs:sequence>
        <xs:sequence>
          <xs:element name="stag" type="dsl:vlanIDType"/>
          <xs:choice>
            <xs:sequence>
              <xs:element name="ctag" type="dsl:vlanIDType"/>
              <xs:group ref="dsl:dslSlotPort" minOccurs="0"/>
            </xs:sequence>
            <xs:group ref="dsl:dslSlotPort"/>
          </xs:choice>
        </xs:sequence>
        <xs:sequence>
          <xs:element name="vpi" type="bt:byteType"/>
          <xs:element name="vci" type="bt:twoByteType"/>
        </xs:sequence>
        <xs:any namespace="##other" processContents="lax"
          minOccurs="0" maxOccurs="unbounded"/>
      </xs:choice>
      <xs:anyAttribute namespace="##other" processContents="lax"/>
    </xs:restriction>
  </xs:complexContent>
</xs:complexType>
```

```
<xs:simpleType name="vlanIDType">
  <xs:restriction base="xs:nonNegativeInteger">
    <xs:maxInclusive value="4095"/>
  </xs:restriction>
</xs:simpleType>
<xs:group name="dslSlotPort">
  <xs:sequence>
    <xs:element name="slot" type="xs:token"/>
    <xs:element name="port" type="xs:token"/>
  </xs:sequence>
</xs:group>
</xs:schema>
```

DSL measurement schema

9. IANA Considerations

This section creates a registry for GNSS types (Section 5.5) and registers the namespaces and schema defined in Section 8.

9.1. IANA Registry for GNSS Types

This document establishes a new IANA registry for "Global Navigation Satellite System (GNSS) types". The registry includes tokens for the GNSS type and for each of the signals within that type. Referring to [RFC5226], this registry operates under "Specification Required" rules. The IESG will appoint an Expert Reviewer who will advise IANA promptly on each request for a new or updated GNSS type.

Each entry in the registry requires the following information:

GNSS name: the name of the GNSS

Brief description: a brief description of the GNSS

GNSS token: a token that can be used to identify the GNSS

Signals: a set of tokens that represent each of the signals that the system provides

Documentation reference: a reference to one or more stable, public specifications that outline usage of the GNSS, including (but not limited to) signal specifications and time systems

The registry initially includes two registrations:

GNSS name: Global Positioning System (GPS)

Brief description: a system of satellites that use spread-spectrum transmission, operated by the US military for commercial and military applications

GNSS token: gps

Signals: L1, L2, L1C, L2C, L5

Documentation reference: Navstar GPS Space Segment/Navigation User Interface [GPS.ICD]

GNSS name: Galileo

Brief description: a system of satellites that operate in the same spectrum as GPS, operated by the European Union for commercial applications

GNSS Token: galileo

Signals: L1, E5A, E5B, E5A+B, E6

Documentation Reference: Galileo Open Service Signal In Space Interface Control Document (SIS ICD) [Galileo.ICD]

9.2. URN Sub-Namespace Registration for urn:ietf:params:xml:ns:pidf:geopriv10:lmsrc

This section registers a new XML namespace, "urn:ietf:params:xml:ns:pidf:geopriv10:lmsrc", as per the guidelines in [RFC3688].

URI: urn:ietf:params:xml:ns:pidf:geopriv10:lmsrc

Registrant Contact: IETF, GEOPRIV working group,
(geopriv@ietf.org), Martin Thomson (martin.thomson@commscope.com).

XML:

```
BEGIN
<?xml version="1.0"?>
<!DOCTYPE html PUBLIC "-//W3C//DTD XHTML 1.0 Strict//EN"
  "http://www.w3.org/TR/xhtml1/DTD/xhtml1-strict.dtd">
<html xmlns="http://www.w3.org/1999/xhtml" xml:lang="en">
  <head>
    <title>Measurement Source for PIDF-LO</title>
  </head>
  <body>
    <h1>Namespace for Location Measurement Source</h1>
```

```
<h2>urn:ietf:params:xml:ns:pidf:geopriv10:lmsrc</h2>
[[NOTE TO IANA/RFC-EDITOR: Please update RFC URL and replace XXXX
with the RFC number for this specification.]]
  <p>See <a href="[[RFC URL]]">RFCXXXX</a>.</p>
</body>
</html>
END
```

9.3. URN Sub-Namespace Registration for urn:ietf:params:xml:ns:geopriv:lm

This section registers a new XML namespace,
"urn:ietf:params:xml:ns:geopriv:lm", as per the guidelines in
[RFC3688].

URI: urn:ietf:params:xml:ns:geopriv:lm

Registrant Contact: IETF, GEOPRIV working group,
(geopriv@ietf.org), Martin Thomson (martin.thomson@commscope.com).

XML:

```
BEGIN
<?xml version="1.0"?>
<!DOCTYPE html PUBLIC "-//W3C//DTD XHTML 1.0 Strict//EN"
  "http://www.w3.org/TR/xhtml1/DTD/xhtml1-strict.dtd">
<html xmlns="http://www.w3.org/1999/xhtml" xml:lang="en">
  <head>
    <title>Measurement Container</title>
  </head>
  <body>
    <h1>Namespace for Location Measurement Container</h1>
    <h2>urn:ietf:params:xml:ns:geopriv:lm</h2>
[[NOTE TO IANA/RFC-EDITOR: Please update RFC URL and replace XXXX
with the RFC number for this specification.]]
    <p>See <a href="[[RFC URL]]">RFCXXXX</a>.</p>
  </body>
</html>
END
```

9.4. URN Sub-Namespace Registration for urn:ietf:params:xml:ns:geopriv:lm:basetypes

This section registers a new XML namespace,
"urn:ietf:params:xml:ns:geopriv:lm:basetypes", as per the guidelines
in [RFC3688].

URI: urn:ietf:params:xml:ns:geopriv:lm:basetypes

Registrant Contact: IETF, GEOPRIV working group,
(geopriv@ietf.org), Martin Thomson (martin.thomson@commscope.com).

XML:

```
BEGIN
<?xml version="1.0"?>
<!DOCTYPE html PUBLIC "-//W3C//DTD XHTML 1.0 Strict//EN"
  "http://www.w3.org/TR/xhtml1/DTD/xhtml1-strict.dtd">
<html xmlns="http://www.w3.org/1999/xhtml" xml:lang="en">
  <head>
    <title>Base Device Types</title>
  </head>
  <body>
    <h1>Namespace for Base Types</h1>
    <h2>urn:ietf:params:xml:ns:geopriv:lm:basetypes</h2>
    [[NOTE TO IANA/RFC-EDITOR: Please update RFC URL and replace XXXX
      with the RFC number for this specification.]]
    <p>See <a href="[[RFC URL]]">RFCXXXX</a>.</p>
  </body>
</html>
END
```

9.5. URN Sub-Namespace Registration for urn:ietf:params:xml:ns:geopriv:lm:lldp

This section registers a new XML namespace,
"urn:ietf:params:xml:ns:geopriv:lm:lldp", as per the guidelines in
[RFC3688].

URI: urn:ietf:params:xml:ns:geopriv:lm:lldp

Registrant Contact: IETF, GEOPRIV working group,
(geopriv@ietf.org), Martin Thomson (martin.thomson@commscope.com).

XML:

```
BEGIN
<?xml version="1.0"?>
<!DOCTYPE html PUBLIC "-//W3C//DTD XHTML 1.0 Strict//EN"
  "http://www.w3.org/TR/xhtml1/DTD/xhtml1-strict.dtd">
<html xmlns="http://www.w3.org/1999/xhtml" xml:lang="en">
  <head>
    <title>LLDP Measurement Set</title>
  </head>
  <body>
    <h1>Namespace for LLDP Measurement Set</h1>
    <h2>urn:ietf:params:xml:ns:geopriv:lm:lldp</h2>
    [[NOTE TO IANA/RFC-EDITOR: Please update RFC URL and replace XXXX
      with the RFC number for this specification.]]
    <p>See <a href="[[RFC URL]]">RFCXXXX</a>.</p>
  </body>
</html>
END
```

9.6. URN Sub-Namespace Registration for urn:ietf:params:xml:ns:geopriv:lm:dhcp

This section registers a new XML namespace,
"urn:ietf:params:xml:ns:geopriv:lm:dhcp", as per the guidelines in
[RFC3688].

URI: urn:ietf:params:xml:ns:geopriv:lm:dhcp

Registrant Contact: IETF, GEOPRIV working group,
(geopriv@ietf.org), Martin Thomson (martin.thomson@commscope.com).

XML:

```
BEGIN
<?xml version="1.0"?>
<!DOCTYPE html PUBLIC "-//W3C//DTD XHTML 1.0 Strict//EN"
  "http://www.w3.org/TR/xhtml1/DTD/xhtml1-strict.dtd">
<html xmlns="http://www.w3.org/1999/xhtml" xml:lang="en">
  <head>
    <title>DHCP Measurement Set</title>
  </head>
  <body>
    <h1>Namespace for DHCP Measurement Set</h1>
    <h2>urn:ietf:params:xml:ns:geopriv:lm:dhcp</h2>
    [[NOTE TO IANA/RFC-EDITOR: Please update RFC URL and replace XXXX
      with the RFC number for this specification.]]
    <p>See <a href="[[RFC URL]]">RFCXXXX</a>.</p>
  </body>
```

```
</html>
END
```

9.7. URN Sub-Namespace Registration for
urn:ietf:params:xml:ns:geopriv:lm:wifi

This section registers a new XML namespace,
"urn:ietf:params:xml:ns:geopriv:lm:wifi", as per the guidelines in
[RFC3688].

URI: urn:ietf:params:xml:ns:geopriv:lm:wifi

Registrant Contact: IETF, GEOPRIV working group,
(geopriv@ietf.org), Martin Thomson (martin.thomson@commscope.com).

XML:

```
BEGIN
<?xml version="1.0"?>
<!DOCTYPE html PUBLIC "-//W3C//DTD XHTML 1.0 Strict//EN"
  "http://www.w3.org/TR/xhtml1/DTD/xhtml1-strict.dtd">
<html xmlns="http://www.w3.org/1999/xhtml" xml:lang="en">
  <head>
    <title>WiFi Measurement Set</title>
  </head>
  <body>
    <h1>Namespace for WiFi Measurement Set</h1>
    <h2>urn:ietf:params:xml:ns:geopriv:lm:wifi</h2>
    [[NOTE TO IANA/RFC-EDITOR: Please update RFC URL and replace XXXX
      with the RFC number for this specification.]]
    <p>See <a href="[[RFC URL]]">RFCXXXX</a>.</p>
  </body>
</html>
END
```

9.8. URN Sub-Namespace Registration for
urn:ietf:params:xml:ns:geopriv:lm:cell

This section registers a new XML namespace,
"urn:ietf:params:xml:ns:geopriv:lm:cell", as per the guidelines in
[RFC3688].

URI: urn:ietf:params:xml:ns:geopriv:lm:cell

Registrant Contact: IETF, GEOPRIV working group,
(geopriv@ietf.org), Martin Thomson (martin.thomson@commscope.com).

XML:

```

BEGIN
<?xml version="1.0"?>
<!DOCTYPE html PUBLIC "-//W3C//DTD XHTML 1.0 Strict//EN"
  "http://www.w3.org/TR/xhtml1/DTD/xhtml1-strict.dtd">
<html xmlns="http://www.w3.org/1999/xhtml" xml:lang="en">
  <head>
    <title>Cellular Measurement Set</title>
  </head>
  <body>
    <h1>Namespace for Cellular Measurement Set</h1>
    <h2>urn:ietf:params:xml:ns:geopriv:lm:cell</h2>
    [[NOTE TO IANA/RFC-EDITOR: Please update RFC URL and replace XXXX
      with the RFC number for this specification.]]
    <p>See <a href="[[RFC URL]]">RFCXXXX</a>.</p>
  </body>
</html>
END

```

9.9. URN Sub-Namespace Registration for urn:ietf:params:xml:ns:geopriv:lm:gns

This section registers a new XML namespace,
"urn:ietf:params:xml:ns:geopriv:lm:gns", as per the guidelines in
[RFC3688].

URI: urn:ietf:params:xml:ns:geopriv:lm:gns

Registrant Contact: IETF, GEOPRIV working group,
(geopriv@ietf.org), Martin Thomson (martin.thomson@commscope.com).

XML:

```

BEGIN
<?xml version="1.0"?>
<!DOCTYPE html PUBLIC "-//W3C//DTD XHTML 1.0 Strict//EN"
  "http://www.w3.org/TR/xhtml1/DTD/xhtml1-strict.dtd">
<html xmlns="http://www.w3.org/1999/xhtml" xml:lang="en">
  <head>
    <title>GNSS Measurement Set</title>
  </head>
  <body>
    <h1>Namespace for GNSS Measurement Set</h1>
    <h2>urn:ietf:params:xml:ns:geopriv:lm:gns</h2>
    [[NOTE TO IANA/RFC-EDITOR: Please update RFC URL and replace XXXX
      with the RFC number for this specification.]]

```

```
    <p>See <a href="[[RFC URL]]">RFCXXXX</a>.</p>
  </body>
</html>
END
```

9.10. URN Sub-Namespace Registration for urn:ietf:params:xml:ns:geopriv:lm:dsl

This section registers a new XML namespace,
"urn:ietf:params:xml:ns:geopriv:lm:dsl", as per the guidelines in
[RFC3688].

URI: urn:ietf:params:xml:ns:geopriv:lm:dsl

Registrant Contact: IETF, GEOPRIV working group,
(geopriv@ietf.org), Martin Thomson (martin.thomson@commscope.com).

XML:

```
  BEGIN
  <?xml version="1.0"?>
  <!DOCTYPE html PUBLIC "-//W3C//DTD XHTML 1.0 Strict//EN"
    "http://www.w3.org/TR/xhtml1/DTD/xhtml1-strict.dtd">
  <html xmlns="http://www.w3.org/1999/xhtml" xml:lang="en">
  <head>
    <title>DSL Measurement Set</title>
  </head>
  <body>
    <h1>Namespace for DSL Measurement Set</h1>
    <h2>urn:ietf:params:xml:ns:geopriv:lm:dsl</h2>
  [[NOTE TO IANA/RFC-EDITOR: Please update RFC URL and replace XXXX
    with the RFC number for this specification.]]
    <p>See <a href="[[RFC URL]]">RFCXXXX</a>.</p>
  </body>
  </html>
  END
```

9.11. XML Schema Registration for Measurement Source Schema

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

URI: urn:ietf:params:xml:schema:pidf:geopriv10:lm:src

Registrant Contact: IETF, GEOPRIV working group, (geopriv@ietf.org),
Martin Thomson (martin.thomson@commscope.com).

Schema: The XML for this schema can be found in Section 8.2 of this document.

9.12. XML Schema Registration for Measurement Container Schema

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

URI: urn:ietf:params:xml:schema:lm

Registrant Contact: IETF, GEOPRIV working group, (geopriv@ietf.org),
Martin Thomson (martin.thomson@commscope.com).

Schema: The XML for this schema can be found in Section 8.1 of this document.

9.13. XML Schema Registration for Base Types Schema

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

URI: urn:ietf:params:xml:schema:lm:basetypes

Registrant Contact: IETF, GEOPRIV working group, (geopriv@ietf.org),
Martin Thomson (martin.thomson@commscope.com).

Schema: The XML for this schema can be found in Section 8.3 of this document.

9.14. XML Schema Registration for LLDP Schema

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

URI: urn:ietf:params:xml:schema:lm:lldp

Registrant Contact: IETF, GEOPRIV working group, (geopriv@ietf.org),
Martin Thomson (martin.thomson@commscope.com).

Schema: The XML for this schema can be found in Section 8.4 of this document.

9.15. XML Schema Registration for DHCP Schema

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

URI: urn:ietf:params:xml:schema:lm:dhcp

Registrant Contact: IETF, GEOPRIV working group, (geopriv@ietf.org),
Martin Thomson (martin.thomson@commscope.com).

Schema: The XML for this schema can be found in Section 8.5 of this document.

9.16. XML Schema Registration for WiFi Schema

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

URI: urn:ietf:params:xml:schema:lm:wifi

Registrant Contact: IETF, GEOPRIV working group, (geopriv@ietf.org),
Martin Thomson (martin.thomson@commscope.com).

Schema: The XML for this schema can be found in Section 8.6 of this document.

9.17. XML Schema Registration for Cellular Schema

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

URI: urn:ietf:params:xml:schema:lm:cellular

Registrant Contact: IETF, GEOPRIV working group, (geopriv@ietf.org),
Martin Thomson (martin.thomson@commscope.com).

Schema: The XML for this schema can be found in Section 8.7 of this document.

9.18. XML Schema Registration for GNSS Schema

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

URI: urn:ietf:params:xml:schema:lm:gnss

Registrant Contact: IETF, GEOPRIV working group, (geopriv@ietf.org),
Martin Thomson (martin.thomson@commscope.com).

Schema: The XML for this schema can be found in Section 8.8 of this document.

9.19. XML Schema Registration for DSL Schema

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

URI: urn:iETF:params:xml:schema:lm:dsl

Registrant Contact: IETF, GEOPRIV working group, (geopriv@ietf.org),
Martin Thomson (martin.thomson@commscope.com).

Schema: The XML for this schema can be found in Section 8.9 of this document.

10. Acknowledgements

Thanks go to Simon Cox for his comments relating to terminology that have helped ensure that this document is aligned with ongoing work in the Open Geospatial Consortium (OGC). Thanks to Neil Harper for his review and comments on the GNSS sections of this document. Thanks to Noor-E-Gagan Singh, Gabor Bajko, Russell Priebe, and Khalid Al-Mufti for their significant input to and suggestions for improving the 802.11 measurements. Thanks to Cullen Jennings for feedback and suggestions. Bernard Aboba provided review and feedback on a range of measurement data definitions. Mary Barnes and Geoff Thompson provided a review and corrections. David Waitzman and John Bressler both noted shortcomings with 802.11 measurements. Keith Drage, Darren Pawson provided expert LTE knowledge.

11. References

11.1. Normative References

[ASCII] , "US-ASCII. Coded Character Set - 7-Bit American Standard Code for Information Interchange. Standard ANSI X3.4-1986, ANSI, 1986.", .

[GPS.ICD] , "Navstar GPS Space Segment/Navigation User Interface", ICD GPS-200, Apr 2000.

[Galileo.ICD]
GJU, "Galileo Open Service Signal In Space Interface Control Document (SIS ICD)", May 2006.

[IANA.enterprise]
IANA, "Private Enterprise Numbers", 2011,
<<http://www.iana.org/assignments/enterprise-numbers>>.

[IEEE.80211]

IEEE, "Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications", IEEE Std 802.11-2012, March 2012.

[IEEE.8021AB]

IEEE, "IEEE Standard for Local and Metropolitan area networks, Station and Media Access Control Connectivity Discovery", IEEE Std 802.1AB-2009, September 2009.

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

[RFC3046] Patrick, M., "DHCP Relay Agent Information Option", RFC 3046, January 2001.

[RFC3315] Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, July 2003.

[RFC3629] Yergeau, F., "UTF-8, a transformation format of ISO 10646", STD 63, RFC 3629, November 2003.

[RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, January 2005.

[RFC3993] Johnson, R., Palaniappan, T., and M. Stapp, "Subscriber-ID Suboption for the Dynamic Host Configuration Protocol (DHCP) Relay Agent Option", RFC 3993, March 2005.

[RFC4119] Peterson, J., "A Presence-based GEOPRIV Location Object Format", RFC 4119, December 2005.

[RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, February 2006.

[RFC4580] Volz, B., "Dynamic Host Configuration Protocol for IPv6 (DHCPv6) Relay Agent Subscriber-ID Option", RFC 4580, June 2006.

[RFC4649] Volz, B., "Dynamic Host Configuration Protocol for IPv6 (DHCPv6) Relay Agent Remote-ID Option", RFC 4649, August 2006.

[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.

- [RFC5952] Kawamura, S. and M. Kawashima, "A Recommendation for IPv6 Address Text Representation", RFC 5952, August 2010.
- [RFC5985] Barnes, M., "HTTP-Enabled Location Delivery (HELD)", RFC 5985, September 2010.
- [RFC5986] Thomson, M. and J. Winterbottom, "Discovering the Local Location Information Server (LIS)", RFC 5986, September 2010.
- [TIA-2000.5] TIA/EIA, "Upper Layer (Layer 3) Signaling Standard for cdma2000(R) Spread Spectrum Systems", TIA-2000.5-D, March 2004.
- [TS.3GPP.23.003] 3GPP, "Numbering, addressing and identification", 3GPP TS 23.003 9.4.0, September 2010.

11.2. Informative References

- [ANSI-TIA-1057] ANSI/TIA, "Link Layer Discovery Protocol for Media Endpoint Devices", TIA 1057, April 2006.
- [DSL.TR025] Wang, R., "Core Network Architecture Recommendations for Access to Legacy Data Networks over ADSL", September 1999.
- [DSL.TR101] Cohen, A. and E. Shrum, "Migration to Ethernet-Based DSL Aggregation", April 2006.
- [GPS.SPOOF] Scott, L., "Anti-Spoofing and Authenticated Signal Architectures for Civil Navigation Signals", ION-GNSS Portland, Oregon, 2003.
- [HARPER] Harper, N., Dawson, M., and D. Evans, "Server-side spoofing and detection for Assisted-GPS", Proceedings of International Global Navigation Satellite Systems Society (IGNSS) Symposium 2009 16, December 2009, <<http://ignss.org/files/Paper16.pdf>>.
- [RFC2661] Townsley, W., Valencia, A., Rubens, A., Pall, G., Zorn, G., and B. Palter, "Layer Two Tunneling Protocol "L2TP"", RFC 2661, August 1999.

- [RFC2865] Rigney, C., Willens, S., Rubens, A., and W. Simpson, "Remote Authentication Dial In User Service (RADIUS)", RFC 2865, June 2000.
- [RFC3688] Mealling, M., "The IETF XML Registry", BCP 81, RFC 3688, January 2004.
- [RFC3693] Cuellar, J., Morris, J., Mulligan, D., Peterson, J., and J. Polk, "Geopriv Requirements", RFC 3693, February 2004.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.
- [RFC6155] Winterbottom, J., Thomson, M., Tschofenig, H., and R. Barnes, "Use of Device Identity in HTTP-Enabled Location Delivery (HELD)", RFC 6155, March 2011.
- [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.

Authors' Addresses

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

Phone: +1 650-353-1925
Email: martin.thomson@skype.net

James Winterbottom
Unaffiliated
AU

Email: a.james.winterbottom@gmail.com

GEOPRIV
Internet-Draft
Intended status: Standards Track
Expires: February 22, 2013

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

M. Thomson
Microsoft
August 21, 2012

Geolocation Policy: A Document Format for Expressing Privacy Preferences
for Location Information
draft-ietf-geopriv-policy-27

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 while 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 Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any

time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on February 22, 2013.

Copyright Notice

Copyright (c) 2012 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	5
2.	Terminology	7
3.	Generic Processing	8
3.1.	Structure of Geolocation Authorization Documents	8
3.2.	Rule Transport	8
4.	Location-specific Conditions	9
4.1.	Geodetic Location Condition Profile	9
4.2.	Civic Location Condition Profile	10
5.	Actions	11
6.	Transformations	12
6.1.	Set Retransmission-Allowed	12
6.2.	Set Retention-Expiry	12
6.3.	Set Note-Well	12
6.4.	Keep Ruleset Reference	13
6.5.	Provide Location	13
6.5.1.	Civic Location Profile	14
6.5.2.	Geodetic Location Profile	15
7.	Examples	18
7.1.	Rule Example with Civic Location Condition	18
7.2.	Rule Example with Geodetic Location Condition	19
7.3.	Rule Example with Civic and Geodetic Location Condition	19
7.4.	Rule Example with Location-based Transformations	20
7.5.	Location Obfuscation Example	22
8.	XML Schema for Basic Location Profiles	26
9.	XML Schema for Geolocation Policy	27
10.	XCAP Usage	29
10.1.	Application Unique ID	29
10.2.	XML Schema	29
10.3.	Default Namespace	29
10.4.	MIME Media Type	29
10.5.	Validation Constraints	29
10.6.	Data Semantics	29
10.7.	Naming Conventions	29
10.8.	Resource Interdependencies	30
10.9.	Authorization Policies	30
11.	IANA Considerations	31
11.1.	Geolocation Policy XML Schema Registration	31
11.2.	Geolocation Policy Namespace Registration	31
11.3.	Geolocation Policy Location Profile Registry	32
11.4.	Basic Location Profile XML Schema Registration	32
11.5.	Basic Location Profile Namespace Registration	33
11.6.	XCAP Application Usage ID	34
12.	Internationalization Considerations	35
13.	Security Considerations	36
13.1.	Introduction	36
13.2.	Obfuscation	36

13.3. Algorithm Limitations 38
13.4. Usability 38
13.5. Location Obscuring Limitations 39
14. References 41
14.1. Normative References 41
14.2. Informative References 41
Appendix A. Acknowledgments 44
Appendix B. Pseudo-Code 45
Authors' Addresses 49

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 access to location information. 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 requestor's identity will likely be revealed as part of this request for location information. 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. Transformations regulate in a location information context 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). It is worth pointing out that the geodetic transformation algorithm Section 6.5.2 deals with privacy risks related to targets that are stationary, as well as to moving targets. However, with respect to movement there are restriction as to what information can be hidden from 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.

In this document we use the term Location Servers as the entities that evaluate the geolocation authorization policies. 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 which 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 Section 4.1 and Section 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. An extension that 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, but allows a `>location-condition>` to be TRUE if an child `>location>` is not understood as long as an understood `>location>` 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 GML [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 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 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 <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 a 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 a 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 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 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 'profile' attribute MUST match with the contained child elements.

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', '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.

1. Choose a geodesic projection with Cartesian coordinates and a surface you want to cover. The maximal distortion of the map may not be too much (see notes below).
2. Given uncertainty "d", choose a grid of so called "landmarks" at a distance (maximal) d of each other.
3. Given a measured location $M=(m,n)$ in 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 local coordinates of the point M in the small grid square that contains it. $0 \leq x, y < 1$.

5. Let $p = 0.2887$ ($=\sqrt{3}/6$) and $q = 0.7113$ ($=1-p$), determine which of the following 8 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 on (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 in 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 in 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 Step3:

SW is mnemonic for south-west, b for bottom, l for left (SW=(l,b)), etc, but the directions of the geodesic projection may be arbitrary, and thus SW may be not south-west 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 equal 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 in 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 short-hand 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 approximately to 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 for our purposes.

We will set up a grid not only for the continental US, 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 US. (On the south 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 largest distance at the latitude of the origin).

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

The same procedure, basically, produces grids for

- * 45 degrees south to 45 degrees north Tropics and subtropics
- * 25 to 50 degrees (both north or south) Continental US
- * 35 to 55 degrees (both north or south) South and Central Europe
- * 45 to 60 degrees (both north or south) Central and North Europe
- * 55 to 65 degrees (both north or south) Scandinavia

* 60 to 70 degrees (both north or south)

Since we do not want to often change grid system (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 points, we start with our chosen origin and place the along the main axes (NS and EW) grid points at a distance d of each other.

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 on east-west direction and to $d_2 = 100/110.6 = 0.904$ degrees in 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 has as we go further away from it.

The distance between two points in the grid at 25 degrees latitude is indeed approx 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). (On 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 south-west of $M=(m,n)$, we calculate

$$i = \text{floor}(m/d1) = \text{floor}(-105/0.993) = -106$$

$$j = \text{floor}(n-o/d2) = \text{floor}(15/0.904) = 16$$

Those are the indexes of SW on the grid. The coordinates of SW are then: $(d1*i, 25+d2*j) = (-105.242, 39.467)$.

Thus:

$$l = d1 * \text{floor}(m/d1) = -105.243$$

$$r = l + d1 = -105.243 + 0.993 = -104.250$$

$$b = o + d2 * \text{floor}(n-o/d2) = 39.467$$

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

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 (0.7113). x is smaller than p . Therefore, only cases 1,4 or 6 could hold.

Also compare y with p (0.2887) and (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:restriction>
    </xs:complexContent>
  </xs:complexType>

```

```
        <xs:attribute name="label" type="xs:string" />
        <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

The following 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 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 usage 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="[URL of published RFC]">RFCXXXX
  [NOTE TO IANA/RFC-EDITOR:
    Please replace XXXX with the RFC number of this
    specification.]</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="[URL of published RFC]">RFCXXXX
  [NOTE TO IANA/RFC-EDITOR:
    Please replace XXXX with the RFC number of this
    specification.]</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 (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, and 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 geolocation (via 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 of these techniques.

The privacy protection requirements for altering location information vary. The two obfuscation algorithms in this document provide a basis for protecting against unauthorized disclosure of location information 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 the drawing in 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.

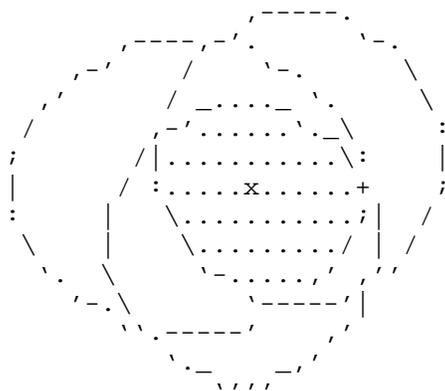


Figure 1: Obfuscation: A Static Target

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 amount of information.

13.3. Algorithm Limitations

The algorithm presented in Section 6.5.2 has some issues where information is leaked: when moving, 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 render a 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 information that the target is, was, or regularly visits, a region of size A_1 , smaller than A . The quotient of the sizes A_1/A should be, even in the worst case, larger than a fixed known number, in order that the user knows what is the maximal information leakage he has. 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 km², he will be leaking, in worst case, the location within a region of size 130 km².

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 [I-D.thomson-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 non-sensical. 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. Location Obscuring Limitations

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.

A caution: 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",
http://portal.opengeospatial.org/files/?artifact_id=4700,
July 2004.
- [NIMA.TR8350.2-3e] OpenGIS, "US National Imagery and Mapping Agency,
"Department of Defense (DoD) World Geodetic System 1984
(WGS 84), Third Edition, NIMA TR8350.2", , January 2000.
- [OGC-06-103r4] OpenGIS, "OpenGIS Implementation Standard for Geographic
information - Simple feature access - Part 1: Common
architecture",
<http://www.opengeospatial.org/docs/06-103r4.pdf>,
May 2011.
- [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

- [I-D.thomson-geopriv-geo-shape] Thomson, M., "Geodetic Shapes for the Representation of
Uncertainty in PIDF-LO",
draft-thomson-geopriv-geo-shape-03 (work in progress),

December 2006.

- [I-D.thomson-geopriv-uncertainty]
Thomson, M. and J. Winterbottom, "Representation of Uncertainty and Confidence in PIDF-LO", draft-thomson-geopriv-uncertainty-07 (work in progress), March 2012.
- [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.
- [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 GIS Workshop on Security and Privacy in GIS and LBS (SPRINGL), ACM.", Nov 2010.

- [ifip07] Ardagna, C., Cremonini, M., Damiani, E., De Capitani di Vimercati, S., and S. Samarati, "Location-privacy protection through obfuscation-based techniques, in: Proceedings of the 21st Annual IFIP WG 11.3 Working Conference on Data and Applications Security, Redondo Beach, CA, USA", July 2007.

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>, 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 [I-D.thomson-geopriv-geo-shape]. Therefore, we would like to thank Martin Thomson for his work in [I-D.thomson-geopriv-geo-shape]. We would also like to thank Martin Thomson, Matt Lepinski and Richard Barnes for their comments regarding the geodetic location transformation procedure. Richard provided us with a detailed text proposal.

Robert Sparks, Martin Thomson, 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, Eric Gray, Russ Housley, Carl Reed, Martin Thomson, Lisa Dusseault, Chris Newman, Jon Peterson, Sam Hartman, Cullen Jennings, Tim Polk, and Brian Rosen.

Appendix B. Pseudo-Code

This section provides an informal description for the algorithm described in Section 6.5.2 in form of pseudo-code.

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   | max   |               |               |
//      |               |               |               |
```

//	lat	lat	Examples	o
//	-45	45	Tropics and subtropics Africa Australia	0
//	25	50	Continental US Mediterranean most of China	25
//	35	55	South and Central Europe	35
//	45	60	Central and North 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 := (l,b);
SE := (r,b);
NW := (l,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

Email: jmpolk@cisco.com

John B. Morris, Jr.

Email: ietf@jmorris.org

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

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

Network Working Group
Internet-Draft
Intended status: Informational
Expires: July 23, 2010

B. Rosen
NeuStar
January 19, 2010

Prefix elements for Road and House Numbers in PIDF-LO
draft-ietf-geopriv-prefix-00

Abstract

RFC4119 updated by RFC5139 defines suffixes for street names and house numbers, but does not define prefixes. Both occur regularly in addresses and CATypes are needed for them. This memo defines STP Street Prefix and HNP house number prefix CATypes.

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and BCP 79.

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

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

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

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

This Internet-Draft will expire on July 23, 2010.

Copyright Notice

Copyright (c) 2010 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 BSD License.

Table of Contents

- 1. Terminology 3
 - 1.1. Introduction 3
 - 1.2. Prefixes in addressing 3
 - 1.3. Security Considerations 3
 - 1.4. IANA Considerations 3
- 2. References 4
 - 2.1. Normative References 4
 - 2.2. Informative References 4
- Author's Address 4

1. Terminology

This document uses terminology described in [RFC4119].

1.1. Introduction

RFC4119 updated by RFC5139 defines suffixes for street names and house numbers, but does not define prefixes. Both occur regularly in addresses and CATypes are needed for them. This memo defines STP Street Prefix and HNP house number prefix CATypes.

1.2. Prefixes in addressing

In [RFC4119] one can define a PIDF for 123 Main Street but not 123 Boulevard Coronado. There is an STS CAtype for a suffix, but no corresponding prefix. [RFC5139] added PRM Premodifier and POM Postmodifier, but those are not suitable for the purpose.

Similarly, one can express 123B Main, but not H123 Main. Although one can include such letters in the house number, most addressing authorities keep the number numeric only to facilitate sorting, and have prefix and suffix fields for alphanumerics that appear in front of or following the numeric house number.

To remedy this situation, new CATypes are required: STP for a street (road) prefix, and HNP for a house number prefix.

1.3. Security Considerations

The XML representation described in this document is designed for inclusion in a PIDF-LO document. As such, it is subject to the same security considerations as are described in [RFC4119]. Considerations relating to the inclusion of this representation in other XML documents are outside the scope of this document.

1.4. IANA Considerations

This document updates the civic address type registry established by [RFC4776]. Two additional value are added:

41	STP	Street (Road) Prefix
42	HNP	House Number Prefix

2. References

2.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC4119] Peterson, J., "A Presence-based GEOPRIV Location Object Format", RFC 4119, December 2005.

2.2. Informative References

- [RFC4776] Schulzrinne, H., "Dynamic Host Configuration Protocol (DHCPv4 and DHCPv6) Option for Civic Addresses Configuration Information", RFC 4776, November 2006.
- [RFC5139] Thomson, M. and J. Winterbottom, "Revised Civic Location Format for Presence Information Data Format Location Object (PIDF-LO)", RFC 5139, February 2008.

Author's Address

Brian Rosen
NeuStar, Inc.
470 Conrad Dr
Mars, PA 16046
US

Email: br@brianrosen.net

GEOPRIV
Internet-Draft
Intended status: Standards Track
Expires: March 10, 2014

M. Thomson
Microsoft
B. Rosen
Neustar
D. Stanley
Aruba Networks
G. Bajko
Nokia
A. Thomson
Cisco Systems, Inc.
September 06, 2013

Relative Location Representation
draft-ietf-geopriv-relative-location-08

Abstract

This document defines an extension to PIDF-LO (RFC4119) for the expression of location information that is defined relative to a reference point. The reference point may be expressed as a geodetic or civic location, and the relative offset may be one of several shapes. An alternative binary representation is described.

Optionally, a reference to a secondary document (such as a map image) can be included, along with the relationship of the map coordinate system to the reference/offset coordinate system to allow display of the map with the reference point and the relative offset.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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

This Internet-Draft will expire on March 10, 2014.

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	3
2.	Conventions used in this document	3
3.	Overview	3
4.	Relative Location	7
4.1.	Relative Coordinate System	7
4.2.	Placement of XML Elements	8
4.3.	Binary Format	8
4.4.	Distances and Angles	9
4.5.	Value Encoding	9
4.6.	Relative Location Restrictions	9
4.7.	Baseline TLVs	9
4.8.	Reference TLV	9
4.9.	Shapes	10
4.9.1.	Point	10
4.9.2.	Circle or Sphere Shape	11
4.9.3.	Ellipse or Ellipsoid Shape	12
4.9.4.	Polygon or Prism Shape	14
4.9.5.	Arc-Band Shape	16
4.10.	Dynamic Location TLVs	18
4.10.1.	Orientation	18
4.10.2.	Speed	18
4.10.3.	Heading	18
4.11.	Secondary Map Metadata	19
4.11.1.	Map URL	19
4.11.2.	Map Coordinate Reference System	19
4.11.3.	Map Example	22
5.	Examples	22
5.1.	Civic PIDF with Polygon Offset	22
5.2.	Geo PIDF with Circle Offset	24
5.3.	Civic TLV with Point Offset	25
6.	Schema Definition	25
7.	Security Considerations	28
8.	IANA Considerations	28

8.1.	Relative Location Registry	29
8.2.	URN Sub-Namespace Registration	30
8.3.	XML Schema Registration	31
8.4.	Geopriv Identifiers Registry	31
8.4.1.	Registration of Two-Dimensional Relative Coordinate Reference System URN	32
8.4.2.	Registration of Three-Dimensional Relative Coordinate Reference System URN	32
9.	Acknowledgements	33
10.	References	33
10.1.	Normative References	33
10.2.	Informative References	35

1. Introduction

This document describes a format for the expression of relative location information.

A relative location is formed of a reference location, plus a relative offset from that reference location. The reference location can be represented in either civic or geodetic form. The reference location can also have dynamic components such as velocity. The relative offset is specified in meters using a Cartesian coordinate system.

In addition to the relative location, an optional URI can be provided to a document that contains a map, floorplan or other spatially oriented information. Applications could use this information to display the relative location. Additional fields allow the map to be oriented and scaled correctly.

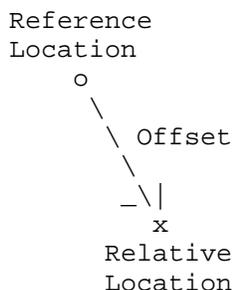
Two formats are included: an XML form that is intended for use in PIDF-LO [RFC4119] and a TLV format for use in other protocols such as those that already convey binary representation of location information defined in [RFC4776].

2. Conventions used in this document

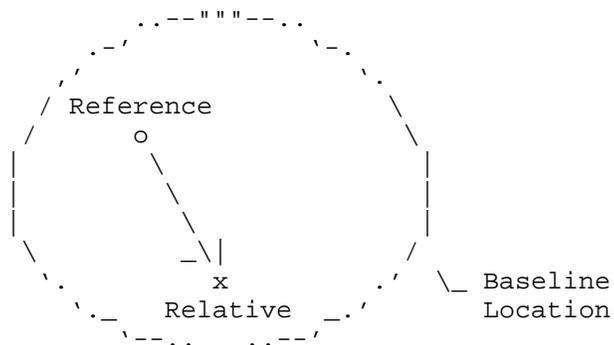
The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

3. Overview

This document describes an extension to PIDF-LO [RFC4119] as updated by [RFC5139] and [RFC5491], to allow the expression of a location as an offset relative to a reference.



This extension allows the creator of a location object to include two location values plus an offset. The two location values, named "baseline" and "reference", combine to form the origin of the offset. The final, relative location is described relative to this reference point.



The "baseline" location is included outside of the <relative-location> element. The baseline location is visible to a client that does not understand relative location (i.e., it ignores the <relative-location> element).

A client that does understand relative location will interpret the location within the relative element as a refinement of the baseline location. This document defines both a "reference" location, which serves as a refinement of the baseline location and the starting point; and an offset, which describes the location of the Target based on this starting point.

Creators of location objects with relative location thus have a choice of how much information to put into the "baseline" location and how much to put into the "reference" location. For example, the baseline location value could be precise enough to specify a building

that contains the relative location, and the reference location could specify a point within the building from which the offset is measured.

Location objects SHOULD NOT have all location information in the baseline location. Doing this would cause clients that do not understand relative location to incorrectly interpret the baseline location (i.e., the reference point) as the actual, precise location of the client. The baseline location is intended to carry a location that encompasses both the reference location and the relative location (i.e., the reference location plus offset).

It is possible to provide a valid relative location with no information in the baseline. However, this provides recipients who do not understand relative location with no information. A baseline location SHOULD include sufficient information to encompass both the reference and relative locations while providing a baseline that is as accurate as possible.

Both the baseline and the reference location are defined either as a geodetic location [OGC.GeoShape] or a civic address [RFC4776]. If the baseline location was expressed as a geodetic location, the reference MUST be geodetic. If the baseline location was expressed as a civic address, the reference MUST be a civic.

Baseline and reference locations MAY also include dynamic location information [RFC5962].

The relative location can be expressed using a point (2- or 3-dimensional), or a shape that includes uncertainty: circle, sphere, ellipse, ellipsoid, polygon, prism or arc-band. Descriptions of these shapes can be found in [RFC5491].

Optionally, a reference to a 'map' document can be provided. The reference is a URI [RFC3986]. The document could be an image or dataset that represents a map, floorplan or other form. The type of document the URI points to is described as a MIME media type [RFC2046]. Metadata in the relative location can include the location of the reference point in the map as well as an orientation (angle from North) and scale to align the document Co-ordinate Reference System (CRS) with the WGS84 [WGS84] CRS. The document is assumed to be useable by the application receiving the PIDF with the relative location to locate the reference point in the map. This document does not describe any mechanisms for displaying or manipulating the document other than providing the reference location, orientation and scale.

As an example, consider a relative location expressed as a point,

relative to a civic location:

```
<presence xmlns="urn:ietf:params:xml:ns:pidf"
  xmlns:dm="urn:ietf:params:xml:ns:pidf:data-model"
  xmlns:gp="urn:ietf:params:xml:ns:pidf:geopriv10"
  xmlns:ca="urn:ietf:params:xml:ns:pidf:geopriv10:civicAddr"
  xmlns:rel="urn:ietf:params:xml:ns:pidf:geopriv10:relative"
  xmlns:gml="http://www.opengis.net/gml"
  xmlns:gs="http://www.opengis.net/pidflo/1.0"
  entity="pres:relative@example.com">
  <dm:device id="relative1">
    <gp:geopriv>
      <gp:location-info>
        <ca:civicAddress xml:lang="en-AU">
          <ca:country>AU</ca:country>
          <ca:A1>NSW</ca:A1>
          <ca:A3>Wollongong</ca:A3>
          <ca:A4>North Wollongong</ca:A4>
          <ca:RD>Flinders</ca:RD>
          <ca:STS>Street</ca:STS>
          <ca:HNO>123</ca:HNO>
        </ca:civicAddress>
        <rel:relative-location>
          <rel:reference>
            <ca:civicAddress xml:lang="en-AU">
              <ca:LMK>Front Door</ca:LMK>
            </ca:civicAddress>
          </rel:reference>
          <rel:offset>
            <gml:Point xmlns:gml="http://www.opengis.net/gml"
              srsName="urn:ietf:params:geopriv:relative:2d">
              <gml:pos>100 50</gml:pos>
            </gml:Point>
          </rel:offset>
        </rel:relative-location>
      </gp:location-info>
      <gp:usage-rules/>
      <gp:method>GPS</gp:method>
      <rel:map>
        <rel:url type="image/png">
          http://example.com/location/map.png
        </rel:url>
        <rel:offset>20. 120.</rel:offset>
        <rel:orientation>29.</rel:orientation>
        <rel:scale>20. -20.</rel:scale>
      </rel:map>
    </gp:geopriv>
  </dm:deviceID>mac:1234567890ab</dm:deviceID>
```

```
<dm:timestamp>2007-06-22T20:57:29Z</dm:timestamp>
</dm:device>
</presence>
```

4. Relative Location

Relative location is a shape (e.g., point, circle, ellipse). The shape is defined with a CRS that has a datum defined as the reference (which appears as a civic address or geodetic location in the tuple), and the shape coordinates as meter offsets North/East of the datum measured in meters (with an optional Z offset relative to datum altitude). An optional angle allows the reference CRS be to rotated with respect to North.

4.1. Relative Coordinate System

The relative coordinate reference system uses a coordinate system with two or three axes.

The baseline and reference locations are used to define a relative datum. The reference location defines the origin of the coordinate system. The centroid of the reference location is used when the reference location contains any uncertainty.

The axes in this coordinate system are originally oriented based on the directions of East, North and Up from the reference location: the first (x) axis increases to the East, the second (y) axis points North, and the optional third (z) axis points Up. All axes of the coordinate system use meters as a basic unit.

Any coordinates in the relative shapes use the described Cartesian coordinate system. In the XML form, this uses a URN of "urn:ietf:params:geopriv:relative:2d" for two-dimensional shapes and "urn:ietf:params:geopriv:relative:3d" for three-dimensional shapes. The binary form uses different shape type identifiers for 2D and 3D shapes.

Dynamic location information [RFC5962] in the baseline or reference location alters relative coordinate system. The resulting Cartesian coordinate system axes are rotated so that the "y" axis is oriented along the direction described by the <orientation> element. The coordinate system also moves as described by the <speed> and <heading> elements.

The single timestamp included in the tuple (or equivalent) element applies to all location elements, including all three components of a relative location: baseline, reference and relative. This is

particularly important when there are dynamic components to these items. A location generator is responsible for ensuring the consistency of these fields.

4.2. Placement of XML Elements

The baseline of the reference location is represented as <location-info> like a normal PIDF-LO. Relative location adds a new <relative-location> element to <location-info>. Within <relative-location>, <reference> and <offset> elements are described. Within <offset> are the shape elements described below. This document extends PIDF-LO as described in [RFC6848].

4.3. Binary Format

This document describes a way to encode the relative location in a binary TLV form for use in other protocols that use TLVs to represent location.

A type-length-value encoding is used.

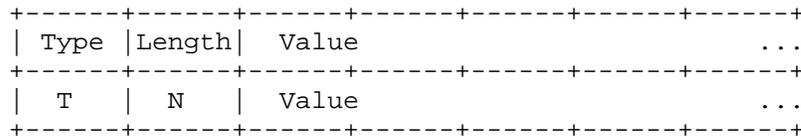


Figure 1: TLV-tuple format

Type field (T) is an 8-bit unsigned integer. The type codes used are registered an IANA-managed "Relative Location Parameters" registry defined by this document, and restricted to not include the values defined by the "CAtypes" registry. This restriction permits a location reference and offset to be coded within the same object without type collisions.

The Length field (N) is defined as an 8-bit unsigned integer. This field can encode values from 0 to 255. The length field describes the number of bytes in the Value. Length does not count the bytes used for the Type or Length.

The Value field is defined separately for each type.

Each element of the relative location has a unique TLV assignment. A relative location encoded in TLV form includes both baseline and reference location TLVs and a reference location TLVs. The reference TLVs are followed by the relative offset, and optional map TLDs described in this document.

4.4. Distances and Angles

All distance measures used in shapes are expressed in meters.

All orientation angles used in shapes are expressed in degrees. Orientation angles are measured from WGS84 Northing to Easting with zero at Northing. Orientation angles in the relative coordinate system start from the second coordinate axis (y or Northing) and increase toward the first axis (x or Easting).

4.5. Value Encoding

The binary form uses single-precision floating point values IEEE 754 [IEEE.754] to represent coordinates, distance and angle measures. Single precision values are 32-bit values with a sign bit, 8 exponent bits and 23 fractional bits. This uses the interchange format defined in [IEEE.754] and Section 3.6 of [RFC1014], that is: sign, biased exponent and significand, with the most significant bit first.

Binary-encoded coordinate values are considered to be a single value without uncertainty. When encoding a value that cannot be exactly represented, the best approximation MUST be selected according to [Clinger1990].

4.6. Relative Location Restrictions

More than one relative shape MUST NOT be included in either a PIDF-LO or TLV encoding of location for a given reference point.

Any error in the reference point transfers to the location described by the relative location. Any errors arising from an implementation not supporting or understanding elements of the reference point directly increases the error (or uncertainty) in the resulting location.

4.7. Baseline TLVs

Baseline locations are described using the formats defined in [RFC4776] or [RFC6225].

4.8. Reference TLV

When a reference is encoded in binary form, the baseline and reference locations are combined in a reference TLV. This TLV is identified with the code 111 and contains civic address TLVs (if the baseline was a civic) or geo TLVs (if the baseline was a geo).

```

+-----+-----+-----+-----+-----+-----+
| 111 |Length| Reference TLVs           |
+-----+-----+-----+-----+-----+-----+

```

Reference TLV

4.9. Shapes

Shape data is used to represent regions of uncertainty for the reference and relative locations. Shape data in the reference location uses a WGS84 [WGS84] CRS. Shape data in the relative location uses a relative CRS.

The XML form for shapes uses Geography Markup Language (GML) [OGC.GML-3.1.1], consistent with the rules in [RFC5491]. Reference locations use the CRS URNs specified in [RFC5491]; relative locations use either a 2D CRS (urn:ietf:params:geopriv:relative:2d), or a 3D (urn:ietf:params:geopriv:relative:3d), depending on the shape type.

The binary form of each shape uses a different shape type for 2d and 3d shapes.

Nine shape type codes are defined.

4.9.1. Point

A point "shape" describes a single point with unknown uncertainty. It consists of a single set of coordinates.

In a two-dimensional CRS, the coordinate includes two values; in a three-dimensional CRS, the coordinate includes three values.

4.9.1.1. XML encoding

A point is represented in GML using the following template:

```

<gml:Point xmlns:gml="http://www.opengis.net/gml"
           srsName="$CRS-URN$">
  <gml:pos>$Coordinate-1 $Coordinate-2$ $Coordinate-3$</gml:pos>
</gml:Point>

```

GML Point Template

Where "\$CRS-URN\$" is replaced by a urn:ietf:params:geopriv:relative:2d or urn:ietf:params:geopriv:relative:3d and "\$Coordinate-3\$" is omitted if the CRS is two-dimensional.

4.9.1.2. TLV encoding

The point shape is introduced by a TLV of 113 for a 2D point and 114 for a 3D point.

```
+-----+-----+
| 113/4|Length|
+-----+-----+-----+-----+
|  Coordinate-1                |
+-----+-----+-----+-----+
|  Coordinate-2                |
+-----+-----+-----+-----+
| (3D-only) Coordinate-3      |
+-----+-----+-----+-----+
```

Point Encoding

4.9.2. Circle or Sphere Shape

A circle or sphere describes a single point with a single uncertainty value in meters.

In a two-dimensional CRS, the coordinate includes two values and the resulting shape forms a circle. In a three-dimensional CRS, the coordinate includes three values and the resulting shape forms a sphere.

4.9.2.1. XML encoding

A circle is represented in and converted from GML using the following template:

```
<gs:Circle xmlns:gml="http://www.opengis.net/gml"
           xmlns:gs="http://www.opengis.net/pidflo/1.0"
           srsName="urn:ietf:params:geopriv:relative:2d">
  <gml:pos>${Coordinate-1} ${Coordinate-2}</gml:pos>
  <gs:radius uom="urn:ogc:def:uom:EPSG::9001">
    $Radius$
  </gs:radius>
</gs:Circle>
```

GML Circle Template

A sphere is represented in and converted from GML using the following template:

```

<gs:Sphere xmlns:gml="http://www.opengis.net/gml"
            xmlns:gs="http://www.opengis.net/pidflo/1.0"
            srsName="urn:ietf:params:geopriv:relative:3d">
  <gml:pos>${Coordinate-1} ${Coordinate-2} ${Coordinate-3}</gml:pos>
  <gs:radius uom="urn:ogc:def:uom:EPSG::9001">
    ${Radius}
  </gs:radius>
</gs:Sphere>

```

GML Sphere Template

4.9.2.2. TLV encoding

A circular shape is introduced by a type code of 115. A spherical shape is introduced by a type code of 116.

```

+-----+-----+
| 115/6|Length|
+-----+-----+-----+-----+
|  Coordinate-1                |
+-----+-----+-----+-----+
|  Coordinate-2                |
+-----+-----+-----+-----+
| (3D-only) Coordinate-3      |
+-----+-----+-----+-----+
|  Radius                      |
+-----+-----+-----+-----+

```

Circle or Sphere Encoding

4.9.3. Ellipse or Ellipsoid Shape

An ellipse or ellipsoid describes a point with an elliptical or ellipsoidal uncertainty region.

In a two-dimensional CRS, the coordinate includes two values, plus a semi-major axis, a semi-minor axis, a semi-major axis orientation (clockwise from North). In a three-dimensional CRS, the coordinate includes three values and in addition to the two-dimensional values, an altitude uncertainty (semi-vertical) is added.

4.9.3.1. XML encoding

An ellipse is represented in and converted from GML using the following template:

```

<gs:Ellipse xmlns:gml="http://www.opengis.net/gml"
             xmlns:gs="http://www.opengis.net/pidflo/1.0"
             srsName="urn:ietf:params:geopriv:relative:2d">
  <gml:pos>$Coordinate-1 $Coordinate-2$</gml:pos>
  <gs:semiMajorAxis uom="urn:ogc:def:uom:EPSG::9001">
    $Semi-Major$
  </gs:semiMajorAxis>
  <gs:semiMinorAxis uom="urn:ogc:def:uom:EPSG::9001">
    $Semi-Minor$
  </gs:semiMinorAxis>
  <gs:orientation uom="urn:ogc:def:uom:EPSG::9102">
    $Orientation$
  </gs:orientation>
</gs:Ellipse>

```

GML Ellipse Template

An ellipsoid is represented in and converted from GML using the following template:

```

<gs:Ellipsoid xmlns:gml="http://www.opengis.net/gml"
              xmlns:gs="http://www.opengis.net/pidflo/1.0"
              srsName="urn:ietf:params:geopriv:relative:3d">
  <gml:pos>$Coordinate-1 $Coordinate-2$ $Coordinate-3$</gml:pos>
  <gs:semiMajorAxis uom="urn:ogc:def:uom:EPSG::9001">
    $Semi-Major$
  </gs:semiMajorAxis>
  <gs:semiMinorAxis uom="urn:ogc:def:uom:EPSG::9001">
    $Semi-Minor$
  </gs:semiMinorAxis>
  <gs:verticalAxis uom="urn:ogc:def:uom:EPSG::9001">
    $Semi-Vertical$
  </gs:verticalAxis>
  <gs:orientation uom="urn:ogc:def:uom:EPSG::9102">
    $Orientation$
  </gs:orientation>
</gs:Ellipsoid>

```

GML Ellipsoid Template

4.9.3.2. TLV encoding

An ellipse is introduced by a type code of 117 and an ellipsoid is introduced by a type code of 118.

```

+-----+-----+
| 117/8|Length|
+-----+-----+

```

Coordinate-1	
Coordinate-2	
(3D-only) Coordinate-3	
Semi-Major Axis	Semi-Minor Axis
Orientation	(3D) Semi-Vertical Axis

Ellipse or Ellipsoid Encoding

4.9.4. Polygon or Prism Shape

A polygon or prism include a number of points that describe the outer boundary of an uncertainty region. A prism also includes an altitude for each point and prism height.

At least 3 points MUST be included in a polygon. In order to interoperate with existing systems, an encoding SHOULD include 15 or fewer points, unless the recipient is known to support larger numbers.

4.9.4.1. XML Encoding

A polygon is represented in and converted from GML using the following template:

```
<gml:Polygon xmlns:gml="http://www.opengis.net/gml"
  srsName="urn:ietf:params:geopriv:relative:2d">
  <gml:exterior>
    <gml:LinearRing>
      <gml:posList>
        $Coordinate1-1$ $Coordinate1-2$
        $Coordinate2-1$ $Coordinate2-2$
        $Coordinate3-1$ ...
        ...
        $CoordinateN-1$ $CoordinateN-2$
        $Coordinate1-1$ $Coordinate1-2$
      </gml:posList>
    </gml:LinearRing>
  </gml:exterior>
</gml:Polygon>
```

GML Polygon Template

Alternatively, a series of "pos" elements can be used in place of the single "posList". Each "pos" element contains two or three coordinate values.

Note that the first point is repeated at the end of the sequence of coordinates and no explicit count of the number of points is provided.

A GML polygon that includes altitude cannot be represented perfectly in TLV form. When converting to the binary representation, a two dimensional CRS is used and altitude is removed from each coordinate.

A prism is represented in and converted from GML using the following template:

```
<gs:Prism xmlns:gml="http://www.opengis.net/gml"
  xmlns:gs="http://www.opengis.net/pidflo/1.0"
  srsName="urn:ietf:params:geopriv:relative:3d">
  <gs:base>
    <gml:Polygon>
      <gml:exterior>
        <gml:LinearRing>
          <gml:posList>
            $Coordinate1-1$ $Coordinate1-2$ $Coordinate1-3$
            $Coordinate2-1$ $Coordinate2-2$ $Coordinate2-3$
            $Coordinate2-1$ ... ..
            ...
            $CoordinateN-1$ $CoordinateN-2$ $CoordinateN-3$
            $Coordinate1-1$ $Coordinate1-2$ $Coordinate1-3$
          </gml:posList>
        </gml:LinearRing>
      </gml:exterior>
    </gml:Polygon>
  </gs:base>
  <gs:height uom="urn:ogc:def:uom:EPSG::9001">
    $Height$
  </gs:height>
</gs:Prism>
```

GML Prism Template

Alternatively, a series of "pos" elements can be used in place of the single "posList". Each "pos" element contains three coordinate values.

4.9.4.2. TLV Encoding

A polygon containing 2D points uses a type code of 119. A polygon with 3D points uses a type code of 120. A prism uses a type code of 121. The number of points can be inferred from the length of the TLV.

```

+-----+-----+
|119-21|Length|
+-----+-----+-----+
| (3D-only) Height |
+-----+-----+-----+
| Coordinate1-1 |
+-----+-----+-----+
| Coordinate1-2 |
+-----+-----+-----+
| (3D-only) Coordinate1-3 |
+-----+-----+-----+
| Coordinate2-1 |
+-----+-----+-----+
...
+-----+-----+-----+
| CoordinateN-1 |
+-----+-----+-----+
| CoordinateN-2 |
+-----+-----+-----+
| (3D-only) CoordinateN-3 |
+-----+-----+-----+

```

Polygon or Prism Encoding

Note that unlike the polygon representation in GML, the first and last points are not the same point in the TLV representation. The duplicated point is removed from the binary form.

4.9.5. Arc-Band Shape

An arc-band describes a region constrained by a range of angles and distances from a predetermined point. This shape can only be provided for a two-dimensional CRS.

Distance and angular measures are defined in meters and degrees respectively. Both are encoded as single precision floating point values.

4.9.5.1. XML encoding

An arc-band is represented in and converted from GML using the following template:

```
<gs:ArcBand xmlns:gml="http://www.opengis.net/gml"
            xmlns:gs="http://www.opengis.net/pidflo/1.0"
            srsName="urn:ietf:params:geopriv:relative:2d">
  <gml:pos>$Coordinate-1$ $Coordinate-2$</gml:pos>
  <gs:innerRadius uom="urn:ogc:def:uom:EPSG::9001">
    $Inner-Radius$
  </gs:innerRadius>
  <gs:outerRadius uom="urn:ogc:def:uom:EPSG::9001">
    $Outer-Radius$
  </gs:outerRadius>
  <gs:startAngle uom="urn:ogc:def:uom:EPSG::9102">
    $Start-Angle$
  </gs:startAngle>
  <gs:openingAngle uom="urn:ogc:def:uom:EPSG::9102">
    $Opening-Angle$
  </gs:openingAngle>
</gs:ArcBand>
```

GML Arc-Band Template

4.9.5.2. TLV Encoding

An arc-band is introduced by a type code of 122.

```
+-----+-----+
| 122 |Length|
+-----+-----+
|  Coordinate                |
+-----+-----+
|  Coordinate                |
+-----+-----+-----+-----+
|  Inner Radius              | Outer Radius                |
+-----+-----+-----+-----+
|  Start Angle              | Opening Angle              |
+-----+-----+-----+-----+
```

Arc-Band Encoding

4.10. Dynamic Location TLVs

Dynamic location elements use the definitions in [RFC5962].

4.10.1. Orientation

The orientation of the target is described using one or two angles. Orientation uses a type code of 123.

```
+-----+-----+
| 123  |Length|
+-----+-----+
|           Angle           |
+-----+-----+
| (Optional) Angle         |
+-----+-----+
```

Dynamic Orientation TLVs

4.10.2. Speed

The speed of the target is a scalar value in meters per second. Speed uses a type code of 124.

```
+-----+-----+
| 124  |Length|
+-----+-----+
|           Speed           |
+-----+-----+
```

Dynamic Speed TLVs

4.10.3. Heading

The heading, or direction of travel, is described using one or two angles. Heading uses a type code of 125.

```
+-----+-----+
| 125  |Length|
+-----+-----+
|           Angle           |
+-----+-----+
| (Optional) Angle         |
+-----+-----+
```

Dynamic Heading TLVs

4.11. Secondary Map Metadata

The optional "map" URL can be used to provide a user of relative location with a visual reference for the location information. This document does not describe how the recipient uses the map nor how it locates the reference or offset within the map. Maps can be simple images, vector files, 2-D or 3-D geospatial databases, or any other form of representation understood by both the sender and recipient.

4.11.1. Map URL

In XML, the map is a <map> element defined within <relative-location> and contains the URL. The URL is encoded as a UTF-8 encoded string. An "http:" ([RFC2616]) or "https:" ([RFC2818]) URL MUST be used unless the entity creating the PIDF-LO is able to ensure that authorized recipients of this data are able to use other URI schemes. A "type" attribute MUST be present and specifies the kind of map the URL points to. Map types are specified as MIME media types as recorded in the IANA Media Types registry. For example <map type="image/png">https://www.example.com/floorplans/123South/floor-2</map>.

In binary, the map type is a separate TLV from the map URL. The media type uses a type code of 126; the URL uses a type code of 127.

```
+-----+-----+-----+-----+-----+--  --+-----+
| 126 |Length|   Map Media Type           ...
+-----+-----+-----+-----+-----+--  --+-----+
| 127 |Length|   Map Image URL           ...
+-----+-----+-----+-----+-----+--  --+-----+
```

Map URL TLVs

Note that the binary form restricts data to 255 octets. This restriction could be problematic for URLs in particular. Applications that use the XML form, but cannot guarantee that a binary form won't be used, are encouraged to limit the size of the URL to fit within this restriction.

4.11.2. Map Coordinate Reference System

The CRS used by the map depends on the type of map. For example, a map described by a 3-D geometric model of the building may contain a complete CRS description in it. For some kinds of maps, typically described as images, the CRS used within the map must define the following:

- o The CRS origin

- o The CRS axes used and their orientation
- o The unit of measure used

This document provides elements that allow for a mapping between the local coordinate reference system used for the relative location and the coordinate reference system used for the map where they are not the same.

4.11.2.1. Map Reference Point Offset

This optional element identifies the coordinates of the reference point as it appears in the map. This value is measured in a map-type dependent manner, using the coordinate system of the map.

For image maps, coordinates start from the upper left corner and coordinates are first counted by column with positive values to the right; then rows are counted with positive values toward the bottom of the image. For such an image, the first item is columns, the second rows and any third value applies to any third dimension used in the image coordinate space.

The <offset> element contains 2 (or 3) coordinates similar to a GML "pos". For example:

```
<offset> 2670.0 1124.0 1022.0</offset>
```

Map Reference Point Example XML

The map reference point uses a type code of 129.

```
+-----+-----+
| 129 |Length|
+-----+-----+-----+
|  Coordinate-1 |
+-----+-----+-----+
|  Coordinate-2 |
+-----+-----+-----+
| (3D-only) Coordinate-3 |
+-----+-----+-----+
```

Map Reference Point Coordinates TLV

If omitted, a value containing all zeros is assumed. If the coordinates provided contain fewer values than are needed, the first value from the set is applied in place of any absent values. Thus, if a single value is provided, that value is used for Coordinate-2 and Coordinate-3 (if required). If two values are provided and three

are required, the value of Coordinate-1 is used in place of Coordinate-3.

4.11.2.2. Map Orientation

The map orientation includes the orientation of the map direction in relation to the Earth. Map orientation is expressed relative to the orientation of the relative coordinate system. This means that map orientation with respect to WGS84 North is the sum of the orientation field, plus any orientation included in a dynamic portion of the reference location. Both values default to zero if no value is specified.

This type uses a single precision floating point value of degrees relative to North.

In XML, the <orientation> element contains a single floating point value, example <orientation>67.00</orientation>. In TLV form map orientation uses the code 130:

```
+-----+-----+-----+-----+-----+
|  130 |Length|  Angle  |                |
+-----+-----+-----+-----+-----+
```

Map Orientation TLV

4.11.2.3. Map Scale

The optional map scale describes the relationship between the units of measure used in the map, relative to the meters unit used in the relative coordinate system.

This type uses a sequence of IEEE 754 [IEEE.754] single precision floating point values to represent scale as a sequence of numeric values. The units of these values are dependent on the type of map, and could for example be pixels per meter for an image.

A scaling factor is provided for each axis in the coordinate system. For a two-dimensional coordinate system, two values are included to allow for different scaling along the x and y axes independently. For a three-dimensional coordinate system, three values are specified for the x, y and z axes. Decoders can determine the number of scaling factors by examining the length field.

Alternatively, a single scaling value MAY be used to apply the same scaling factor to all coordinate components.

Images that use a rows/columns coordinate system often use a left-handed coordinate system. A negative value for the y/rows-axis scaling value can be used to account for any change in direction between the y-axis used in the relative coordinate system and the rows axis of the image coordinate system.

In XML, the <scale> element MAY contain a single scale value, or MAY contain 2 (or 3) values in XML list form. In TLV form, scale uses a type code of 131. The length of the TLV determines how many scale values are present:

```
+-----+-----+-----+-----+-----+-----+
| 131 |Length|  Scale(s)          ...
+-----+-----+-----+-----+-----+-----+
```

Map Scale TLV

4.11.3. Map Example

An example of expressing a map is:

```
<rel:map>
  <rel:url type="image/jpeg">
    http://example.com/map.jpg
  </rel:url>
  <rel:offset>200 210</rel:offset>
  <rel:orientation>68</rel:orientation>
  <rel:scale>2.90 -2.90</rel:scale>
</rel:map>
```

Map Example

5. Examples

The examples in this section combine elements from [RFC3863], [RFC4119], [RFC4479], [RFC5139], and [OGC.GeoShape].

5.1. Civic PIDF with Polygon Offset

```
<presence xmlns="urn:ietf:params:xml:ns:pidf"
  xmlns:dm="urn:ietf:params:xml:ns:pidf:data-model"
  xmlns:gp="urn:ietf:params:xml:ns:pidf:geopriv10"
  xmlns:ca="urn:ietf:params:xml:ns:pidf:geopriv10:civicAddr"
  xmlns:rel="urn:ietf:params:xml:ns:pidf:geopriv10:relative"
  xmlns:gml="http://www.opengis.net/gml"
  xmlns:gs="http://www.opengis.net/pidflo/1.0"
  entity="pres:ness@example.com">
  <dm:device id="nesspc-1">
```

```

<gp:geopriv>
  <gp:location-info>
    <ca:civicAddress xml:lang="en-AU">
      <ca:country>AU</ca:country>
      <ca:A1>NSW</ca:A1>
      <ca:A3>Wollongong</ca:A3>
      <ca:A4>North Wollongong</ca:A4>
      <ca:RD>Flinders</ca:RD>
      <ca:STS>Street</ca:STS>
      <ca:HNO>123</ca:HNO>
    </ca:civicAddress>
    <rel:relative-location>
      <rel:reference>
        <ca:civicAddress xml:lang="en-AU">
          <ca:LMK>Front Door</ca:LMK>
          <ca:BLD>A</ca:BLD>
          <ca:FLR>I</ca:FLR>
          <ca:ROOM>113</ca:ROOM>
        </ca:civicAddress>
      </rel:reference>
      <rel:offset>
        <gml:Polygon xmlns:gml="http://www.opengis.net/gml"
          srsName="urn:ietf:params:geopriv:relative:2d">
          <gml:exterior>
            <gml:LinearRing>
              <gml:pos>433.0 -734.0</gml:pos> <!--A-->
              <gml:pos>431.0 -733.0</gml:pos> <!--F-->
              <gml:pos>431.0 -732.0</gml:pos> <!--E-->
              <gml:pos>433.0 -731.0</gml:pos> <!--D-->
              <gml:pos>434.0 -732.0</gml:pos> <!--C-->
              <gml:pos>434.0 -733.0</gml:pos> <!--B-->
              <gml:pos>433.0 -734.0</gml:pos> <!--A-->
            </gml:LinearRing>
          </gml:exterior>
        </gml:Polygon>
      </rel:offset>
    </rel:relative-location>
  </gp:location-info>
  <gp:usage-rules/>
  <gp:method>GPS</gp:method>
</gp:geopriv>
<dm:deviceID>mac:1234567890ab</dm:deviceID>
<dm:timestamp>2007-06-22T20:57:29Z</dm:timestamp>
</dm:device>
</presence>

```

5.2. Geo PIDF with Circle Offset

```
<?xml version="1.0" encoding="UTF-8"?>
  <presence xmlns="urn:ietf:params:xml:ns:pidf"
    xmlns:dm="urn:ietf:params:xml:ns:pidf:data-model"
    xmlns:gp="urn:ietf:params:xml:ns:pidf:geopriv10"
    xmlns:rel="urn:ietf:params:xml:ns:pidf:geopriv10:relative"
    xmlns:gml="http://www.opengis.net/gml"
    xmlns:gs="http://www.opengis.net/pidflo/1.0"
    entity="pres:point2d@example.com">
    <dm:device id="point2d">
      <gp:geopriv>
        <gp:location-info>
          <gs:Circle srsName="urn:ogc:def:crs:EPSG::4326">
            <gml:pos>-34.407 150.883</gml:pos>
            <gs:radius uom="urn:ogc:def:uom:EPSG::9001">
              50.0
            </gs:radius>
          </gs:Circle>
          <rel:relative-location>
            <rel:reference>
              <gml:Point srsName="urn:ogc:def:crs:EPSG::4326">
                <gml:pos>-34.407 150.883</gml:pos>
              </gml:Point>
            </rel:reference>
            <rel:offset>
              <gs:Circle xmlns:gml="http://www.opengis.net/gml"
                srsName="urn:ietf:params:geopriv:relative:2d">
                <gml:pos>500.0 750.0</gml:pos>
                <gs:radius uom="urn:ogc:def:uom:EPSG::9001">
                  5.0
                </gs:radius>
              </gs:Circle>
            </rel:offset>
            <rel:map>
              <rel:url type="image/png">
                https://www.example.com/flrpln/123South/flr-2
              </rel:url>
              <rel:offset>2670.0 1124.0 1022.0</rel:offset>
              <rel:orientation>67.00</rel:orientation>
              <rel:scale>10 -10</rel:scale>
            </rel:map>
          </rel:relative-location>
        </gp:location-info>
        <gp:usage-rules/>
        <gp:method>Wiremap</gp:method>
      </gp:geopriv>
    </dm:deviceID>mac:1234567890ab</dm:deviceID>
```

```
<dm:timestamp>2007-06-22T20:57:29Z</dm:timestamp>
</dm:device>
</presence>
```

5.3. Civic TLV with Point Offset

Type	Value
0	en
1	IL
3	Chicago
34	Wacker
18	Drive
19	3400
112	Reference
25	Building A
27	Floor 6
26	Suite 213
28	Reception Area
115	100 70
126	image/png
127	http://maps.example.com/3400Wacker/A6
129	0.0 4120.0
130	113.0
131	10.6

6. Schema Definition

Note: The pattern value for "mimeType" has been folded onto multiple lines. Whitespace has been added to conform to comply with document formatting restrictions. Extra whitespace around the line endings MUST be removed before using this schema.

```
<?xml version="1.0"?>
<xs:schema
  xmlns:rel="urn:ietf:params:xml:ns:pidf:geopriv10:relative"
  xmlns:xs="http://www.w3.org/2001/XMLSchema"
  xmlns:gml="http://www.opengis.net/gml"
  targetNamespace="urn:ietf:params:xml:ns:pidf:geopriv10:relative"
  elementFormDefault="qualified"
  attributeFormDefault="unqualified">

  <!-- [[NOTE TO RFC-EDITOR: Please replace all instances of the URL
        'http://ietf.org/rfc/rfcXXXX.txt' with the URL of published
        document and remove this note.]] -->

  <xs:annotation>
    <xs:appinfo
      source="urn:ietf:params:xml:schema:pidf:geopriv10:relative">
      Relative Location for PIDF-LO
    </xs:appinfo>
    <xs:documentation source="http://ietf.org/rfc/rfcXXXX.txt">
      This schema defines a location representation that allows for
      the description of locations that are relative to another.
      An optional map reference is also defined.
    </xs:documentation>
  </xs:annotation>

  <xs:import namespace="http://www.opengis.net/gml"/>

  <xs:element name="relative-location" type="rel:relativeType"/>

  <xs:complexType name="relativeType">
    <xs:complexContent>
      <xs:restriction base="xs:anyType">
        <xs:sequence>
          <xs:element name="reference" type="rel:referenceType"/>
          <xs:element name="offset" type="rel:offsetType"/>
          <xs:any namespace="##any" processContents="lax"
            minOccurs="0" maxOccurs="unbounded"/>
        </xs:sequence>
        <xs:anyAttribute namespace="##other" processContents="lax"/>
      </xs:restriction>
    </xs:complexContent>
  </xs:complexType>
```

```
<xs:complexType name="referenceType">
  <xs:complexContent>
    <xs:restriction base="xs:anyType">
      <xs:sequence>
        <xs:any namespace="##other" processContents="lax"
          minOccurs="0" maxOccurs="unbounded"/>
      </xs:sequence>
    </xs:restriction>
  </xs:complexContent>
</xs:complexType>

<xs:complexType name="offsetType">
  <xs:complexContent>
    <xs:restriction base="xs:anyType">
      <xs:sequence>
        <xs:element ref="gml:_Geometry"/>
        <xs:any namespace="##other" processContents="lax"
          minOccurs="0" maxOccurs="unbounded"/>
      </xs:sequence>
    </xs:restriction>
  </xs:complexContent>
</xs:complexType>

<xs:element name="map" type="rel:mapType"/>
<xs:complexType name="mapType">
  <xs:complexContent>
    <xs:restriction base="xs:anyType">
      <xs:sequence>
        <xs:element name="url" type="rel:mapUrlType"/>
        <xs:element name="offset" type="rel:doubleList"
          minOccurs="0"/>
        <xs:element name="orientation" type="rel:doubleList"
          minOccurs="0"/>
        <xs:element name="scale" type="rel:doubleList"
          minOccurs="0"/>
      </xs:sequence>
    </xs:restriction>
  </xs:complexContent>
</xs:complexType>

<xs:complexType name="mapUrlType">
  <xs:simpleContent>
    <xs:extension base="xs:anyURI">
      <xs:attribute name="type" type="rel:mimeType"
        default="application/octet-stream"/>
    </xs:extension>
  </xs:simpleContent>
</xs:complexType>
```

```

<xs:simpleType name="mimeType">
  <xs:restriction base="xs:token">
    <xs:pattern value="(!#$$%&#x27;\*\+\-\.\dA-Z^_`a-z\|~]+
      /(!#$$%&#x27;\*\+\-\.\dA-Z^_`a-z\|~]+([\t ]*([\t ])*(!#$$%&#x27;
      \*\+\-\.\dA-Z^_`a-z\|~]+=(!#$$%&#x27;\*\+\-\.\dA-Z^_`a-z\|~]+|
      &quot;(!#-\[\\]-~)|[\t ]*|\\[\\t !-~])*&quot;)*"/>
    </xs:restriction>
  </xs:simpleType>

  <xs:simpleType name="doubleList">
    <xs:list itemType="xs:double"/>
  </xs:simpleType>

</xs:schema>

```

xml schema relative-location

7. Security Considerations

This document describes a data format. To a large extent, security properties of this depend on how this data is used.

Privacy for location data is typically important. Adding relative location may increase the precision of the location, but does not otherwise alter its privacy considerations, which are discussed in [RFC4119].

The map URL provided in a relative location could accidentally reveal information if a Location Recipient uses the URL to acquire the map. The coverage area of a map, or parameters of the URL itself, could provide information about the location of a Target. In combination with other information that could reveal the set of potential Targets that the Location Recipient has location information for, acquiring a map could leak significant information. In particular, it is important to note that the Target and Location Recipient are often the same entity.

Access to map URLs MUST be secured with TLS [RFC5246] (that is, restricting the map URL to be an https URI), unless the map URL cannot leak information about the Target's location. This restricts information about the map URL to the entity serving the map request. If the map URL conveys more information about a target than a map server is authorized to receive, that URL MUST NOT be included in the PIDF-LO.

8. IANA Considerations

8.1. Relative Location Registry

This document creates a new registry called "Relative Location Parameters". This shares a page, entitled "Civic and Relative Location Parameters" with the existing "Civic Address Types Registry (CATypes)" registry. As defined in [RFC5226], this new registry operates under "IETF Review" rules.

The content of this registry includes:

Relative Location Code: Numeric identifier, assigned by IANA.

Brief description: Short description identifying the meaning of the element.

Reference to published specification: A stable reference to an RFC which describes the value in sufficient detail so that interoperability between independent implementations is possible.

Values requested to be assigned into this registry MUST NOT conflict with values assigned in the "Civic Address Types Registry (CATypes)" registry or vice versa, unless the IANA considerations section for the new value explicitly overrides this prohibition and the document defining the value describes how conflicting TLV codes will be interpreted by implementations. To ensure this, the CATypes entries are explicitly reserved in the initial values table below. Those reserved entries can be changed, but only with caution as explained here.

To make this clear for future users of the registry, the following note is added to the "Civic Address Types Registry (CATypes)": The registration of new values should be accompanied by a corresponding reservation in the "Relative Location Parameters" registry. Similarly, the "Relative Location Parameters" registry bears the note: The registration of new values should be accompanied by a corresponding reservation in the "Civic Address Types Registry (CATypes)" registry.

The values defined are:

RLtype	description	Reference
0-40 128	RESERVED by CATypes registry	this RFC & RFC4776
111	relative location reference	this RFC
113	relative location shape 2D point	this RFC

114	relative location shape 3D point	this RFC
115	relative location shape circular	this RFC
116	relative location shape spherical	this RFC
117	relative location shape elliptical	this RFC
118	relative location shape ellipsoid	this RFC
119	relative location shape 2D polygon	this RFC
120	relative location shape 3D polygon	this RFC
121	relative location shape prism	this RFC
122	relative location shape arc-band	this RFC
123	relative location dynamic orientation	this RFC
124	relative location dynamic speed	this RFC
125	relative location dynamic heading	this RFC
126	relative location map type	this RFC
127	relative location map URI	this RFC
129	relative location map coordinates	this RFC
130	relative location map angle	this RFC
131	relative location map scale	this RFC

8.2. URN Sub-Namespace Registration

This document registers a new XML namespace, as per the guidelines in [RFC3688]).

URI: urn:ietf:params:xml:ns:pidf:geopriv10:relative

Registrant Contact: IETF, GEOPRIV working group (geopriv@ietf.org),
Martin Thomson (martin.thomson@skype.net).

XML:

```

BEGIN
<?xml version="1.0"?>
<!DOCTYPE html PUBLIC "-//W3C//DTD XHTML 1.0 Strict//EN"
    "http://www.w3.org/TR/xhtml1/DTD/xhtml1-strict.dtd">
<html xmlns="http://www.w3.org/1999/xhtml" xml:lang="en">
  <head>
    <title>GEOPRIV Relative Location</title>
  </head>
  <body>
    <h1>Format for representing relative location</h1>
    <h2>urn:ietf:params:xml:ns:pidf:geopriv10:relative</h2>
    <p>See <a href="http://www.rfc-editor.org/rfc/rfcXXXX.txt">
      RFCXXXX</a>.</p>
  </body>
</html>
<!--

```

[[NOTE TO RFC-EDITOR: Please replace all instances of RFCXXXX
with the number of the published document and remove this note.]]
-->
END

8.3. XML Schema Registration

This section registers an XML schema as per the procedures in [RFC3688].

URI: urn:ietf:params:xml:schema:pidf:geopriv10:relative

Registratant Contact: IETF, GEOPRIV working group
(geopriv@ietf.org), Martin Thomson (martin.thomson@skype.net).

Schema The XML for this schema is found in Section 6 of this document.

8.4. Geopriv Identifiers Registry

This section registers two URNs for use in identifying relative coordinate reference systems. These are added to a new "Geopriv Identifiers" registry according to the procedures in Section 4 of [RFC3553]. The "Geopriv Identifiers" registry is entered under the "Uniform Resource Name (URN) Namespace for IETF Use" category.

Registrations in this registry follow the IETF Review [RFC5226] policy.

Registry name: Geopriv Identifiers

URN Prefix: urn:ietf:params:geopriv:

Specification: RFCXXXX (this document)

Repository: [Editor/IANA note: please include a link to the registry location.]

Index value: Values in this registry are URNs or URN prefixes that start with the prefix "urn:ietf:params:geopriv:". Each is registered independently.

Each registration in the "Geopriv Identifiers" registry requires the following information:

URN The complete URN that is used, or the prefix for that URN.

Description: A summary description for the URN or URN prefix.

Specification: A reference to a specification describing the URN or URN prefix.

Contact: Email for the person or groups making the registration.

Index value: As described in [RFC3553], URN prefixes that are registered include a description of how the URN is constructed. This is not applicable for specific URNs.

The "Geopriv Identifiers" registry has two initial registrations, included in the following sections.

8.4.1. Registration of Two-Dimensional Relative Coordinate Reference System URN

This section registers the "urn:ietf:params:geopriv:relative:2d" URN in the "Geopriv Identifiers" registry.

URN urn:ietf:params:geopriv:relative:2d

Description: A two-dimensional relative coordinate reference system

Specification: RFCXXXX (this document)

Contact: IETF, GEOPRIV working group (geopriv@ietf.org), Martin Thomson (martin.thomson@skype.net).

Index value: N/A.

8.4.2. Registration of Three-Dimensional Relative Coordinate Reference System URN

This section registers the "urn:ietf:params:geopriv:relative:3d" URN in the "Geopriv Identifiers" registry.

URN urn:ietf:params:geopriv:relative:3d

Description: A three-dimensional relative coordinate reference system

Specification: RFCXXXX (this document)

Contact: IETF, GEOPRIV working group (geopriv@ietf.org), Martin Thomson (martin.thomson@skype.net).

Index value: N/A.

9. Acknowledgements

This is the product of a design team on relative location. Besides the authors, this team included: Marc Linsner, James Polk, and James Winterbottom.

10. References

10.1. Normative References

- [RFC1014] Sun Microsystems, Inc., "XDR: External Data Representation standard", RFC 1014, June 1987.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2046] Freed, N. and N. Borenstein, "Multipurpose Internet Mail Extensions (MIME) Part Two: Media Types", RFC 2046, November 1996.
- [RFC2616] Fielding, R., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., and T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1", RFC 2616, June 1999.
- [RFC2818] Rescorla, E., "HTTP Over TLS", RFC 2818, May 2000.
- [RFC3553] Mealling, M., Masinter, L., Hardie, T., and G. Klyne, "An IETF URN Sub-namespace for Registered Protocol Parameters", BCP 73, RFC 3553, June 2003.
- [RFC3688] Mealling, M., "The IETF XML Registry", BCP 81, RFC 3688, January 2004.
- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, January 2005.
- [RFC4119] Peterson, J., "A Presence-based GEOPRIV Location Object Format", RFC 4119, December 2005.
- [RFC4776] Schulzrinne, H., "Dynamic Host Configuration Protocol (DHCPv4 and DHCPv6) Option for Civic Addresses Configuration Information", RFC 4776, November 2006.
- [RFC5139] Thomson, M. and J. Winterbottom, "Revised Civic Location Format for Presence Information Data Format Location Object (PIDF-LO)", RFC 5139, February 2008.

- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.
- [RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, August 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.
- [RFC5962] Schulzrinne, H., Singh, V., Tschofenig, H., and M. Thomson, "Dynamic Extensions to the Presence Information Data Format Location Object (PIDF-LO)", RFC 5962, September 2010.
- [RFC6225] Polk, J., Linsner, M., Thomson, M., and B. Aboba, "Dynamic Host Configuration Protocol Options for Coordinate-Based Location Configuration Information", RFC 6225, July 2011.
- [RFC6848] Winterbottom, J., Thomson, M., Barnes, R., Rosen, B., and R. George, "Specifying Civic Address Extensions in the Presence Information Data Format Location Object (PIDF-LO)", RFC 6848, January 2013.
- [OGC.GML-3.1.1]
Cox, S., Daisey, P., Lake, R., Portele, C., and A. Whiteside, "Geographic information - Geography Markup Language (GML)", OpenGIS 03-105r1, April 2004, <http://portal.opengeospatial.org/files/?artifact_id=4700>.
- [OGC.GeoShape]
Thomson, M. and C. Reed, "GML 3.1.1 PIDF-LO Shape Application Schema for use by the Internet Engineering Task Force (IETF)", OGC Best Practice 06-142r1, Version: 1.0, April 2007.
- [IEEE.754]
IEEE, "IEEE Standard for Binary Floating-Point Arithmetic", IEEE Standard 754-1985, January 2003.
- [Clinger1990]
Clinger, W., "How to Read Floating Point Numbers Accurately", Proceedings of Conference on Programming Language Design and Implementation pp. 92-101, 1990, <<ftp://ftp.ccs.neu.edu/pub/people/will/howtoread.ps>>.

[WGS84] US National Imagery and Mapping Agency, "Department of Defense (DoD) World Geodetic System 1984 (WGS 84), Third Edition ", NIMA TR8350.2, January 2000.

10.2. Informative References

[RFC3863] Sugano, H., Fujimoto, S., Klyne, G., Bateman, A., Carr, W., and J. Peterson, "Presence Information Data Format (PIDF)", RFC 3863, August 2004.

[RFC4479] Rosenberg, J., "A Data Model for Presence", RFC 4479, July 2006.

Authors' Addresses

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

Phone: +1 650-353-1925
EMail: martin.thomson@skype.net

Brian Rosen
Neustar
470 Conrad Dr
Mars, PA 16046
US

EMail: br@brianrosen.net

Dorothy Stanley
Aruba Networks
1322 Crossman Ave
Sunnyvale, CA 94089
US

EMail: dstanley@arubanetworks.com

Gabor Bajko
Nokia
323 Fairchild Drive
Mountain View, CA 94043
US

EMail: gabor.bajko@nokia.com

Allan Thomson
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134
US

EMail: althomso@cisco.com

Network Working Group
Internet Draft
Expires: Mar 4, 2012
Intended Status: Standards Track (PS)

James Polk
Cisco Systems
Brian Rosen
Jon Peterson
NeuStar
Sept 4, 2011

Location Conveyance for the Session Initiation Protocol
draft-ietf-sipcore-location-conveyance-09.txt

Abstract

This document defines an extension to the Session Initiation Protocol (SIP) to convey geographic location information from one SIP entity to another SIP entity. The SIP extension covers end-to-end conveyance as well as location-based routing, where SIP intermediaries make routing decisions based upon the location of the Location Target.

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and BCP 79.

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

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

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

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

This Internet-Draft will expire on Mar 4, 2012.

Copyright Notice

Copyright (c) 2011 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 BSD License.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

- 1. Conventions and Terminology used in this document 3
- 2. Introduction 3
- 3. Overview of SIP Location Conveyance 4
 - 3.1 Location Conveyed by Value 4
 - 3.2 Location Conveyed as a Location URI 5
 - 3.3 Location Conveyed through a SIP Intermediary 5
 - 3.4 SIP Intermediary Replacing Bad Location 7
- 4. SIP Modifications for Geolocation Conveyance 8
 - 4.1 The Geolocation Header Field 8
 - 4.2 The Geolocation-Routing Header Field 10
 - 4.2.1 Explaining Geolocation-Routing header-value States . . 11
 - 4.3 424 (Bad Location Information) Response Code 13
 - 4.4 The Geolocation-Error Header Field 14
 - 4.5 Location URIs in Message Bodies 17
 - 4.6 Location Profile Negotiation 17
- 5. Geolocation Examples 18
 - 5.1 Location-by-value (Coordinate Format) 18
 - 5.2 Two Locations Composed in Same Location Object Example . 20
- 6. Geopriv Privacy Considerations 22
- 7. Security Considerations 22
- 8. IANA Considerations 24
 - 8.1 IANA Registration for New SIP Geolocation Header Field . 24
 - 8.2 IANA Registration for New SIP Geolocation-Routing Header Field 24
 - 8.3 IANA Registration for New SIP Option Tags 25
 - 8.4 IANA Registration for New 424 Response Code 25
 - 8.5 IANA Registration for New SIP Geolocation-Error Header Field 26
 - 8.6 IANA Registration for New SIP Geolocation-Error Codes . . 26
- 9. Acknowledgements 27

10. References	27
10.1 Normative References	27
10.2 Informative References	28
Author Information	29
Appendix A. Requirements for SIP Location Conveyance	29

1. Conventions and Terminology used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119]. This document furthermore uses numerous terms defined in RFC 3693 [RFC3693], including Location Object, Location Recipient, Location Server, Target, Rulemaker and Using Protocol.

2. Introduction

Session Initiation Protocol (SIP) [RFC3261] creates, modifies and terminates multimedia sessions. SIP carries certain information related to a session while establishing or maintaining calls. This document defines how SIP conveys geographic location information of a Target to a Location Recipient (LR). SIP acts as a Using Protocol of location information, as defined in RFC 3693.

In order to convey location information, this document specifies three new SIP header fields, Geolocation, Geolocation-Routing and Geolocation-Error, which carry a reference to a Location Object (LO), grant permission to route a SIP request based on the location-value and provide error notifications specific to location errors respectively. The Location Object (LO) may appear in a MIME body attached to the SIP request, or it may be a remote resource in the network.

A Target is an entity whose location is being conveyed, per RFC 3693. Thus, a Target could be a SIP user agent (UA), some other IP device (a router or a PC) that does not have a SIP stack, a non-IP device (a person or a black phone) or even a non-communications device (a building or store front). In no way does this document assume that the SIP user agent client which sends a request containing a location object is necessarily the Target. The location of a Target conveyed within SIP typically corresponds to that of a device controlled by the Target, for example, a mobile phone, but such devices can be separated from their owners, and moreover, in some cases the user agent may not know its own location.

In the SIP context, a location recipient will most likely be a SIP UA, but due to the mediated nature of SIP architectures, location information conveyed by a single SIP request may have multiple recipients, as any SIP proxy server in the signaling path that inspects the location of the Target must also be considered a Location Recipient. In presence-like architectures, an intermediary

that receives publications of location information and distributes them to watchers acts as a Location Server per RFC 3693. This location conveyance mechanism can also be used to deliver URIs pointing to such Location Servers where prospective Location Recipients can request Location Objects.

3. Overview of SIP Location Conveyance

An operational overview of SIP location conveyance can be shown in 4 basic diagrams, with most applications falling under one of the following basic use cases. Each is separated into its own subsection here in section 3.

Each diagram has Alice and Bob as UAs. Alice is the Target, and Bob is an LR. A SIP intermediary appears in some of the diagrams. Any SIP entity that receives and inspects location information is an LR, therefore in any of the diagrams the SIP intermediary that receives a SIP request is potentially an LR - though that does not mean such an intermediary necessarily has to route the SIP request based on the location information. In some use cases, location information passes through the LS on the right of each diagram.

3.1 Location Conveyed by Value

We start with the simplest diagram of Location Conveyance, Alice to Bob, where no other layer 7 entities are involved.

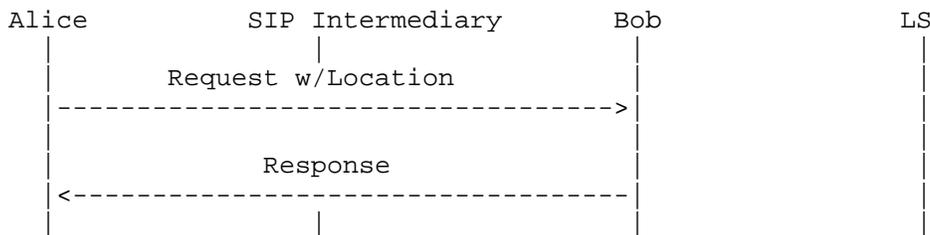


Figure 1. Location Conveyed by Value

In Figure 1, Alice is both the Target and the LS that is conveying her location directly to Bob, who acts as an LR. This conveyance is point-to-point - it does not pass through any SIP-layer intermediary. A Location Object appears by-value in the initial SIP request as a MIME body, and Bob responds to that SIP request as appropriate. There is a 'Bad Location Information' response code introduced within this document to specifically inform Alice if she conveys bad location to Bob (e.g., Bob "cannot parse the location provided", or "there is not enough location information to determine where Alice is").

3.2 Location Conveyed as a Location URI

Here we make Figure 1 a little more complicated by showing a diagram of indirect Location Conveyance from Alice to Bob, where Bob's entity has to retrieve the location object from a 3rd party server.

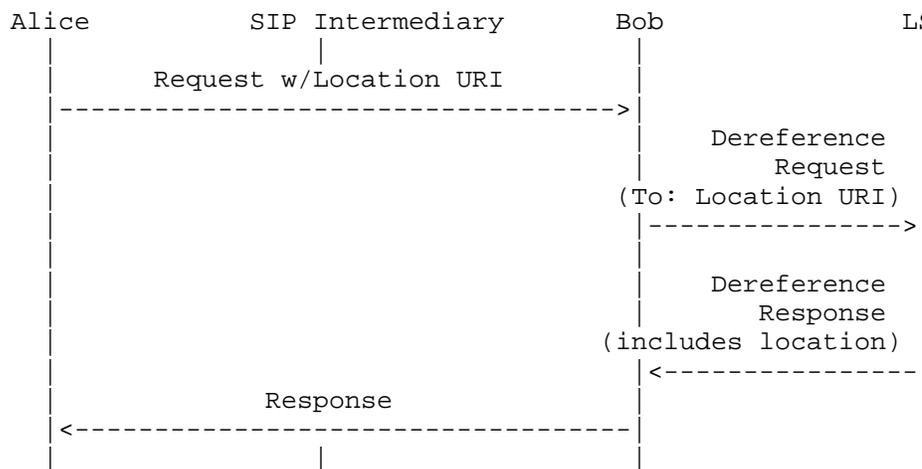


Figure 2. Location Conveyed as a Location URI

In Figure 2, location is conveyed indirectly, via a Location URI carried in the SIP request (more of those details later). If Alice sends Bob this Location URI, Bob will need to dereference the URI - analogous to Content Indirection [RFC4483] - in order to request the location information. In general, the LS provides the location value to Bob instead of Alice directly for conveyance to Bob. From a user interface perspective, Bob the user won't know that this information was gathered from an LS indirectly rather than culled from the SIP request, and practically this does not impact the operation of location-based applications.

The example given in this section is only illustrative, not normative. In particular, applications can choose to dereference a location URI at any time, possibly several times, or potentially not at all. Applications receiving a Location URI in a SIP transaction need to be mindful of timers used by different transactions. In particular, if the means of dereferencing the Location URI might take longer than the SIP transaction timeout (Timer C for INVITE transactions, Timer F for non-INVITE transactions), then it needs to rely on mechanisms other than the transaction's response code to convey location errors, if returning such errors are necessary.

3.3 Location Conveyed though a SIP Intermediary

In Figure 3, we introduce the idea of a SIP intermediary into the example to illustrate the role of proxying in the location architecture. This intermediary can be a SIP proxy or it can be a back-to-back-user-agent (B2BUA). In this message flow, the SIP intermediary could act as a LR, in addition to Bob. The primary use case for intermediaries consuming location information is location-based routing. In this case, the intermediary chooses a next hop for the SIP request by consulting a specialized location service which selects forwarding destinations based on geographical location.

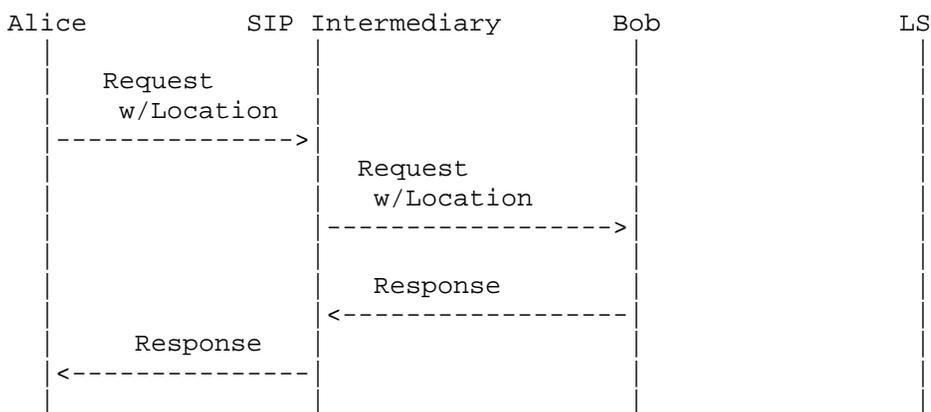


Figure 3. Location Conveyed through a SIP Intermediary

However, the most common case will be one in which the SIP intermediary receives a request with location information (conveyed either by-value or by-reference) and does not know or care about Alice’s location, or support this extension, and merely passes it on to Bob. In this case, the intermediary does not act as a Location Recipient. When the intermediary is not an LR, this use case is the same as the one described in Section 3.1.

Note that an intermediary does not have to perform location-based routing in order to be a Location Recipient. It could be the case that a SIP intermediary which does not perform location-based routing does care when Alice includes her location; for example, it could care that the location information is complete or that it correctly identifies where Alice is. The best example of this is intermediaries that verify location information for emergency calling, but it could also be for any location based routing - e.g., contacting your favorite local pizza delivery service, making sure that organization has Alice’s proper location in the initial SIP request.

There is another scenario in which the SIP intermediary cares about location and is not an LR, one in which the intermediary inserts another location of the Target, Alice in this case, into the request, and forwards it. This secondary insertion is generally not

advisable because downstream SIP entities will not be given any guidance about which location to believe is better, more reliable, less prone to error, more granular, worse than the other location or just plain wrong.

This document takes a "you break it, you bought it" approach to dealing with second locations placed into a SIP request by an intermediary entity. That entity becomes completely responsible for all location within that SIP request (more on this in Section 4).

3.4 SIP Intermediary Replacing Bad Location

If the SIP intermediary rejects the message due to unsuitable location information, the SIP response will indicate there was 'Bad Location Information' in the SIP request, and provide a location specific error code indicating what Alice needs to do to send an acceptable request (see Figure 4 for this scenario).

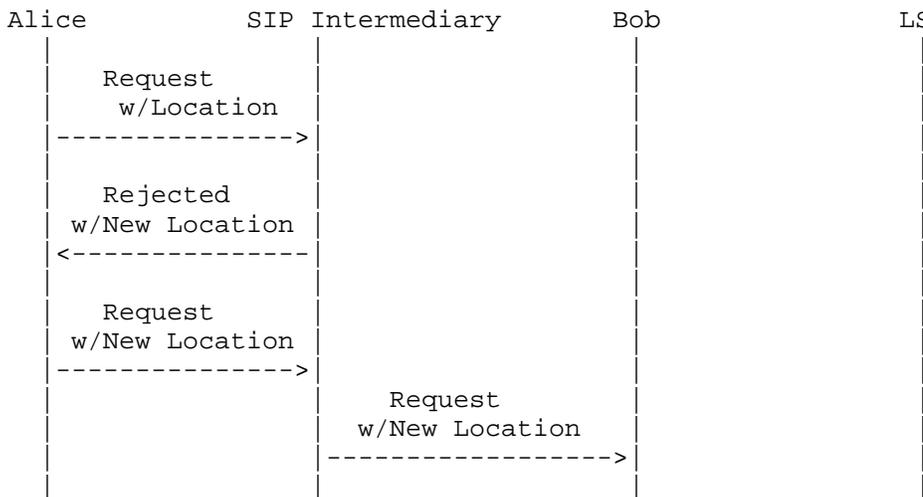


Figure 4. SIP Intermediary Replacing Bad Location

In this last use case, the SIP intermediary wishes to include a Location Object indicating where it understands Alice to be. Thus, it needs to inform her user agent what location it will include in any subsequent SIP request that contains her location. In this case, the intermediary can reject Alice's request and, through the SIP response, convey to her the best way to repair the request in order for the intermediary to accept it.

Overriding location information provided by the user requires a deployment where an intermediary necessarily knows better than an end user - after all, it could be that Alice has an on-board GPS, and the SIP intermediary only knows her nearest cell tower. Which is more accurate location information? Currently, there is no way to

tell which entity is more accurate, or which is wrong - for that matter. This document will not specify how to indicate which location is more accurate than another.

As an aside, it is not envisioned that any SIP-based emergency services request (i.e., IP-911, or 112 type of call attempt) will receive a corrective 'Bad Location Information' response from an intermediary. Most likely, the SIP intermediary would in that scenario act as a B2BUA and insert into the request by-value any appropriate location information for the benefit of Public Safety Answering Point (PSAP) call centers to expedite call reception by the emergency services personnel; thereby, minimizing any delay in call establishment time. The implementation of these specialized deployments is, however, outside the scope of this document.

4. SIP Extensions for Geolocation Conveyance

The following sections detail the extensions to SIP for location conveyance.

4.1 The Geolocation Header Field

This document defines "Geolocation" as a new SIP header field registered by IANA, with the following ABNF [RFC5234]:

```

message-header    /= Geolocation-header ; (message-header from 3261)
Geolocation-header = "Geolocation" HCOLON locationValue
                    *( COMMA locationValue )
locationValue     = LAQUOT locationURI RAQUOT
                    *(SEMI geoloc-param)
locationURI       = sip-URI / sips-URI / pres-URI
                    / http-URI / https-URI
                    / cid-url ; (from RFC 2392)
                    / absoluteURI ; (from RFC 3261)
geoloc-param      = generic-param; (from RFC 3261)

```

HCOLON, COMMA, LAQUOT, RAQUOT, and SEMI are defined in RFC3261 [RFC3261].

sip-URI, sips-URI and absoluteURI are defined according to [RFC3261].

The pres-URI is defined in [RFC3859].

http-URI and https-URI are defined according to [RFC2616] and [RFC2818], respectively.

The cid-url is defined in [RFC2392] to locate message body parts. This URI type is present in a SIP request when location is conveyed as a MIME body in the SIP message.

GEO-URIs [RFC5870] are not appropriate for usage in the SIP

Geolocation header, because it does not include retention and re-transmission flags as part of the location information. Other URI schemes used in the location URI MUST be reviewed against the RFC 3693 [RFC3693] criteria for a Using Protocol. Section 4.6 discusses how URI schemes are communicated using this SIP extension, and what to do if a URI scheme is received that cannot be supported.

The generic-param in the definition of locationValue is included as a mechanism for future extensions that might require parameters. This document defines no parameters for use with locationValue. If a Geolocation header field is received that contains generic-params, each parameter SHOULD be ignored, and SHOULD NOT be removed when forwarding the locationValue. If a need arises to define parameters for use with locationValue, a revision/extension to this document is required.

The Geolocation header field MUST have at least one locationValue. A SIP intermediary SHOULD NOT add location to a SIP request that already contains location. This will quite often lead to confusion within LRs. However, if a SIP intermediary adds location, even if location was not previously present in a SIP request, that SIP intermediary is fully responsible for addressing the concerns of any 424 (Bad Location Information) SIP response it receives about this location addition, and MUST NOT pass on (upstream) the 424 response. A SIP intermediary that adds a locationValue MUST position the new locationValue as the last locationValue within the Geolocation header field of the SIP request.

This document defines the Geolocation header field as valid in the following SIP requests:

INVITE [RFC3261],	REGISTER [RFC3261],
OPTIONS [RFC3261],	BYE [RFC3261],
UPDATE [RFC3311],	INFO [RFC6086],
MESSAGE [RFC3428],	REFER [RFC3515],
SUBSCRIBE [RFC3265],	NOTIFY [RFC3265],
PUBLISH [RFC3903]	

The Geolocation header field MAY be included in any one of the above listed requests by a UA, and a 424 response to any one of the requests sent above. Fully appreciating the caveats/warnings mentioned above, a SIP intermediary MAY add the Geolocation header field.

A SIP intermediary MAY add a Geolocation header field if one is not present - for example, when a user agent does not support the Geolocation mechanism but their outbound proxy does and knows the Target's location, or any of a number of other use cases (see Section 3).

The Geolocation header field MAY be present in a SIP request or response without the presence of a Geolocation-Routing header

(defined in Section 4.2). As stated in Section 4.2, the default value of Geolocation-Routing header-value is "no", meaning SIP intermediaries MUST NOT view (i.e., process, inspect or actively dereference) any direct or indirect location within this SIP message. This is for at least two fundamental reasons,

- 1) to make the possibility of retention of the Target's location moot (because it was not viewed in the first place); and
- 2) to prevent a different treatment of this SIP request based on the contents of the Location Information in the SIP request.

Any locationValue MUST be related to the original Target. This is equally true for the location information in a SIP response, i.e., from a SIP intermediary back to the Target as explained in Section 3.4. SIP intermediaries SHOULD NOT modify or delete any existing locationValue(s). A use-case in which this would not apply would be where the SIP intermediary is an anonymizer. The problem with this scenario is that the geolocation included by the Target then becomes useless for the purpose or service they wanted to use (include) it for. For example, 911/emergency calling or finding the nearest (towing company/pizza delivery/dry cleaning) service(s) will not yield intended results if the Location Information were to be modified or deleted from the SIP request.

4.2 The Geolocation-Routing Header Field

This document defines "Geolocation-Routing" as a new SIP header field registered by IANA, with the following ABNF [RFC5234]:

```
message-header    /= Georouting-header ; (message-header from 3261)
Georouting-header = "Geolocation-Routing" HCOLON
                  ( "yes" / "no" / generic-value )
generic-value     = generic-param; (from RFC 3261)
```

HCOLON is defined in RFC3261 [RFC3261].

The only defined values for the Geolocation-Routing header field are "yes" or "no". When the value is "yes", the locationValue can be used for routing decisions along the downstream signaling path by intermediaries. Values other than "yes" or "no" are permitted for future extensions. Implementations not aware of an extension MUST treat any other received value the same as "no".

If no Geolocation-Routing header field is present in a SIP request, a SIP intermediary MAY insert this header. Without knowledge from a Rulemaker, the SIP intermediary inserting this header-value SHOULD NOT set the value to "yes", as this may be more permissive than the originating party intends. An easy way around this is to have the Target always insert this header-value as "no".

When this Geolocation-Routing header-value is set to "no", this means no locationValue (inserted by the originating UAC or any intermediary along the signaling path) can be used by any SIP intermediary to make routing decisions. Intermediaries that attempt to use the location information for routing purposes in spite of this counter indication could end up routing the request improperly as a result. Section 4.4 describes the details on what a routing intermediary does if it determines it needs to use the location in the SIP request in order to process the message further. The practical implication is that when the Geolocation-Routing header-value is set to "no", if a cid:url is present in the SIP request, intermediaries MUST NOT view the location (because it is not for intermediaries to consider when processing the request), and if a location URI is present, intermediaries MUST NOT dereference it. UAs are allowed to view location in the SIP request even when the Geolocation-Routing header-value is set to "no". An LR MUST by default consider the Geolocation-Routing header-value as set to "no", with no exceptions, unless the header field value is set to "yes".

A Geolocation-Routing header-value that is set to "no" has no special security properties. It is at most a request for behavior within SIP intermediaries. That said, if the Geolocation-Routing header-value is set to "no", SIP intermediaries are still to process the SIP request and send it further downstream within the signaling path if there are no errors present in this SIP request.

The Geolocation-Routing header field satisfies the recommendations made in section 3.5 of RFC 5606 [RFC5606] regarding indication of permission to use location-based routing in SIP.

SIP implementations are advised to pay special attention to the policy elements for location retransmission and retention described in RFC 4119.

The Geolocation-Routing header field cannot appear without a header-value in a SIP request or response (i.e., a null value is not allowed). The absence of a Geolocation-Routing header-value in a SIP request is always the same as the following header field:

```
Geolocation-Routing: no
```

The Geolocation-Routing header field MAY be present without a Geolocation header field in the same SIP request. This concept is further explored in Section 4.2.1.

4.2.1 Explaining Geolocation-Routing header-value States

The Geolocation header field contains a Target's location, and MUST NOT be present if there is no location information in this SIP request. The location information is contained in one or more

locationValues. These locationValues MAY be contained in a single Geolocation header field, or distributed among multiple Geolocation header fields. (See section 7.3.1 of RFC3261.)

The Geolocation-Routing header field indicates whether or not SIP intermediaries can view and then route this SIP request based on the included (directly or indirectly) location information. The Geolocation-Routing header field MUST NOT appear more than once in any SIP request, and MUST NOT lack a header-value. The default or implied policy of a SIP request that does not have a Geolocation-Routing header field is the same as if one were present and the header-value were set to "no".

There are only 3 possible states regarding the Geolocation-Routing header field

- "no"
- "yes"
- no header-field present in this SIP request

The expected results in each state are:

If the Geolocation-Routing -----	Only possible interpretations: -----
"no"	SIP intermediaries MUST NOT process included geolocation information within this SIP request. SIP intermediaries inserting a locationValue into a Geolocation header field (whether adding to an existing header-value or inserting the Geolocation header field for the first time) MUST NOT modify or delete the received "no" header-value.
"yes"	SIP intermediaries can process included geolocation information within this SIP request, and can change the policy to "no" for intermediaries further downstream.
Geolocation-Routing absent	If a Geolocation header field exists (meaning a locationValue is already present), a SIP intermediary MUST interpret the lack of a Geolocation-Routing header field as if there were one present and the header-value is set to "no".

If there is no Geolocation header field in this SIP request, the default Geolocation-Routing is open and can be set by a SIP intermediary or not at all.

4.3 424 (Bad Location Information) Response Code

This SIP extension creates a new location-specific response code, defined as follows,

424 (Bad Location Information)

The 424 (Bad Location Information) response code is a rejection of the request due to its location contents, indicating location information that was malformed or not satisfactory for the recipient's purpose, or could not be dereferenced.

A SIP intermediary can also reject a location it receives from a Target when it understands the Target to be in a different location. The proper handling of this scenario, described in Section 3.4, is for the SIP intermediary to include the proper location in the 424 Response. This SHOULD be included in the response as a MIME message body (i.e., a location value), rather than as a URI; however, in cases where the intermediary is willing to share location with recipients but not with a user agent, a reference might be necessary.

As mentioned in Section 3.4, it might be the case that the intermediary does not want to chance providing less accurate location information than the user agent; thus it will compose its understanding of where the user agent is in a separate <geopriv> element of the same PIDF-LO [RFC4119] message body in the SIP response (which also contains the Target's version of where it is). Therefore, both locations are included - each with different <method> elements. The proper reaction of the user agent is to generate a new SIP request that includes this composed location object, and send it towards the original LR. SIP intermediaries can verify that subsequent requests properly insert the suggested location information before forwarding said requests.

SIP intermediaries that are forwarding (as opposed to generating) a 424 response MUST NOT add, modify, or delete any location appearing in that response. This specifically applies to intermediaries that are between the 424 response generator and the original UAC. Geolocation and Geolocation-Error header fields and PIDF-LO body parts MUST remain unchanged, never added to or deleted.

Section 4.4 describes a Geolocation-Error header field to provide more detail about what was wrong with the location information in the request. This header field MUST be included in the 424 response.

It is only appropriate to generate a 424 response when the responding entity needs a locationValue and there are no values in the request that are usable by the responder, or when the responder has additional location information to provide. The latter case is shown in Figure 4 of section 3.4. There, a SIP intermediary is informing the upstream UA which location to include in the next SIP request.

A 424 MUST NOT be sent in response to a request that lacks a Geolocation header entirely, as the user agent in that case may not support this extension at all. If a SIP intermediary inserted a locationValue into a SIP request where one was not previously present, it MUST take any and all responsibility for the corrective action if it receives a 424 to a SIP request it sent.

A 424 (Bad Location Information) response is a final response within a transaction, and MUST NOT terminate an existing dialog.

4.4 The Geolocation-Error Header Field

As discussed in Section 4.3, more granular error notifications specific to location errors within a received request are required if the location inserting entity is to know what was wrong within the original request. The Geolocation-Error header field is used for this purpose.

The Geolocation-Error header field is used to convey location-specific errors within a response. The Geolocation-Error header field has the following ABNF [RFC5234]:

```

message-header      /= Geolocation-Error
                    ; (message-header from 3261)
Geolocation-Error   = "Geolocation-Error" HCOLON
                    locationErrorValue
locationErrorValue  = location-error-code
                    *(SEMI location-error-params)
location-error-code = 1*3DIGIT
location-error-params = location-error-code-text
                    / generic-param ; from RFC3261
location-error-code-text = "code" EQUAL quoted-string ; from RFC3261

```

HCOLON, SEMI, and EQUAL are defined in RFC3261 [RFC3261]. DIGIT is defined in RFC5234 [RFC5234].

The Geolocation-Error header field MUST contain only one locationErrorValue to indicate what was wrong with the locationValue the Location Recipient determined was bad. The locationErrorValue contains a 3-digit error code indicating what was wrong with the

location in the request. This error code has a corresponding quoted error text string that is human understandable. The text string is OPTIONAL, but RECOMMENDED for human readability, similar to the string phrase used for SIP response codes. That said, the strings are complete enough for rendering to the user, if so desired. The strings in this document are recommendations, and are not standardized - meaning an operator can change the strings - but MUST NOT change the meaning of the error code. Similar to how RFC 3261 specifies, there MUST NOT be more than one string per error code.

The Geolocation-Error header field MAY be included in any response to one of the SIP Methods mentioned in Section 4.1, so long as a locationValue was in the request part of the same transaction. For example, Alice includes her location in an INVITE to Bob. Bob can accept this INVITE, thus creating a dialog, even though his UA determined the location contained in the INVITE was bad. Bob merely includes a Geolocation-Error header value in the 200 OK to the INVITE informing Alice the INVITE was accepted but the location provided was bad.

If, on the other hand, Bob cannot accept Alice's INVITE without a suitable location, a 424 (Bad Location Information) is sent. This message flow is shown in Figures 1, 2 or 3 in Sections 3.1, 3.2 and 3.3 respectively.

If Alice is deliberately leaving location information out of the LO because she does not want Bob to have this additional information, implementations should be aware that Bob could error repeatedly in order to receive more location information about Alice in a subsequent SIP request. Implementations MUST be on guard for this, by not allowing continually more information to be revealed unless it is clear that any LR is permitted by Alice to know all that Alice knows about her location. A limit on the number of such rejections to learn more location information SHOULD be configurable, with a RECOMMENDED maximum of 3 times for each related transaction.

A SIP intermediary that requires Alice's location in order to properly process Alice's INVITE also sends a 424 with a Geolocation-Error code. This message flow is shown in Figure 4 of Section 3.4.

If more than one locationValue is present in a SIP request and at least one locationValue is determined to be valid by the LR, the location in that SIP request MUST be considered good as far as location is concerned, and no Geolocation-Error is to be sent.

Here is an initial list of location based error code ranges for any SIP response, including provisional responses (other than 100 Trying) and the new 424 (Bad Location Information) response. These error codes are divided into 3 categories, based on how the response receiver should react to these errors. There MUST be no more than one Geolocation-Error code in a SIP response, regardless of how many

locationValues there are in the correlating SIP request. There is no guidance given in this document as to which locationValue, when more than one was present in the SIP request, is related to the Geolocation-Error code; meaning that, somehow not defined here, the LR just picks one to error.

- o 1XX errors mean the LR cannot process the location within the request

A non-exclusive list of reasons for returning a 1XX is

- the location was not present or could not be found,
- there was not enough location information to determine where the Target was,
- the location information was corrupted or known to be inaccurate,

- o 2XX errors mean some specific permission is necessary to process the included location information.
- o 3XX errors mean there was trouble dereferencing the Location URI sent.

Dereference attempts to the same request SHOULD be limited to 10 attempts within a few minutes. This number SHOULD be configurable, but result in a Geolocation-Error: 300 error once reached.

It should be noted that for non-INVITE transactions, the SIP response will likely be sent before the dereference response has been received. This document does not alter that SIP protocol reality. This means the receiver of any non-INVITE response to a request containing location SHOULD NOT consider a 200 OK to mean the act of dereferencing has concluded and the dereferencer (i.e., the LR) has successfully received and parsed the PIDF-LO for errors and found none. The end of section 3.2 discusses how transaction timing considerations lead to this requirement.

Additionally, if an LR cannot or chooses not to process location from a SIP request, a 500 (Server Internal Error) SHOULD be used with or without a configurable Retry-After header field. There is no special location error code for what already exists within SIP today.

Within each of these ranges, there is a top level error as follows:

Geolocation-Error: 100 ; code="Cannot Process Location"

Geolocation-Error: 200 ; code="Permission To Use Location
Information"

Geolocation-Error: 300 ; code="Dereference Failure"

If an error recipient cannot process a specific error code (such as the 201 or 202 below), perhaps because it does not understand that specific error code, the error recipient SHOULD process the error code as if it originally were a top level error code where the X in X00 matches the specific error code. If the error recipient cannot process a non-100 error code, for whatever reason, then the error code 100 MUST be processed.

There are two specific Geolocation-Error codes necessary to include in this document, both have to do with permissions necessary to process the SIP request; they are

Geolocation-Error: 201 ; code="Permission To Retransmit Location Information to a Third Party"

This location error is specific to having the Presence Information Data Format (PIDF-LO) [RFC4119] <retransmission-allowed> element set to "no". This location error is stating it requires permission (i.e., PIDF-LO <retransmission-allowed> element set to "yes") to process this SIP request further. If the LS sending the location information does not want to give this permission, it will not change this permission in a new request. If the LS wants this message processed with the <retransmission-allowed> element set to "yes" it MUST choose another logical path (if one exists) for this SIP request.

Geolocation-Error: 202 ; code="Permission to Route based on Location Information"

This location error is specific to having the Geolocation-Routing header value set to "no". This location error is stating it requires permission (i.e., the Geolocation-Routing header value set to "yes") to process this SIP request further. If the LS sending the location information does not want to give this permission, it will not change this permission in a new request. If the LS wants this message processed with the <retransmission-allowed> element set to "yes" it MUST choose another logical path (if one exists) for this SIP request.

4.5 Location URIs in Message Bodies

In the case where an LR sends a 424 response and wishes to communicate suitable location by reference rather than by value, the 424 MUST include a content-indirection body per RFC 4483.

4.6 Location Profile Negotiation

The following is part of the discussion started in Section 3, Figure 2, which introduced the concept of sending location indirectly.

If a location URI is included in a SIP request, the sending user agent MUST also include a Supported header field indicating which location profiles it supports. Two option tags for location profiles are defined by this document: "geolocation-sip" and "geolocation-http". Future specifications MAY define further location profiles per the IANA policy described in Section 8.3.

The "geolocation-sip" option tag signals support for acquiring location information via the presence event package of SIP ([RFC3856]). A location recipient who supports this option can send a SUBSCRIBE request and parse a resulting NOTIFY containing a PIDF-LO object. The URI schemes supported by this option include "sip", "sips" and "pres".

The "geolocation-http" option tag signals support for acquiring location information via an HTTP ([RFC2616]). A location recipient who supports this option can request location with an HTTP GET and parse a resulting 200 response containing a PIDF-LO object. The URI schemes supported by this option include "http" and "https". A failure to parse the 200 response, for whatever reason, will return a "Dereference Failure" indication to the original location sending user agent to inform it that location was not delivered as intended.

If the location URI receiver does not understand the URI scheme sent to it, it will return an Unsupported header value of the option-tag from the SIP request, and include the option-tag of the preferred URI scheme in the response's Supported header field.

See [ID-GEO-FILTERS] or [ID-HELD-DEREF] for more details on dereferencing location information.

5. Geolocation Examples

5.1 Location-by-value (in Coordinate Format)

This example shows an INVITE message with a coordinate location. In this example, the SIP request uses a sips-URI [RFC3261], meaning this message is protected using TLS on a hop-by-hop basis.

```
INVITE sips:bob@biloxi.example.com SIP/2.0
Via: SIP/2.0/TLS pc33.atlanta.example.com;branch=z9hG4bK74bf9
Max-Forwards: 70
To: Bob <sips:bob@biloxi.example.com>
From: Alice <sips:alice@atlanta.example.com>;tag=9fxced76s1
Call-ID: 3848276298220188511@atlanta.example.com
Geolocation: <cid:target123@atlanta.example.com>
Geolocation-Routing: no
Accept: application/sdp, application/pidf+xml
CSeq: 31862 INVITE
Contact: <sips:alice@atlanta.example.com>
Content-Type: multipart/mixed; boundary=boundary1
```

Content-Length: ...

--boundary1

Content-Type: application/sdp

...SDP goes here

--boundary1

Content-Type: application/pidf+xml

Content-ID: <target123@atlanta.example.com>

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

<presence

 xmlns="urn:ietf:params:xml:ns:pidf"

 xmlns:gp="urn:ietf:params:xml:ns:pidf:geopriv10"

 xmlns:gbp="urn:ietf:params:xml:ns:pidf:geopriv10:basicPolicy"

 xmlns:cl="urn:ietf:params:xml:ns:pidf:geopriv10:civicAddr"

 xmlns:gml="http://www.opengis.net/gml"

 xmlns:dm="urn:ietf:params:xml:ns:pidf:data-model"

 entity="pres:alice@atlanta.example.com">

<dm:device id="target123-1">

 <gp:geopriv>

 <gp:location-info>

 <gml:location>

 <gml:Point srsName="urn:ogc:def:crs:EPSG::4326">

 <gml:pos>32.86726 -97.16054</gml:pos>

 </gml:Point>

 </gml:location>

 </gp:location-info>

 <gp:usage-rules>

 <gbp:retransmission-allowed>>false

 </gbp:retransmission-allowed>

 <gbp:retention-expiry>2010-11-14T20:00:00Z

 </gbp:retention-expiry>

 </gp:usage-rules>

 <gp:method>802.11</gp:method>

</gp:geopriv>

<dm:deviceID>mac:1234567890ab</dm:deviceID>

<dm:timestamp>2010-11-04T20:57:29Z</dm:timestamp>

</dm:device>

</presence>

--boundary1--

The Geolocation header field from the above INVITE:

Geolocation: <cid:target123@atlanta.example.com>

... indicates the content-ID location [RFC2392] within the multipart message body of where location information is. The other message body part is SDP. The "cid:" eases message body parsing and disambiguates multiple parts of the same type.

If the Geolocation header field did not contain a "cid:" scheme, for example, it could look like this location URI:

```
Geolocation: <sips:target123@server5.atlanta.example.com>
```

... the existence of a non-"cid:" scheme indicates this is a location URI, to be dereferenced to learn the Target's location. Any node wanting to know where the target is located would subscribe to the SIP presence event package [RFC3856] at

```
sips:target123@server5.atlanta.example.com
```

(see Figure 2 in Section 3.2 for this message flow).

5.2 Two Locations Composed in Same Location Object Example

This example shows the INVITE message after a SIP intermediary rejected the original INVITE (say, the one in section 5.1). This INVITE contains the composed LO sent by the SIP intermediary which includes where the intermediary understands Alice to be. The rules of RFC 5491 [RFC5491] are followed in this construction.

This example is here, but ought not be taken as occurring very often. In fact, this example is believed to be a corner case of location conveyance applicability.

```
INVITE sips:bob@biloxi.example.com SIP/2.0
Via: SIP/2.0/TLS pc33.atlanta.example.com;branch=z9hG4bK74bf0
Max-Forwards: 70
To: Bob <sips:bob@biloxi.example.com>
From: Alice <sips:alice@atlanta.example.com>;tag=9fxced76s1
Call-ID: 3848276298220188512@atlanta.example.com
Geolocation: <cid:target123@atlanta.example.com>
Geolocation-Routing: no
Accept: application/sdp, application/pidf+xml
CSeq: 31863 INVITE
Contact: <sips:alice@atlanta.example.com>
Content-Type: multipart/mixed; boundary=boundary1
Content-Length: ...
```

```
--boundary1
```

```
Content-Type: application/sdp
```

```
...SDP goes here
```

```
--boundary1
```

```
Content-Type: application/pidf+xml
```

```
Content-ID: <target123@atlanta.example.com>
```

```
<?xml version="1.0" encoding="UTF-8"?>
<presence
  xmlns="urn:ietf:params:xml:ns:pidf"
  xmlns:gp="urn:ietf:params:xml:ns:pidf:geopriv10"
  xmlns:gbp="urn:ietf:params:xml:ns:pidf:geopriv10:basicPolicy"
  xmlns:dm="urn:ietf:params:xml:ns:pidf:data-model"
  xmlns:cl="urn:ietf:params:xml:ns:pidf:geopriv10:civicAddr"
  xmlns:gml="http://www.opengis.net/gml"
  entity="pres:alice@atlanta.example.com">
  <dm:device id="target123-1">
    <gp:geopriv>
      <gp:location-info>
        <gml:location>
          <gml:Point srsName="urn:ogc:def:crs:EPSG::4326">
            <gml:pos>32.86726 -97.16054</gml:pos>
          </gml:Point>
        </gml:location>
      </gp:location-info>
      <gp:usage-rules>
        <gbp:retransmission-allowed>>false
        </gbp:retransmission-allowed>
        <gbp:retention-expiry>2010-11-14T20:00:00Z
        </gbp:retention-expiry>
      </gp:usage-rules>
      <gp:method>802.11</gp:method>
    </gp:geopriv>
    <dm:deviceID>mac:1234567890ab</dm:deviceID>
    <dm:timestamp>2010-11-04T20:57:29Z</dm:timestamp>
  </dm:device>
  <dm:person id="target123">
    <gp:geopriv>
      <gp:location-info>
        <cl:civicAddress>
          <cl:country>US</cl:country>
          <cl:A1>Texas</cl:A1>
          <cl:A3>Colleyville</cl:A3>
          <cl:RD>Treemont</cl:RD>
          <cl:STS>Circle</cl:STS>
          <cl:HNO>3913</cl:HNO>
          <cl:FLR>1</cl:FLR>
          <cl:NAM>Haley's Place</cl:NAM>
          <cl:PC>76034</cl:PC>
        </cl:civicAddress>
      </gp:location-info>
      <gp:usage-rules>
        <gbp:retransmission-allowed>>false
        </gbp:retransmission-allowed>
        <gbp:retention-expiry>2010-11-14T20:00:00Z
        </gbp:retention-expiry>
      </gp:usage-rules>
      <gp:method>triangulation</gp:method>
    </gp:geopriv>
  </dm:person>
</presence>
```

```
<dm:timestamp>2010-11-04T12:28:04Z</dm:timestamp>
</dm:person>
</presence>
--boundary1--
```

6. Geopriv Privacy Considerations

Location information is considered by most to be highly sensitive information, requiring protection from eavesdropping and altering in transit. [RFC3693] originally articulated rules to be followed by any protocol wishing to be considered a "Using Protocol", specifying how a transport protocol meets those rules. [RFC6280] updates the guidance in RFC3693 to include subsequently introduced entities and concepts in the geolocation architecture.

RFC5606 explores the difficulties inherent in mapping the GEOPRIV architecture onto SIP elements. In particular, the difficulties of defining and identifying recipients of location information are given in that document, along with guidance in Section 3.3.2 on the use of location by-reference mechanisms to preserve confidentiality of location information from unauthorized recipients.

In a SIP deployment, location information may be added by any of several elements, including the originating user agent or a proxy server. In all cases, the Rule Maker associated with that location information decides which entity adds location information and what access control rules apply. For example, a SIP user agent that does not support the Geolocation header may rely on a proxy server under the direction of the Rule Maker adding a Geolocation header with a reference to location information. The manner in which the Rule Maker operates on these devices is outside the scope of this document.

The manner in which SIP implementations honor the Rule Maker's stipulations for access control rules (including retention and retransmission) is application-specific and not within the scope of SIP protocol operations. Entities in SIP networks that fulfill the architectural roles of the Location Server or Location Recipient treat the privacy rules associated with location information per the guidance in [RFC6280] section 4.2.1. In particular, RFC4119 (especially 2.2.2) gives guidance for handling access control rules; SIP implementations should furthermore consult the emendations in RFC5606.

7. Security Considerations

Conveyance of physical location of a UA raises privacy concerns, and depending on use, there probably will be authentication and integrity concerns. This document calls for conveyance to be accomplished through secure mechanisms, like S/MIME encrypting

message bodies (although this is not widely deployed), TLS protecting the overall signaling or conveyance location by-reference and requiring all entities that dereference location to authenticate themselves. In location-based routing cases, encrypting the location payload with an end-to-end mechanism such as S/MIME is problematic, because one or more proxies on the path need the ability to read the location information to retarget the message to the appropriate new destination UAS. Data can only be encrypted to a particular, anticipated target, and thus if multiple recipients need to inspect a piece of data, and those recipients cannot be predicted by the sender of data, encryption is not a very feasible choice. Securing the location hop-by-hop, using TLS, protects the message from eavesdropping and modification in transit, but exposes the information to all proxies on the path as well as the endpoint. In most cases, the UA has no trust relationship with the proxy or proxies providing location-based routing services, so such end-to-middle solutions might not be appropriate either.

When location information is conveyed by reference, however, one can properly authenticate and authorize each entity that wishes to inspect location information. This does not require that the sender of data anticipate who will receive data, and it does permit multiple entities to receive it securely, but it does not however obviate the need for pre-association between the sender of data and any prospective recipients. Obviously, in some contexts this pre-association cannot be presumed; when it is not, effectively unauthenticated access to location information must be permitted. In this case, choosing pseudo-random URIs for location by-reference, coupled with path encryption like SIPS, can help to ensure that only entities on the SIP signaling path learn the URI, and thus restores rough parity with sending location by-value.

Location information is especially sensitive when the identity of its Target is obvious. Note that there is the ability, according to [RFC3693] to have an anonymous identity for the Target's location. This is accomplished by use of an unlinkable pseudonym in the "entity=" attribute of the <presence> element [RFC4479]. Though, this can be problematic for routing messages based on location (covered in the document above). Moreover, anyone fishing for information would correlate the identity at the SIP layer with that of the location information referenced by SIP signaling.

When a UA inserts location, the UA sets the policy on whether to reveal its location along the signaling path - as discussed in Section 4, as well as flags in the PIDF-LO [RFC4119]. UAC implementations MUST make such capabilities conditional on explicit user permission, and MUST alert the user that location is being conveyed.

This SIP extension offers the default ability to require permission to process location while the SIP request is in transit. The default for this is set to "no". There is an error explicitly

describing how an intermediary asks for permission to view the Target's location, plus a rule stating the user has to be made aware of this permission request.

There is no end-to-end integrity on any locationValue or locationErrorValue header field parameter (or middle-to-end if the value was inserted by a intermediary), so recipients of either header field need to implicitly trust the header field contents, and take whatever precautions each entity deems appropriate given this situation.

8. IANA Considerations

The following are the IANA considerations made by this SIP extension. Modifications and additions to all these registrations require a standards track RFC (Standards Action).

[Editor's Note: RFC-Editor - within the IANA section, please replace "this doc" with the assigned RFC number, if this document reaches publication.]

8.1 IANA Registration for the SIP Geolocation Header Field

The SIP Geolocation Header Field is created by this document, with its definition and rules in Section 4.1 of this document, and should be added to the IANA sip-parameters registry with the following actions

1. Update the Header Fields registry with

Registry:

Header Name	compact	Reference
-----	-----	-----
Geolocation		[this doc]

8.2 IANA Registration for the SIP Geolocation-Routing Header Field

The SIP Geolocation-Routing Header Field is created by this document, with its definition and rules in Section 4.2 of this document, and should be added to the IANA sip-parameters registry with the following action

1. Update the Header Fields registry with

Registry:

Header Name	compact	Reference
-----	-----	-----
Geolocation-Routing		[this doc]

8.3 IANA Registration for Location Profiles

This document defines two new SIP option tags: "geolocation-sip" and "geolocation-http" to be added to the IANA sip-parameters Options Tags registry.

Name	Description	Reference
geolocation-sip	The "geolocation-sip" option tag signals support for acquiring location information via the presence event package of SIP (RFC 3856). A location recipient who supports this option can send a SUBSCRIBE request and parse a resulting NOTIFY containing a PIDF-LO object. The URI schemes supported by this option include "sip", "sips" and "pres".	[this doc]
geolocation-http	The "geolocation-http" option tag signals support for acquiring location information via an HTTP ([RFC2616]). A location recipient who supports this option can request location with an HTTP GET and parse a resulting 200 response containing a PIDF-LO object. The URI schemes supported by this option include "http" and "https".	[this doc]

The names of profiles are SIP option-tags, and the guidance in this document does not supersede the option-tag assignment guidance in [RFC3261] (which requires a Standards Action for the assignment of a new option tag). This document does however stipulate that option-tags included to convey the name of a location profile per this definition MUST begin with the string "geolocation" followed by a dash. All such option tags should describe protocols used to acquire location by reference: these tags have no relevance to location carried in SIP requests by value, which use standard MIME typing and negotiation.

8.4 IANA Registration for 424 Response Code

In the SIP Response Codes registry, the following is added

Reference: RFC-XXXX (i.e., this document)
 Response code: 424 (recommended number to assign)
 Default reason phrase: Bad Location Information

Registry:

Response Code	Reference
Request Failure 4xx	

This SIP Response code is defined in section 4.3 of this document.

8.5 IANA Registration of New Geolocation-Error Header Field

The SIP Geolocation-error header field is created by this document, with its definition and rules in Section 4.4 of this document, to be added to the IANA sip-parameters registry with two actions

1. Update the Header Fields registry with

Registry:

Header Name	compact	Reference
Geolocation-Error		[this doc]

2. In the portion titled "Header Field Parameters and Parameter Values", add

Header Field	Parameter Name	Predefined Values	Reference
Geolocation-Error	code	yes	[this doc]

8.6 IANA Registration for the SIP Geolocation-Error Codes

This document creates a new registry for SIP, called "Geolocation-Error Codes." Geolocation-Error codes provide reason for the error discovered by Location Recipients, categorized by action to be taken by error recipient. The initial values for this registry are shown below.

Registry Name: Geolocation-Error Codes

Reference: [this doc]

Registration Procedures: Specification Required

Code	Default Reason Phrase	Reference
100	"Cannot Process Location"	[this doc]
200	"Permission To Use Location Information"	[this doc]
201	"Permission To Retransmit Location Information to a Third Party"	[this doc]
202	"Permission to Route based on Location Information"	[this doc]
300	"Dereference Failure"	[this doc]

Details of these error codes are in Section 4.4 of this document.

9. Acknowledgements

To Dave Oran for helping to shape this idea.

To Dean Willis for guidance of the effort.

To Allison Mankin, Dick Knight, Hannes Tschofenig, Henning Schulzrinne, James Winterbottom, Jeroen van Bommel, Jean-Francois Mule, Jonathan Rosenberg, Keith Drage, Marc Linsner, Martin Thomson, Mike Hammer, Ted Hardie, Shida Shubert, Umesh Sharma, Richard Barnes, Dan Wing, Matt Lepinski, John Elwell, Thomas Stach, Jacqueline Lee and Adam Roach for constructive feedback and nits checking.

Special thanks to Paul Kyzivat for his help with the ABNF in this document and to Robert Sparks for many helpful comments and the proper construction of the Geolocation-Error header field.

And finally, to Spencer Dawkins for giving this doc a good scrubbing to make it more readable.

10. References

10.1 Normative References

- [RFC3261] J. Rosenberg, H. Schulzrinne, G. Camarillo, A. Johnston, J. Peterson, R. Sparks, M. Handley, and E. Schooler, "SIP: Session Initiation Protocol", RFC 3261, May 2002.
- [RFC4119] J. Peterson, "A Presence-based GEOPRIV Location Object Format", RFC 4119, December 2005
- [RFC2119] S. Bradner, "Key words for use in RFCs to Indicate Requirement Levels", RFC 2119, March 1997
- [RFC2392] E. Levinson, "Content-ID and Message-ID Uniform Resource Locators", RFC 2392, August 1998
- [RFC3856] J. Rosenberg, "A Presence Event Package for the Session Initiation Protocol (SIP)", RFC 3856, August 2004
- [RFC3859] J. Peterson, "Common Profile for Presence (CPP)", RFC 3859, August 2004
- [RFC3428] B. Campbell, Ed., J. Rosenberg, H. Schulzrinne, C. Huitema, D. Gurle, "Session Initiation Protocol (SIP) Extension for

Instant Messaging" , RFC 3428, December 2002

- [RFC3311] J. Rosenberg, "The Session Initiation Protocol (SIP) UPDATE Method", RFC 3311, October 2002
- [RFC3265] Roach, A, "Session Initiation Protocol (SIP)-Specific Event Notification", RFC 3265, June 2002.
- [RFC6086] C. Holmberg, E. Burger, H. Kaplan, "Session Initiation Protocol (SIP) INFO Method and Package Framework", RFC 6086, January 2011
- [RFC3515] R. Sparks, "The Session Initiation Protocol (SIP) Refer Method", RFC 3515, April 2003
- [RFC3903] Niemi, A, "Session Initiation Protocol (SIP) Extension for Event State Publication", RFC 3903, October 2004.
- [RFC5234] Crocker, D., Ed., and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", STD 68, RFC 5234, January 2008.
- [RFC4479] J. Rosenberg, "A Data Model for Presence", RFC 4479, July 2006
- [RFC4483] E. Berger, "A Mechanism for Content Indirection in SIP", RFC 4483, May 2006
- [RFC5491] J. Winterbottom, M. Thomson, H. Tschofenig, "GEOPRIV PIDF-LO Usage Clarification, Considerations, and Recommendations ", RFC 5491, March 2009
- [RFC5870] A. Mayrhofer, C. Spanring, "A Uniform Resource Identifier for Geographic Locations ('geo' URI)", RFC 5870, June 2010
- [RFC2616] R. Fielding, J. Gettys, J., Mogul, H. Frystyk, L., Masinter, P. Leach, T. Berners-Lee, "Hypertext Transfer Protocol - HTTP/1.1", RFC 2616, June 1999

10.2 Informative References

- [RFC3693] J. Cuellar, J. Morris, D. Mulligan, J. Peterson. J. Polk, "Geopriv Requirements", RFC 3693, February 2004
- [RFC2818] E. Rescorla, "HTTP Over TLS", RFC 2818, May 2000
- [RFC5606] J. Peterson, T. Hardie, J. Morris, "Implications of 'retransmission-allowed' for SIP Location Conveyance", RFC5606, Oct 2008
- [ID-GEO-FILTERS] R. Mahy, B. Rosen, H. Tschofenig, "Filtering Location Notifications in SIP", draft-ietf-geopriv-loc-filters, "work

in progress", March 2010

[ID-HELD-DEREF] J. Winterbottom, H. Tschofenig, H. Schulzrinne, M. Thomson, M. Dawson, "A Location Dereferencing Protocol Using HELD", "work in progress", June 2011

[RFC6280] R. Barnes, M. Lepinski, A. Cooper, J. Morris, H. Tschofenig, H. Schulzrinne, "An Architecture for Location and Location Privacy in Internet Applications", draft-ietf-geopriv-arch, "work in progress", October 2010

Authors' Addresses

James Polk
Cisco Systems
3913 Treemont Circle
Colleyville, Texas 76034

33.00111N
96.68142W

Phone: +1-817-271-3552

Email: jmpolk@cisco.com

Brian Rosen
NeuStar, Inc.
470 Conrad Dr.
Mars, PA 16046

40.70497N
80.01252W

Phone: +1 724 382 1051

Email: br@brianrosen.net

Jon Peterson
NeuStar, Inc.

Email: jon.peterson@neustar.biz

Appendix A. Requirements for SIP Location Conveyance

The following subsections address the requirements placed on the UAC, the UAS, as well as SIP proxies when conveying location. This is from the original requirements draft that has since evolved into the solution document (that is above). This has been kept for historical reasons.

If a requirement is not obvious in intent, a motivational statement is included below it.

A.1 Requirements for a UAC Conveying Location

UAC-1 The SIP INVITE Method [RFC3261] must support location conveyance.

UAC-2 The SIP MESSAGE method [RFC3428] must support location conveyance.

UAC-3 SIP Requests within a dialog should support location conveyance.

UAC-4 Other SIP Requests may support location conveyance.

UAC-5 There must be one, mandatory to implement means of transmitting location confidentially.

Motivation: to guarantee interoperability.

UAC-6 It must be possible for a UAC to update location conveyed at any time in a dialog, including during dialog establishment.

Motivation: if a UAC has moved prior to the establishment of a dialog between UAs, the UAC must be able to send location information. If location has been conveyed, and the UA moves, the UAC must be able to update the location previously conveyed to other parties.

UAC-7 The privacy and security rules established within [RFC3693] that would categorize SIP as a 'Using Protocol' MUST be met.

UAC-8 The PIDF-LO [RFC4119] is a mandatory to implement format for location conveyance within SIP.

Motivation: interoperability with other IETF location protocols and Mechanisms.

UAC-9 There must be a mechanism for the UAC to request the UAS send its location.

UAC-9 has been DEPRECATED by the SIP WG, due to the many problems this requirement would have caused if implemented. The solution is for the above UAS to send a new request to the original UAC with the UAS's location.

UAC-10 There must be a mechanism to differentiate the ability of the UAC to convey location from the UACs lack of knowledge of its location

Motivation: Failure to receive location when it is expected can happen because the UAC does not implement this extension, or because the UAC implements the extension, but does not know where the Target is. This may be, for example, due to the failure of the access network to provide a location acquisition mechanism the UAC supports. These cases must be differentiated.

UAC-11 It must be possible to convey location to proxy servers along the path.

Motivation: Location-based routing.

A.2 Requirements for a UAS Receiving Location

The following are the requirements for location conveyance by a UAS:

UAS-1 SIP Responses must support location conveyance.

The SIPCORE WG reached consensus that this be allowed, but not to communicate the UAS's location; rather for a SIP intermediary to inform the UAC which location to include in its next SIP request (as a matter of correcting what was originally sent by the UAC).

UAS-2 There must be a unique 4XX response informing the UAC it did not provide applicable location information.

In addition, requirements UAC-5, 6, 7 and 8 also apply to the UAS.

A.3 Requirements for SIP Proxies and Intermediaries

The following are the requirements for location conveyance by a SIP proxies and intermediaries:

Proxy-1 Proxy servers must be capable of adding a Location header field during processing of SIP requests.

Motivation: Provide network assertion of location when UACs are unable to do so, or when network assertion is more reliable than UAC assertion of location

Note: Because UACs connected to SIP signaling networks can have widely varying access network arrangements, including VPN tunnels and roaming mechanisms, it can be difficult for a network to reliably know the location of the endpoint. Proxies SHOULD NOT assert location of an endpoint unless the SIP signaling network has reliable knowledge of the actual location of the Targets.

Proxy-2 There must be a unique 4XX response informing the UAC it did not provide applicable location information.

GEOPRIV
Internet-Draft
Intended status: Informational
Expires: December 29, 2011

M. Thomson
Andrew Corporation
June 27, 2011

Obscuring Location
draft-thomson-geopriv-location-obscuring-03.txt

Abstract

A method for obscuring location information is described. Both static and changing location information can be obscured. A single distance measure is input to the process; this parameter controls the precision of location information that can be extracted by a recipient.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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

This Internet-Draft will expire on December 29, 2011.

Copyright Notice

Copyright (c) 2011 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	3
2. Method Characteristics and Applicability	3
3. Obscuring Static Locations	4
3.1. Known Point Locations	5
3.2. Known Locations with Uncertainty	5
3.3. Selecting a Offset Vector	6
3.3.1. Angle and Distance Method	6
3.3.2. Square Peg Method	7
3.3.3. Randomness Requirements	8
3.4. Multiple Reported Locations	8
4. Obscuring Changing Locations	9
4.1. Update Conditions	9
4.1.1. Bad Triggers	9
4.1.2. Hidden Trigger	10
4.2. Consecutive Reported Locations	11
4.2.1. Reducing Variation between Offset Vectors	12
4.2.2. Trade-off in Reducing Variation	13
4.3. Returning to the Same Location	14
4.3.1. Positional Stability	14
4.3.2. Triggering with Positional Stability	15
4.3.3. Selecting a Grid	15
4.3.4. Random Grid	16
4.3.5. Linear Interpolation of Random Offsets	17
4.3.5.1. Uniformly Distributed Interpolation	18
4.3.5.2. Applying Uniformly Distributed Interpolation	20
4.3.5.3. Selecting an Appropriate Grid Size	20
4.3.6. The Wonky Grid	21
4.3.6.1. Wonky Grid Points at the Poles	23
4.3.6.2. Interpolation About the 180th Meridian	23
4.3.7. Temporal Interpolation	24
5. Examples	24
6. Acknowledgements	26
7. IANA Considerations	26
8. Security Considerations	26
9. Informative References	27
Appendix A. Sample Implementation	28
Author's Address	32

1. Introduction

A method for obscuring location information is described. This method obscures location information such that it can be provided to recipients without revealing the location of the subject to within the desired distance.

Obscuring location has applications for protecting privacy, as described in [I-D.ietf-geopriv-policy].

This method uses a single configuration parameter as input: an `_obscure_distance_`.

A location recipient (or recipient) is the entity that is given location about a target entity. The goal is to ensure that the recipient is unable to recover location information with better accuracy than is desired. Despite this obscuring the recipient should still be able to use the reported locations.

The obscuring process takes a series of `_known_locations_`, which might have greater accuracy than the recipient is permitted to receive. The obscuring process produces a series of `_reported_locations_`.

2. Method Characteristics and Applicability

The method described here is intended to provide limited protection for location information by constrained degradation. The method has the following characteristics:

Simple Configuration: It might be possible to define a more complete solution for obscuring location information that is more configurable. However, a more configurable option would also demand greater involvement from users so that they would be able to specify a configuration that meets their goals. This method is designed to be easy to understand, which increases the chances that a user is able to successfully choose an appropriate configuration. The method has just one input parameter: the obscuring distance.

A separate parameter for the size of the grid used in the algorithm can affect results; a fixed value is recommended in this document.

Irreversible: Obscuring is intended to be irreversible. Information is lost by applying the process. Multiple applications of this process to the same input location could reduce information more than a single application of the process with the largest obscuring distance.

Increases Uncertainty: A recipient does not need to treat obscured location information any differently to location information that contains uncertainty. The uncertainty of the reported location is increased so that the reported location includes the known location. Thus, the information that is reported is correct, though the accuracy might be reduced. This document relies on a definition of uncertainty for location described in more detail in [I-D.thomson-geopriv-uncertainty].

Two Dimensions: The method described in this document operates in two dimensions only. Many of the principles might be applicable in a higher number of dimensions, though no effort has been made to validate their integrity. A three-dimensional location can be reduced to a two-dimensional form for use in this algorithm. This is not contrary to the goal of reducing the amount of information provided.

Time Invariant: The method described in this document does not use time. An entity performing obscuring does not need to consider time in applying this method. Only the location is protected, not the time that the location was determined. The time from the known location is included in the reported location.

Obscuring Distance Not Secret: No attempt is made to protect the obscuring distance as a secret. It is assumed that a recipient is able to learn this value.

Minimal State: An entity that performs obscuring of locations often performs this service for the combination of many targets and recipients. This process requires only that the obscuring entity hold maintain a trigger location for each recipient. The additional state that an obscuring entity retains in order to apply this obscuring method is a small increment over what is typically required. The current known location does not need to be retained; it need only be reacted to when it changes.

3. Obscuring Static Locations

A static location doesn't change. That is, different locations are not attributed to a single target at different times.

The basic location obscuring case involves a single, isolated instance of location information.

It might be appropriate to apply just this section in protecting the privacy of a single location. A recipient must be unable to acquire multiple location instances for the same entity if this is the only form of obscuring used.

3.1. Known Point Locations

A known point location can be obscured by adding a randomized offset vector to the location. The size of the offset vector is randomly selected so that the reported location could be anywhere within the obscuring distance of the known location, see Section 3.3.

The uncertainty of the reported location is set to the obscuring distance. This ensures that the reported uncertainty region encloses the known location.

Note: It's not sufficient to increase the uncertainty region so that it minimally includes the known location. Doing this reveals that the known location is at the boundary of the reported uncertainty region.

3.2. Known Locations with Uncertainty

A known location with uncertainty is reduced to a circular uncertainty region (see [I-D.thomson-geopriv-uncertainty], Section 4.2). An irregularly shaped uncertainty region is difficult to evaluate against the scalar obscuring distance, and it might inadvertently reveal more information than intended.

A known location with uncertainty greater than the obscuring radius does not require additional obscuring. The radius of the circular uncertainty region is compared to the obscuring distance to determine if further obscuring is necessary. A location with sufficient uncertainty can be directly reported.

Randomization is needed if the known location contains insufficient uncertainty. As for a point location, an offset vector is added and the uncertainty increased to the obscuring distance. A smaller offset vector is necessary where the known location has uncertainty - this vector need only be of a size up to the obscuring distance, less the existing uncertainty.

The reported uncertainty is increased so that the reported location contains an uncertainty radius of at least the obscuring distance. An uncertainty in a known location cannot be recovered by a recipient

of an obscured location unless it is larger than the obscuring distance.

Paradoxically, more accurate location determination methods are better suited to obscuring.

A location that is reported with uncertainty does not always have a uniform probability distribution. A non-uniform distribution is not conducive to obscuring, since a location with an unevenly distributed probability distribution reveals that the location of the target is more likely to be in specific parts of the uncertainty region.

Information on the likely probability distribution cannot be conveyed in many systems, including presence (see [RFC4119], [RFC5491]). The location determination method can be reported, which can reveal characteristics of the probability distribution. Specific measures to counteract this effect are therefore not feasible.

Removing or replacing the location determination method parameter denies a recipient any information about probability distribution.

3.3. Selecting a Offset Vector

There are two methods that can be used to generate a random vector. Both methods produce random vectors that are evenly distributed on the plane within the maximum size.

The angle and distance (polar) method is considerably simpler, but it is less well suited to the complete algorithm. The square peg method is more conducive to the interpolation used.

Both methods take two uniformly distributed random numbers as input.

3.3.1. Angle and Distance Method

In the polar method, the first random value is used to select a random angle, the second to select a random distance.

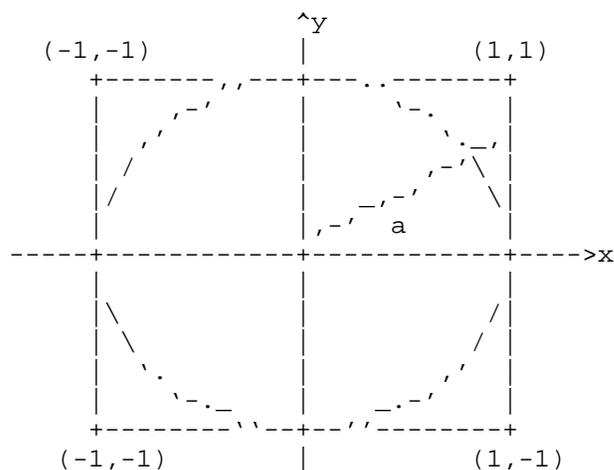
Assuming a "random()" function produces a number distributed between 0 (inclusive) and 1 (exclusive) - that is, the range [0, 1) - the angle and length can be produced by the following:

```
angle = random() * 2 * pi
length = sqrt(random()) * size
or
length = (1 - |random() - random()|) * size
```

...where "sqrt(x)" takes the square root of "x" and "|" takes the absolute value of the enclosed. "size" is the desired size of the random vector, which could be the obscuring distance less any existing uncertainty.

3.3.2. Square Peg Method

In this method, the two random values are used to select a point in a 2x2 square between -1 and +1 on each axis. Vectors that are in the direction of a corner are reduced in length so that the total length of any vector is limited to 1 (as opposed to the square root of 2).



The effect of this is that the probability of finding a value that is toward the corners of the square (angles of $\pi/4$, $3\pi/4$, $-\pi/4$, etc...) is twice the probability of finding a value along the axes (angles of 0, $\pi/2$, π , etc...). This can be corrected by applying the resulting cumulative distribution function to the angle.

```
x      = random() * 2 - 1
y      = random() * 2 - 1
length = sqrt(x*x + y*y)
angle  = atan2(y, x)
a_side = round(angle * 2 / PI);
a_rem  = angle - a_side * PI / 2
length = length * cos(a_rem) * size
angle  = (tan(a_rem) / 8 + a_side / 4) * 2 * PI
```

...where "atan2" produces the angle of the vector, "round(x)" produces the nearest whole number to "x" and the cosine and tangent functions are represented by "cos(x)" and "tan(x)" respectively.

This can be more efficiently calculated without trigonometric functions using:

```
x      = random() * 2 - 1
y      = random() * 2 - 1
length = max(|x|, |y|) * size
if (x == 0) and (y == 0)
    >> return zero length vector
if (|x| > |y|)
    angle = PI * y / x / 4
else
    angle = PI * (2 - x / y) / 4
if (y < -x)
    angle = angle + PI
```

...where "max(x, y)" chooses the more positive value of "x" and "y".

3.3.3. Randomness Requirements

A recipient that is able to learn the state of the random number generator could use this to determine the offset vector. This would reveal the known location based on a given reported location. A secure pseudo-random number generator [RFC4086] provides an assurance that recovering the state of the random number generator is made considerably more difficult.

3.4. Multiple Reported Locations

Multiple applications of this algorithm produce different results. The intersection of multiple reported locations can be used to recover a better estimate of the known location. This recovered estimate has less uncertainty than the obscuring distance, which is not desirable.

Multiple reported locations for the same known location must not be produced. An entity that is responsible for obscuring location might achieve this by storing the reported location with the obscured location.

It is possible to implement obscuring for a static location without retaining state. Seeding a pseudo-random number generator with data that is not available to the recipient can ensure that the same result is produced from the same input. Taking a hash of the known location combined with a secret key ensures that this seed cannot be easily determined by a recipient (see Section 6.2 of [RFC4086] for alternative methods). A hash function that uses the values shown in Section 4.3.4 as input might be sufficient for this task.

4. Obscuring Changing Locations

Applications that use the location of a target over time, such as presence [RFC4079] require additional steps to ensure that the location a recipient acquires does not reveal more information than desired.

The first consideration is the frequency of updates. As the target moves, the known location changes. A frequently updated sequence of reported locations could give a recipient sufficient information to determine the known location with low uncertainty in a fashion close to that described in Section 3.4.

Note: It is not necessary to ensure that a recipient always has accurate location information. Early proposed algorithms wrongly assumed that the reported location was required to cover the known location at all times. Even in the absence of obscuring, changes in location result in a recipient having outdated information. The only necessary constraint is that the location be accurate at the time that it is reported (or the time associated with that report).

4.1. Update Conditions

To limit the amount of information provided to a recipient, new reported locations are not generated in response to all changes in the known location. The trigger for creating a new reported location can be defined.

Any trigger condition needs to be constructed in a way that does not reveal information. At the point that a new reported location is provided to a recipient, the fact that the trigger conditions are met at that point in time provides the recipient with significant information that could - if the trigger conditions were poorly defined - reveal significant information.

The goal is to provide a new reported location when the known location moves by approximately the obscuring distance. This limits the information that a recipient has available with similar accuracy to each individual location.

4.1.1. Bad Triggers

One potential trigger is the movement of the target outside of the reported uncertainty region. At the point that a new reported location is generated, a recipient knows that the target is a) at the boundary of the last uncertainty region, and b) somewhere in the new uncertainty region. The intersection of these two regions produces

an area that is significantly smaller than desired.

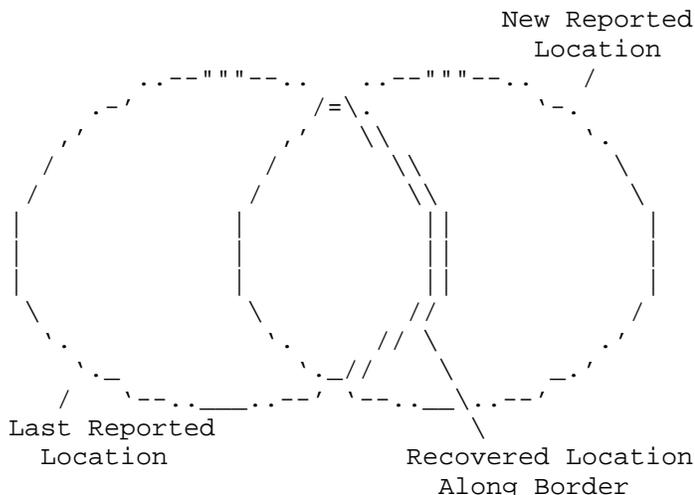


Figure 1: Trigger on Leaving the Reported Location

Similarly, information is revealed if the trigger is movement based on the known location. A new reported location might be produced when the known location moves more than the obscuring distance from the known location from the last report.

That is, when a new location is reported, the corresponding known location is saved. A new reported location is determined when the current known location is more than the obscuring distance from the saved location.

If the recipient is able to assume that the target is moving in a straight line, the speed of the target is revealed.

4.1.2. Hidden Trigger

To limit the information that is revealed at the point that a new reported location is provided, the trigger conditions can be based on information that is not available to the recipient.

Applying randomization to the trigger reduces the ability of a recipient to make assertions about the significance of a new reported location.

A hidden trigger is established using the following process:

- o When a new reported location is generated:
 1. The centroid of the known location is determined.
 2. A random offset vector (Section 3.3) of a maximum size of half the obscuring distance is determined.
 3. The offset vector is added to the centroid and this value is saved as a trigger point.
- o When the known location changes:
 1. The centroid of the (new) known location is determined.
 2. If this centroid is further than the obscuring distance from the saved trigger point, a new reported location is generated.

Each new reported location is randomized using the process described in Section 3.

This algorithm ensures that the centroid of the known location moves between 0.5 and 1.5 times the obscuring distance before a new reported location is produced. As a consequence, the uncertainty in the distance moved is equal to the obscuring distance.

4.2. Consecutive Reported Locations

The obscuring method has a weakness that is as a direct consequence of the triggering conditions. These conditions grant a recipient this information:

For any two consecutive reported locations there is a pair of points that are less than 1.5 times the obscuring distance apart, with one point in the area described by each reported location. The first point is the known location at the time of the first reported location; the second point is the known location at the time of the second reported location.

At the time that a location is reported, the recipient can use this knowledge to determine that the current location of the target is at the intersection of the new reported location and a circle with a radius of 2.5 times the obscuring distance, centered on the last reported location, as shown in Figure 2

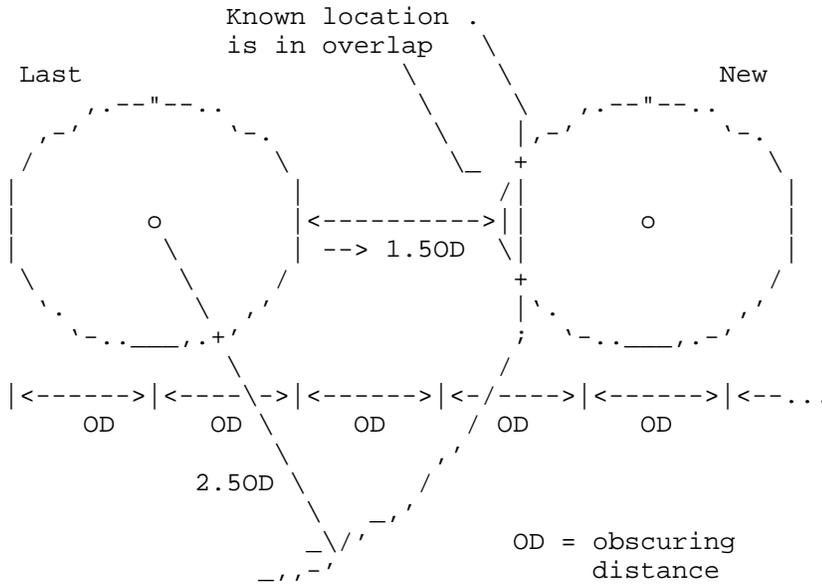


Figure 2: Consecutive Reported Locations

Two consecutive reported locations can have their centers up to 3.5 times the obscuring distance apart; making the closest points on each uncertainty region up to 1.5 times the obscuring distance apart. When consecutive reported locations are maximally distant, a recipient can recover the location of the target almost perfectly.

In general, the known location can be found by taking the intersection of the current reported location and the preceding and antecedent reported locations with their radius increased to 2.5 times the obscuring distance.

This relies on the recipient being able to determine the obscuring distance. As long as the known location has lower uncertainty than the obscuring distance at any point in time, the obscuring distance is trivial to recover.

4.2.1. Reducing Variation between Offset Vectors

This shortcoming can be addressed by reducing the difference between the random offset vector added to consecutive reported locations. The extreme case shown in Figure 2 only arises because the absolute difference between the randomization vector used for in consecutive reported locations is twice the obscuring distance. The problem occurs when the difference between consecutive known locations approaches 1.5 times the obscuring distance in combination with this

large difference between randomization vectors.

Reducing the amount that an offset vector can change between consecutive reported locations ensures that the most extreme configuration cannot occur.

Using the same offset vector for all reported locations removes the problem entirely. However, using the same offset vector increases the probability of that vector being discovered and results in a serious problem if it is discovered. An adversary need only discover a single known location for a specific reported location. For instance, if the target is following a road, reported locations that have a fixed offset from the known location will reveal the shape of the road. From this it is trivial to learn the offset vector and hence all past and future locations can be recovered.

Averaging is one potential approach to this problem. Each time a location is randomized, the offset vector used can be the average of a new random offset vector and the offset vector that was last used. The proportion of old and new vectors determines the trade-off between the probability that a recipient is able to learn a more accurate location with the probability that a recipient is able to learn the offset.

4.2.2. Trade-off in Reducing Variation

It is more difficult to learn an offset vector if additional randomness is added to each new vector. An adversary that learns a known location immediately has less information about subsequent known locations based on the amount of additional randomness. As long as the offset vector is able to change significantly as several locations are reported, learning a limited number of offset vectors is of limited use in recovering future known locations.

Too large a change in the offset vector increases the chances of revealing the known location to a small area. Too small a change provides an adversary that discovers a known location information more information about subsequent known locations. A trade-off is necessary.

The only way that the known location can be guaranteed to be unknown over the entire area is when the offset vector doesn't change at all. If the absolute difference in offset vectors is half the obscuring distance, in the worst case the recipient is able to determine the known location to be within 77 percent of the desired area. This varies based on " $\cos(\text{diff})$ ", as follows:

```
diff    = | offset[x] - offset[x - 1] | / obscuring distance
a(diff) = ((1.5 + diff)^2 - 5.25) / (2*(1.5 + diff))
o(diff) = acos(a(diff)) + 6.25 * acos((1.5 + diff - a(diff)) / 2.5)
          - (1.5 + diff) * sqrt(1 - a(diff)^2)
```

...where "acos(x)" returns the inverse cosine of "x". This only produces a result where "diff" is less than 2.

It might be useful in this case to create a offset vector that is no more than "diff" times the obscuring distance different to the previous vector. This might be done by taking a weighted average of the previous vector with a new random offset vector as follows:

```
o[new] = (o[prev] * (2 - diff) + o[random] * diff) / 2
```

...where "o[new]" is the new offset vector, "o[prev]" is the previous vector, and "o[random]" is a completely random vector of the same magnitude.

4.3. Returning to the Same Location

A moving target might return to the same location several times. The method described thus far produces a different reported location each time. A recipient that is able to observe location over time could intersect reported locations to recover the known location as long as they make the assumption that the known location is the same each time.

This can be extended to reveal a path that is habitually followed in the same way. Each time the path is travelled, changing offset vectors eventually reveal a more accurate view of the path.

4.3.1. Positional Stability

The key to addressing this flaw is to have the randomization of offset vectors based on the known location. If the same known location produced a reported location that was equal or very close to it each time that the location was obscured, this would address the problem.

It might be possible to take the coordinates of the known location and pass them - along with a secret key - through a cryptographic hash function. The resulting bits could be used as randomness that produces an offset vector. This would ensure that the exact same location always produces the same random vector.

The drawback of this sort of method is that the location is obscured inconsistently when the known location changes even slightly. Such

imprecision is commonplace in location determination methods, rendering this approach unsuitable.

The goal is to ensure that two known locations in close proximity produce a constant (or near almost constant) random offset vector. It is also desirable that the random vector change as the locations change. This has the consequence of reducing the difference in randomness between consecutive reported locations, provided that the random values do not vary significantly over shorter distances (see Section 4.2.1). The offset vector needs to change over a longer distance to limit the amount that an adversary benefits from learning both known and reported locations.

An approach similar to that described in [PERLIN] is used to achieve a continuously varying random field. In this, randomness is constrained to a grid of points with interpolation used to determine values for intervening points.

4.3.2. Triggering with Positional Stability

No specific changes are required for the triggering process, though this does require that some state be maintained by the entity that performs obscuring. For a SIP entity that is maintaining a subscription, this is not expected to be onerous.

The advantage of having a specific trigger for providing a new reported location is that it reduces the information provided to a recipient. Providing updates at a higher rate provide a recipient with additional information that could be used to recover the offset.

4.3.3. Selecting a Grid

In selecting an appropriate grid with two dimensions, the curvature of the surface of the Earth presents a challenge. The simplest approach might be to select an origin at latitude 0, longitude 0. Grid points could be placed at increments based on a constant ratio between latitude and longitude and distance; for example, $9e-6$ degrees per meter assumes a spherical planet of 6366197 meter radius, which is slightly smaller than the semi-major axis of the ellipsoid used in most Earth models.

For a two-dimensional grid with a multiple of "m", the following equations identify the latitude and longitude of the four nearest grid points to a given location:

```
grid = m * obscuring distance * 9e-6

latitude[low] = floor(latitude / grid) * grid
latitude[high] = latitude[low] + grid

longitude[low] = floor(longitude / grid) * grid
longitude[high] = longitude[low] + grid
```

...where "floor(x)" produces the nearest whole integer that is more negative than "x".

Grid intervals can be set to a multiple of the obscuring distance that ensures that consecutive reported locations have continuously varying offset vectors. These vectors need to change at a rate that ensures maximum change over multiple reported locations without causing too much information to be revealed from two consecutive locations (as described in Section 4.2). Selecting a grid size is discussed in more detail in Section 4.3.5.3.

The shortcoming of a grid of this nature is that changes in longitude are more rapid as locations get closer to the poles. At approximately 60 degrees of latitude (North or South), grid intervals on the East-West direction are twice as frequent as desired. For this reason, larger intervals between grid points might be chosen for longitudes.

A solution for this problem is described in Section 4.3.6. An alternative solution might use a local tangent plane, though this introduces the problem of selecting an appropriate tangent plane as locations change and providing consistent transitions between different tangent planes.

In three dimensions, conversion to Earth-centered, Earth-fixed Cartesian coordinates renders this problem moot.

4.3.4. Random Grid

At each of the points on the grid, a random offset vector is produced using the method described in Section 3.3.2. Interpolation is used to produce the offset vector for points within each grid cell, as shown in Figure 3.

Rather than use a random number generator, random numbers are produced using a cryptographic hash function. The input to this hash might include:

- o a secret known only by the entity that performs the obscuring with sufficient entropy to render guessing ineffective (a random

- sequence [RFC4086] is suitable for this purpose),
- o the identity of the target,
 - o each individual coordinate of the grid point, and
 - o as necessary, a counter that allows for multiple random values to be generated (for angle and distance, x and y, depending on the method used to generate the random offset vector).

The inclusion of a secret ensures that a recipient is unable to construct the offset vector. This secret is persistent so that later applications of the obscuring formula do not produce a different offset vector for the same location.

Section 3.3 requires that multiple random numbers are produced. The additional identifier produces additional randomness where multiple random (or pseudo-random) numbers are required.

Using a hash in this fashion ensures that each target gets a different set of random offset vectors and that the same grid point coordinates produce the same result.

Though ordering need only be consistent between consequent applications of the obscuring algorithm, the following might be used to produce random bits:

```
random = HMAC(secret key, target identity | identifier
                | coordinate | coordinate | ...)
```

...where "HMAC" is the hash MAC function [RFC2104] and "|" represents concatenation, which might require a delimiter to terminate variable length values.

Alternatively, the same sequence could be used to seed a secure pseudo-random number generator [RFC4086]. Extracting values in the same order makes the "identifier" unnecessary.

One consequence of this approach is that changes to the obscuring distance result in the noise pattern being completely changed. This can result in the same known location producing a significantly different reported location before and after the change.

4.3.5. Linear Interpolation of Random Offsets

Once a grid of random offset vectors is established, an offset vector is calculated based on the centroid of the known location. Figure 3 shows a centroid at the point "(x,y)" and the values that are used in

the interpolation process.

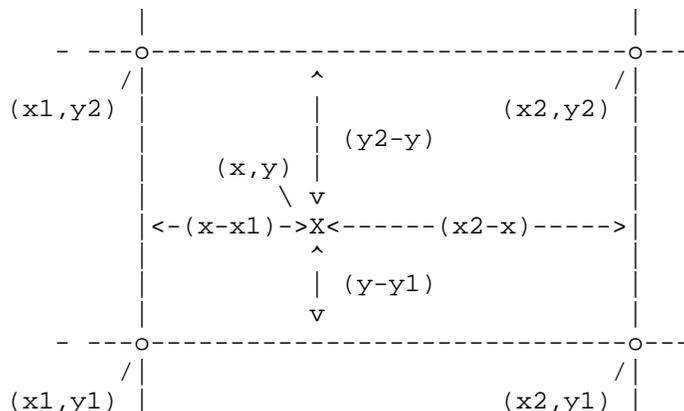


Figure 3: Grid Interpolation

The offset vector at the identified point is produced by taking the weighted average of the offset vectors. Two weighted averages are taken between pairs of adjacent grid points along the same axis, then the weighted average of the two resulting vectors is taken along the other axis.

The following equations produce an linearly interpolated offset vector for any point in this grid cell:

$$\begin{aligned}
 tx &= (x - x1) / (x2 - x1) \\
 ty &= (y - y1) / (y2 - y1) \\
 w1 &= o[x1,y1] * (1 - tx) + o[x2,y1] * tx \\
 w2 &= o[x1,y1] * (1 - ty) + o[x2,y1] * ty \\
 \text{offset} &= w1 * (1 - tx) + w2 * ty
 \end{aligned}$$

...where "o[x1,y1]" is the random offset vector at the grid point "(x1,y1)".

4.3.5.1. Uniformly Distributed Interpolation

A consequence of performing a weighted average is that the resulting value is not uniformly distributed. Depending on the weighting factor (the value "tx" or "ty" in Section 4.3.5), the resulting probability distribution has a higher probability of producing values in the middle of the range of possible values.

For example, the probability distribution for a weighted average of two uniformly distributed random numbers between 0 and 1 is shown in Figure 4. The figure shows the case where "t" is less than 0.5,

though the same distribution is produced for "t" and "(1-t)".

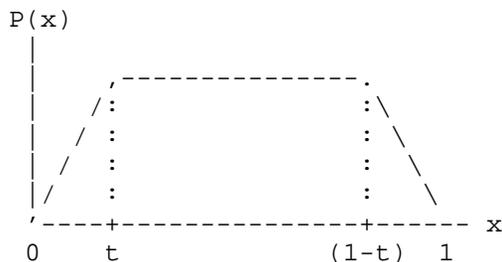


Figure 4

In order to correct for this skewing of results toward the middle of the range, a smoothed interpolation is used.

Over the range from 0 to 1, the following produces a uniformly distributed interpolation between "a" and "b":

```

r = a * (1 - t) + b * t
IF r < t AND r < (1 - t) THEN:
  r = r * r / 2 / t / (1 - t)
ELSE IF r > t AND r > (1 - t) THEN:
  r = 1 - (1 - r) * (1 - r) / t / (1 - t)
ELSE IF t < 0.5 THEN:
  r = (2 * r - t) / 2 / (1 - t)
ELSE:
  r = (2 * r - 1 + t) / 2 / t

```

This maps a linearly interpolated value to a smoothed value, using the cumulative distribution function for the weighted sum of "a" and "b". This mapping produces a value between 0 and 1 for inputs between 0 and 1. The mapping is continuous. The mapping is not monotonically increasing for some values of "a" and "b"; the intent is to have a uniform distribution between 0 and 1, not between "a" and "b".

For convenience, this interpolation function is represented in shorthand throughout the remainder of the document:
"uniformDistInterp(a, b, t)".

Uniform interpolation alters the rate of change of the output. For a proportional movement in "t" of "dt", the absolute change in output is at most:

$$dr = 1 - (1 - dt)^2$$

Toward the middle of the range, for values of "a" and "b" that are at the extents of the possible range and small values of "dt", changes are magnified by up to two times their magnitude.

This interpolation function has similar characteristics to the smoothing function used in [PERLIN], except that the goal is not smoothing, but ensuring a uniform distribution of values in the output. Values are continuous, but their first derivative is not.

4.3.5.2. Applying Uniformly Distributed Interpolation

The methods for producing random vectors described in Section 3.3 produce a result that is uniformly distributed in a circular area. As a result, the cartesian coordinates produced are not evenly distributed on each axis. Similarly, the polar coordinates have a non-uniformly distributed magnitude. Rather than interpolate on the output of this process, the uniformly distributed interpolation is applied to the random inputs.

Interpolation is performed on a set of random numbers that are produced at each grid vertex. This is used to produce a single set of random numbers that are used as input to the random vector algorithm.

A consequence of this process with the simple polar method described in Section 3.3.1 is that the angle of the random vector does not cross 360 degrees (2π) when being interpolated. In the worst case, interpolation between two points requires rotation through almost 360 degrees.

The alternative method of interpolating angles - linear interpolation using the shortest path - does produce an uniformly distributed output, but it also produces a discontinuity that could be exploited by a recipient when interpolation is applied in more than one dimension. It is possible to produce a change in the offset vector of up to twice the obscuring distance in size as the known location moves only a short distance.

The more complicated square peg method (Section 3.3.2) results produces evenly distributed values without this problem.

4.3.5.3. Selecting an Appropriate Grid Size

In the worst case, the polar method of generating a random vector in combination with uniformly distributed interpolation can result in twice the rate of rotation. Interpolation through a complete 360 degrees results in a maximum absolute change of:

$$d[p] = 2 * \sin(\pi * dr))$$

...where "dr" is the distance moved as a proportion of the obscuring distance, which is no more than 0.5.

Using the maximum value from Section 4.3.5.1, the number of multiples required to limit movement can be calculated using:

$$m[p] = 1.5 / (1 - \sqrt{1 - \sin(d[p] / 2) / \pi})$$

For an absolute change in the random vector of no more than the obscuring distance, the grid needs to be at least 17.22 multiples of the obscuring distance. If the absolute change is only half this amount, the grid needs to be larger, at 36.53 multiples of the obscuring distance.

Using such a large grid to deal with a low probability case is suboptimal. The square peg method allows for a much smaller grid, with a maximum absolute change being dependent on only the increased rate of change produced by the interpolation method:

$$\begin{aligned} d[sp] &= 2 * dr \\ &= 2 * (1 - (1 - 1.5 / m[sp])^2) \\ m[sp] &= 1.5 / (1 - \sqrt{1 - d[sp] / 2}) \end{aligned}$$

This means that a grid of 5.12 times the obscuring distance limits absolute difference in the offset vector to obscuring distance; a grid of 11.20 times the obscuring distances limits the difference to half.

Selecting a grid size at 8 times the obscuring distances ensures that the absolute change in offset vector is 0.680 times the obscuring distance. A complete change in offset vector can then occur after linear movement of only 8 times the obscuring distance. In the worst case, movement reveals a location within 66.0% of the area of a circle with a radius of the obscuring distance.

4.3.6. The Wonky Grid

To address the concerns caused by the curvature of the Earth, a modified grid-like structure can be used. It is not strictly necessary that the grid be absolutely grid-like in structure. Therefore, it's possible that different grid intervals could be selected.

This structure uses a different interval for points at different latitudes, at the selected low latitude:

```

grid[l1at] = grid / cos(latitude[low])
longitude[low,l1at] = floor(longitude / grid[l1at]) * grid[l1at]
longitude[high,l1at] = longitude[low,l1at] + grid[l1at]

```

...and at the high latitude:

```

grid[hlat] = grid / cos(latitude[high])
longitude[low,hlat] = floor(longitude / grid[hlat]) * grid[hlat]
longitude[high,hlat] = longitude[low,hlat] + grid[hlat]

```

...where "cos(x)" produces the cosine of "x".

This produces fewer grid points for latitudes that are further from the Equator. At the poles (and above), a single offset vector is sufficient.

Interpolation of these points uses four distinct points, as shown in Figure 5.

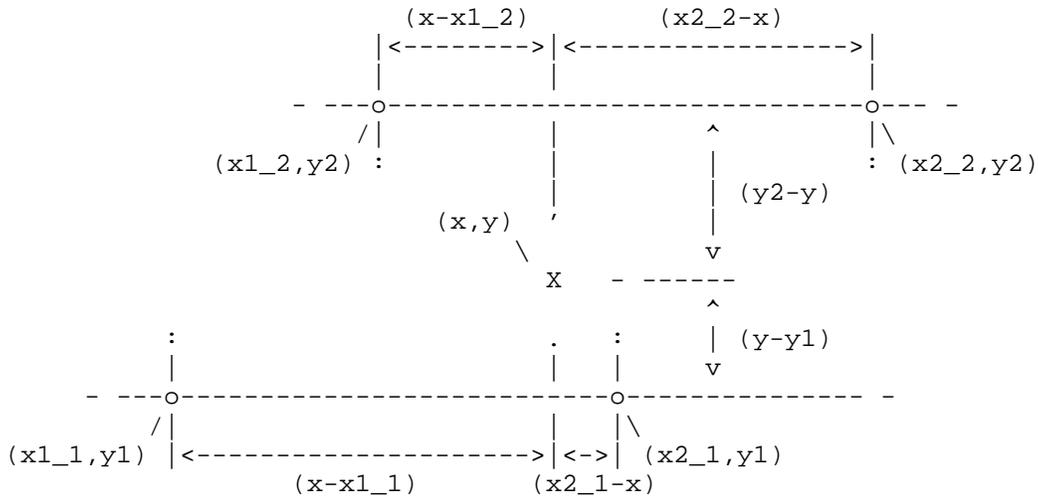


Figure 5: Wonky Grid Interpolation

Linear interpolation uses the amended equations:

```

tx_1 = (x - x1_1) / (x2_1 - x1_1)
w1 = uniformDistInterp(r[x1_1,y1], r[x2_1,y1], tx_1)
tx_2 = (x - x1_2) / (x2_2 - x1_2)
w2 = uniformDistInterp(r[x1_2,y2], r[x2_2,y2], tx_2)

```

Note that this uses the uniformly distributed random values selected at each grid point, rather than the offset vectors. Each random

value is a uniformly distributed random value in the range [0, 1).

4.3.6.1. Wonky Grid Points at the Poles

At 90 degrees North and South, the cosine used to determine the wonky grid produces a zero. This produces an undefined grid spacing.

To avoid this problem, produce a single value at each pole: (90, 0) and (-90, 0). This value replaces "w1" or "w2" in the interpolation equations. Retaining the same weighting (that is, "ty") for determining the final offset is desirable, so that the rate of change is not artificially increased.

4.3.6.2. Interpolation About the 180th Meridian

At 180 degrees East (or West), longitude values cross from positive to negative values. This produces a discontinuity in the values used. This could be exploited to learn when the known location cross the 180th meridian.

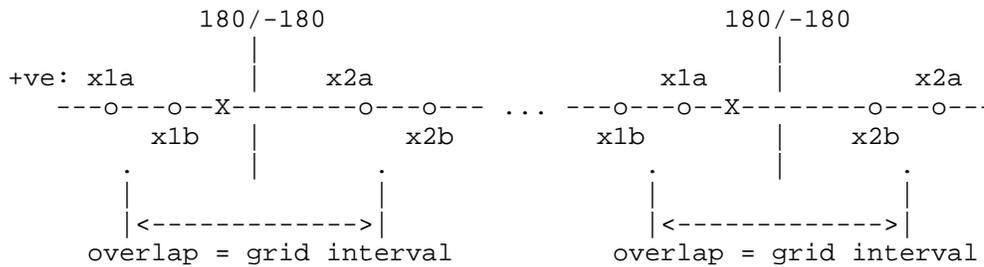


Figure 6: Interpolation About 180 Degrees

This problem might only manifest for one of the two interpolations performed across changing longitude values in a wonky grid. To address this, the values produced by the negative and positive aspects are independently generated, then these values are interpolated over a span of one grid interval.

For any point within half of one grid interval from the 180th meridian, this algorithm is used. Perform interpolation using the selected grid points, then add or subtract 360 degrees from the original value to get a value that is either more than 180 degrees or less than -180 degrees. Perform interpolation on this second point.

The two interpolated values are then interpolated using a different proportion. This interpolation is taken on the overlap interval that crosses the 180th meridian, as shown in the Figure 6. This proportion is produced by taking the positive input value (that is,

the longitude value, with 360 degrees added if it is negative) and applying the following:

```
grid = m * obscuring distance * 9e-6 / cos(latitude)
IF longitude + grid / 2 > 180 OR longitude - grid / 2 < -180 THEN:
  t = ((longitude + 360) % 360 - 180 - grid / 2) / grid
  random[o] = uniformDistInterp(random[+ve], random[-ve], t)
ENDIF
```

...where "%" represents the modulo operation. The final interpolated value is determined using the uniformly distributed weighted average method described in Section 4.3.5.1.

4.3.7. Temporal Interpolation

Providing different values over time is difficult to balance against the need to obscure the same location in the same way. It is possible to add additional dimensions upon which to interpolate the offset vector. Adding time as one such dimension would allow the offset vector to change gradually over time as well as with respect to space.

A form of temporal interpolation might allow the obscuring entity to change the secret key that it maintains over time. However, this does not provide positional stability unless the interpolation is performed over a period that is significantly longer than the period over which the known location might return to the same location. Changing the offset vector applied to the same location would negate much of the benefit derived from the algorithm.

In practice, the period over which the offset would change would have to be significantly longer than the time taken for all potential visited locations to completely change in all aspects. This implies that temporal interpolation is likely only useful on geological time scales.

5. Examples

Obscuring a known location at latitude -34.401072, longitude 150.636361 with 100 meter obscuring distance first requires calculation of the grid size and the grid points:

```
gridsize = 8 * obscuring_distance * 9e-6 = 0.0072
```

Once that is determined, the two latitude values used for the grid are determined:

```
lowlat = floor(-34.401072 / 0.0072) * 0.0072 = -34.4016
highlat = lowlat + 0.0072 = -34.3944
```

For each latitude value, two longitude values are determined using a modified grid size to find the final set of of grid points:

Note: Intermediate values in this example are rounded for presentation purposes.

```
grid[lowlat] = gridsize / cos(lowlat * pi / 180) = 0.0087262
lowlng[lowlat] = floor(150.636361 / 0.0087262) * 0.0087262
                = 150.632339
highlng[lowlat] = lowlng[lowlat] + 0.0087262 = 150.641066
grid[highlat] = 0.0087255
lowlng[highlat] = 150.628105
highlng[highlat] = 150.636831
```

This gives a set of points for which random values are produced. The actual random values used depend on many factors (see Section 4.3.4). The following values are used in this example:

```
random[-34.4016, 150.632339] = 0.4228538586758077
random[-34.4016, 150.641066] = 0.9430289615411311
random[-34.3944, 150.628105] = 0.9174296103883535
random[-34.3944, 150.636831] = 0.008725488356129405
```

The random values are interpolated along the same latitude using a "t" value that is based on the distance from the corresponding low longitude value. The two resulting values are interpolated along the same longitude using a "t" value that is based on the distance from the low latitude value. Uniformly distributed interpolation is used in both cases.

```
t[-34.4016] = (150.636361 - 150.632339) / 0.0087262 = 0.460866
r[-34.4016] = 0.770898
t[-34.3944] = (150.636361 - 150.628105) / 0.0087255 = 0.946145
r[-34.3944] = 0.440578
t = (-34.401072 - -34.4016) / 0.0072 = 0.0733055
r = 0.7661978449732944
```

This first random value is used for the "x" component. A second random value for "y" is chosen using the same process, producing 0.16585607985072537.

These values are then input into the square peg algorithm:

```
d = 100 * max(0.7661978449732944, 0.16585607985072537)
  = 76.61978449732944
-- since |x| > |y|
a = y * pi / x / 4 = 5.3380813420741795
-- no further change since y > -x
```

Therefore, the location is moved 76.62 meters on a bearing of 305.84 degrees. The resulting reported location is moved along the local tangent plane to $\{-34.400719, 150.635772\}$ and a circle of 100 meter radius is described.

Finally, a random point is chosen within 50 meters of the original point. No more location is provided until the known location moves more than 100 meters from that point. In this case, the trigger point is set to $\{-34.401388, 150.636471\}$. If the known location is updated to $\{-34.401816, 150.636361\}$, no new location is reported.

Moving to $\{-34.400621, 150.635717\}$ is more than 100 meters from the trigger point, even though this is very close to the last reported location. This results in a new location being reported at $\{-34.400346, 150.634929\}$.

6. Acknowledgements

Thanks go to Robert Sparks for identifying key shortcomings in early attempts to obscure location. Richard Barnes, Jorge Cuellar, Cullen Jennings, Warren Kumari, and Hannes Tschofenig variously provided input, feedback, criticisms and insightful ideas.

7. IANA Considerations

This document has no IANA actions.

[RFC Editor: please remove this section prior to publication.]

8. Security Considerations

This document describes a method for obscuring location. An effort has been made to ensure that reported locations do not reveal any more information than the input dictates. However, obscuring location is not a substitute for withholding location information if the goal is to ensure that a recipient remains ignorant of the known location. Alternatively, a recipient might be provided with completely falsified location information.

There is little point in obscuring location when other location-related information is included in a composite document, like a presence document [RFC3863]. Removing other information, such as dynamic location information [RFC5965] is necessary to ensure that this cannot be used to recover the known location.

A reported location can inadvertently reveal far more information than intended to a recipient in possession of additional information. A recipient might be able to apply this additional information to determine the location of the target with less uncertainty than desired. Additional information includes information about the reported location or information about the Target.

For instance, a recipient with a map might be able to identify areas on that map that a target is more likely to be found. A recipient can combine any additional information with the knowledge that the reported location is correct at the time it is reported to recover a better estimate of the known location. Aside from map-based data, other information that could be used to acquire a more accurate estimate of the location of a target might include knowledge of the target's past behavior, personality traits, or aggregated demographic data.

Increasing the obscuring distance might increase the uncertainty in the location that a recipient with additional information can ultimately recover. The complexity involved and the large volume of additional data involved makes more specific measures difficult.

9. Informative References

[I-D.ietf-geopriv-arch]

Barnes, R., Lepinski, M., Cooper, A., Morris, J.,
Tschofenig, H., and H. Schulzrinne, "An Architecture for
Location and Location Privacy in Internet Applications",
draft-ietf-geopriv-arch-03 (work in progress),
October 2010.

[I-D.ietf-geopriv-policy]

Schulzrinne, H., Tschofenig, H., Morris, J., Cuellar, J.,
and J. Polk, "Geolocation Policy: A Document Format for
Expressing Privacy Preferences for Location Information",
draft-ietf-geopriv-policy-23 (work in progress),
March 2011.

[I-D.thomson-geopriv-uncertainty]

Thomson, M. and J. Winterbottom, "Representation of
Uncertainty and Confidence in PIDF-LO",

- draft-thomson-geopriv-uncertainty-06 (work in progress),
March 2011.
- [PERLIN] Perlin, K., "An Image Synthesizer", ACM SIGGRAPH Computer Graphics v.19 n.3, p.287-296, July 1985.
- [RFC2104] Krawczyk, H., Bellare, M., and R. Canetti, "HMAC: Keyed-Hashing for Message Authentication", RFC 2104, February 1997.
- [RFC3863] Sugano, H., Fujimoto, S., Klyne, G., Bateman, A., Carr, W., and J. Peterson, "Presence Information Data Format (PIDF)", RFC 3863, August 2004.
- [RFC4079] Peterson, J., "A Presence Architecture for the Distribution of GEOPRIV Location Objects", RFC 4079, July 2005.
- [RFC4086] Eastlake, D., Schiller, J., and S. Crocker, "Randomness Requirements for Security", BCP 106, RFC 4086, June 2005.
- [RFC4119] Peterson, J., "A Presence-based GEOPRIV Location Object Format", RFC 4119, December 2005.
- [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.
- [RFC5965] Shafranovich, Y., Levine, J., and M. Kucherawy, "An Extensible Format for Email Feedback Reports", RFC 5965, August 2010.

Appendix A. Sample Implementation

This javascript implements the obscuring algorithm.

```
/**
 * Location obscurer:
 *   var f = new GeoShape.Fuzzer(100, secret, target);
 *   var reported = f.fuzz(known);
 * This object retains state.
 */
GeoShape.Fuzzer = function(dist, secret, targetIdentity) {
  this.distance = dist;
  var key = Hash.HMAC(secret, targetIdentity, Hash.SHA1);
  this.random = new GeoShape.UIRandom(key, dist);
};
```

```
        this.trigger = null;
        this.used = 0;
        return this;
    };
    GeoShape.Fuzzer.prototype = {
        /**
         * Main obscuring function.
         * @param {GeoShape} a shape
         * @returns {GeoShape.GeoCircle} a fuzzed circle
         */
        fuzz: function(shape) {
            var cu = shape.to2d().calculateCentroid();
            /**
             * cu contains two attributes:
             * centroid: a WGS84 point; uncertainty: a distance in metres
             */
            if (!cu.uncertainty) {
                cu.uncertainty = 0;
            }
            if (this.hasMoved(cu.centroid)) {
                var addunc = Math.max(0, this.distance - cu.uncertainty);
                var centre = this.fuzzPoint(cu.centroid, addunc);
                var unc = Math.max(cu.uncertainty, this.distance);
                this.fuzzed = new GeoShape.GeoCircle(centre, unc);
                var td = this.distance / 2;
                this.trigger = this.randomize(cu.centroid, td);
                this.used = 0;
            }
            this.used++;
            return this.fuzzed;
        },
        /**
         * Determine if the location has moved sufficient distance
         * from the trigger to require fuzzing.
         */
        hasMoved: function(centroid) {
            if (!this.trigger) {
                return true;
            }
            return this.trigger.distanceTo(centroid) > this.distance;
        },
        /**
         * Use a continuously varying random grid to move a point.
         */
        fuzzPoint: function(point, dist) {
            this.random.reset();
            var x = this.random.next(point.lat, point.lng) * 2 - 1;
            var y = this.random.next(point.lat, point.lng) * 2 - 1;
        }
    };
}
```

```

        if (x === 0 && y === 0) {
            return point;
        }
        var d = dist * Math.max(Math.abs(x), Math.abs(y));
        var a;
        if (Math.abs(x) > Math.abs(y)) {
            a = y / x;
        } else {
            a = 2 - x / y;
        }
        if (y < -x) {
            a += 4;
        }
        return point.movePoint(d, a * Math.PI / 4);
    },
    /**
     * Move a point randomly (polar method).
     */
    randomize: function(point, dist) {
        var d = Math.sqrt(Math.random()) * dist;
        var a = Math.random() * 2 * Math.PI;
        return point.movePoint(d, a);
    }
};

/**
 * A uniformly distributed, interpolated, pseudorandom number
 * generator that produces the same value for the same key,
 * location and grid size.
 *
 * @param secret a unique, secret key sequence
 * @param gridSize the desired size of the grid, in metres
 */
GeoShape.UIRandom = function(secret, gridSize) {
    this.key = secret;
    this.grid = 8 * gridSize * 9e-6;
    this.reset();
    return this;
};
GeoShape.UIRandom.prototype = {
    /**
     * Get next pseudorandom value for a latitude and longitude.
     */
    next: function(lat, lng) {
        var lowlat = Math.floor(lat / this.grid) * this.grid;
        var bottom = this.interpLongitude(lowlat, lng);
        var top = this.interpLongitude(lowlat + this.grid, lng);
        var tlat = (lat - lowlat) / this.grid;

```

```
        this.rCount++; /* next time produces a different answer */
        return this.uniformDistInterp(bottom, top, tlat);
    },
    reset: function() {
        this.rCount = 0;
    },
    /* Takes a point and produces a "random" value. */
    hashRandom: function(lat, lng) {
        /* need to fix the lat and lng: 7 decimal places */
        var flat = Math.round(lat * 1e7).toString();
        var flng = Math.round(lng * 1e7).toString();

        var input = [].concat(this.rCount, UTF8(flat),
                               0xff, UTF8(flng));
        var h = Hash.HMAC(this.key, input, Hash.SHA1);
        var r = 0;
        for (var i = 0; i < h.length; ++i) {
            r ^= h[i] << ((i % 4) * 8);
        }
        /* add 0.5 to deal with sign bit */
        return r / Math.pow(2, 32) + 0.5;
    },
    /* interpolate a and b using t, with a uniform distribution */
    uniformDistInterp: function(a, b, t) {
        var r = a * (1 - t) + b * t;
        if (r < t && r < (1 - t)) {
            r = r * r / 2 / t / (1 - t);
        } else if (r > t && r > (1 - t)) {
            r = 1 - (1 - r) * (1 - r) / 2 / t / (1 - t);
        } else {
            r = 0.5 + (r - 0.5) / Math.max(t, 1 - t);
        }
        return r;
    },
    interpLongitude: function(lat, lng) {
        if (Math.abs(lat) >= 90) {
            return this.hashRandom((lat > 0) ? 90 : -90, lng);
        }
        var size = this.grid / Math.cos(lat * Math.PI / 180);
        if ((lng - size / 2) < -180 || (lng + size / 2) > 180) {
            var lngpos = (lng + 360) % 360;
            var rpos = this.interpLongSimple(lat, lngpos, size);
            var lngneg = lngpos - 360;
            var rneg = this.interpLongSimple(lat, lngneg, size);
            var t = ((lng + 360) % 360 - 180 - size / 2) / size;
            return this.uniformDistInterp(rpos, rneg, t);
        }
        return this.interpLongSimple(lat, lng, size);
    }
};
```

```
    },  
    interpLongSimple: function(lat, lng, size) {  
      var lowlng = Math.floor(lng / size) * size;  
      var rlow = this.hashRandom(lat, lowlng);  
      var rhigh = this.hashRandom(lat, lowlng + size);  
      var t = (lng - lowlng) / size;  
      return this.uniformDistInterp(rlow, rhigh, t);  
    }  
  };
```

Author's Address

Martin Thomson
Andrew Corporation
Andrew Building (39)
Wollongong University Campus
Northfields Avenue
Wollongong, NSW 2522
AU

Phone: +61 2 4221 2915
Email: martin.thomson@andrew.com

GEOPRIV
Internet-Draft
Intended status: Informational
Expires: March 11, 2011

M. Thomson
Andrew Corporation
R. Bellis
Nominet UK
September 7, 2010

Location Information Server (LIS) Discovery using IP address and Reverse
DNS
draft-thomson-geopriv-res-gw-lis-discovery-04

Abstract

The residential gateway is a device that has become an integral part of home networking equipment. Discovering a Location Information Server (LIS) is a necessary part of acquiring location information for location-based services. However, discovering a LIS when a residential gateway is present poses a configuration challenge, requiring a method that is able to work around the obstacle presented by the gateway.

This document describes a solution to this problem. The solution provides alternative domain names as input to the LIS discovery process based on the network addresses assigned to a Device.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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

This Internet-Draft will expire on March 11, 2011.

Copyright Notice

Copyright (c) 2010 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	3
2. Conventions used in this document	4
3. Problem Statement	5
3.1. Residential Gateway	6
3.2. Use of Discovery for Third Party Queries	7
3.3. Additional and Optional Constraints	7
4. IP-based DNS Solution	9
4.1. Identification of IP Addresses	9
4.2. Domain Name Selection	10
4.3. When To Use This Method	10
4.4. Necessary Assumptions and Restrictions	11
4.5. Failure Modes	11
4.6. Deployment Considerations	12
5. IANA Considerations	13
6. Security Considerations	14
7. IAB Considerations	15
8. References	17
8.1. Normative References	17
8.2. Informative References	17
Authors' Addresses	19

1. Introduction

A Location Information Server (LIS) is a service provided by an access network. The LIS uses knowledge of the access network topology and other information to generate location for Devices. Devices within an access network are able to acquire location information from a LIS.

The relationship between a Device and an access network might be transient. Configuration of the correct LIS at the Device ensures that accurate location information is available. Without location information, some network services are not available.

The configuration of a LIS address on a Device requires some automated configuration process. This is particularly relevant when it is considered that Devices might move between different access networks. LIS Discovery [I-D.ietf-geopriv-lis-discovery] describes a method that employs the Dynamic Host Configuration Protocol (DHCPv4 [RFC2131], DHCPv6 [RFC3315]) as input to U-NAPTR [RFC4848] discovery.

A residential gateway, or home router, provides a range of networking functions for Devices within the network it serves. In most cases, these functions effectively prevent the successful use of DHCP for LIS discovery.

The drawback with DHCP is that universal deployment of a new option takes a considerable amount of time. Often, networking equipment needs to be updated in order to support the new option. Of particular concern are the millions of residential gateway devices used to provide Internet access to homes and businesses. While [I-D.ietf-geopriv-lis-discovery] describes functions that can be provided by residential gateways to support LIS discovery, gateways built before the publication of this specification do not (and cannot) provide these functions.

This document explores the problem of configuring Devices with a LIS address when a residential gateway is interposed between the Device and access network. Section 3 defines the problem and Section 4 describes a method for determining a domain name that can be used for discovery of the LIS.

In some cases, the solution described in this document is based on a Unilateral Self-Address Fixing (UNSAF) [RFC3424] method. For those cases, this solution is considered transitional until such time as the recommendations for residential gateways in [I-D.ietf-geopriv-lis-discovery] are more widely deployed. Considerations relating to UNSAF applications are described in Section 7.

2. Conventions used in this document

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

This document uses terminology established in [RFC3693] and [RFC5012].

3. Problem Statement

Figure 1 shows a simplified network topology for fixed wire-line Internet access. This arrangement is typical when wired Internet access is provided. The diagram shows two network segments: the access network provided by an internet service provider (ISP), and the residential network served by the residential gateway.

There are a number of variations on this arrangement, as documented in Section 3.1 of [RFC5687]. In each of these variations the goal of LIS discovery is to identify the LIS in the access network.

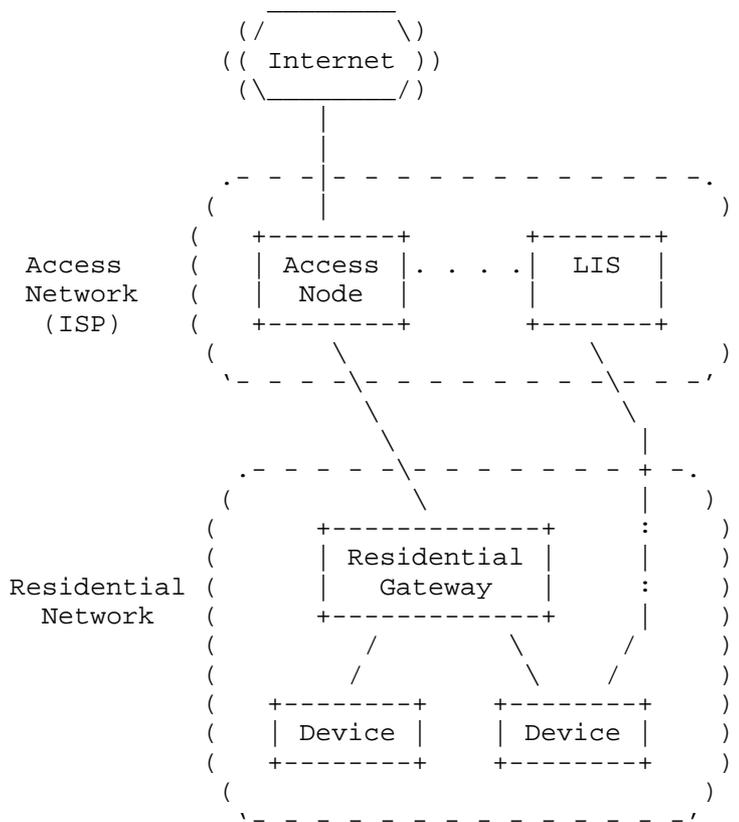


Figure 1: Simplified Network Topology

A particularly important characteristic of this arrangement is the relatively small area served by the residential gateway. Given a small enough area, it is reasonable to delegate the responsibility for providing Devices within the residential network with location

information to the ISP. The ISP is able to provide location information that identifies the residence, which should be adequate for a wide range of purposes.

A residential network that covers a larger area might require a dedicated LIS, a case that is outside the scope of this document.

The goal of LIS discovery is to identify a LIS that is able to provide the Device with accurate location information. In the network topology described, this means identifying the LIS in the access network. The residential gateway is a major obstacle in achieving this goal.

3.1. Residential Gateway

A residential gateway can encompass several different functions including: modem, Ethernet switch, wireless access point, router, network address translation (NAT), DHCP server, DNS relay and firewall. Of the common functions provided, the NAT function of a residential gateway has the greatest impact on LIS discovery.

An ISP is typically parsimonious about their IP address allocations; each customer is allocated a limited number of IP addresses. Therefore, NAT is an extremely common function of gateways. NAT enables the use of multiple Devices within the residential network. However NAT also means that Devices within the residence are not configured by the ISP directly.

When it comes to discovering a LIS, the fact that Devices are not configured by the ISP causes a significant problem. Configuration is the ideal method of conveying the information necessary for discovery. Devices attached to residential gateways are usually given a generic configuration that includes no information about the ISP network. For instance, DNS configuration typically points to a DNS relay on the gateway device. This approach ensures that the local network served by the gateway is able to operate without a connection to the ISP, but it also means that Devices are effectively ignorant of the ISP network.

[I-D.ietf-geopriv-lis-discovery] describes several methods that can be applied by a residential gateway to assist Devices in acquiring location information. For instance, the residential gateway could forward LIS address information to hosts within the network it serves. Such an active involvement in the discovery process only works for new residential gateway devices that implement these recommendations.

Where residential gateways already exist, direct involvement of the

gateway in LIS discovery requires that the residential gateway be updated or replaced. The cost of replacement is difficult to justify to the owner of the gateway, especially when it is considered that the gateway still fills its primary function: Internet access.

Existing residential gateways have proven to be quite reliable devices, some operating continuously for many years without failure. As a result, there are many operational gateways that are of a considerable age, some well outside the period of manufacturer support. Updating the software in such devices is not feasible in many cases. Even if software updates were made available, many residential gateways cannot be updated remotely, inevitably leading to some proportion that is not updated.

This document therefore describes a method which can be used by Devices to discover their LIS without any assistance from the network.

3.2. Use of Discovery for Third Party Queries

It is desirable that any discovery mechanism is capable of being used by hosts outside of the access network. This facilitates third party queries (see [I-D.ietf-geopriv-held-identity-extensions]) by enabling identification of the appropriate LIS.

For example, in some jurisdictions, interim solutions for emergency services require that a voice service provider (VSP) or public safety answering point (PSAP) be able to request location information from the access network provider. These architectures mandate third party queries to accommodate calling devices that are unable to acquire their own location information and subsequently convey [I-D.ietf-sipcore-location-conveyance] that information within call signalling.

This document therefore describes a method which may also be used by third parties to discover the appropriate LIS based on the network address of the Device.

Note that an access network that fully supports DHCP-based LIS discovery [I-D.ietf-geopriv-lis-discovery] might not need to provide a secondary discovery mechanism. However this method SHOULD be provided for the benefit of third parties and for Devices that are unable to use DHCP-based LIS discovery.

3.3. Additional and Optional Constraints

Certain other properties of residential gateways constrain the potential solutions to this problem.

A network firewall function is often provided by residential gateways as a security measure. Security features like intrusion detection systems help protect users from attacks. Amongst these protections is a port filter that prevents both inbound and outbound traffic on certain TCP and UDP ports. Therefore, any solution needs to consider the likelihood of traffic being blocked.

4. IP-based DNS Solution

LIS discovery [I-D.ietf-geopriv-lis-discovery] uses a DNS-based Dynamic Delegation Discovery Service (DDDS) system as the basis of discovery. Input to this process is a domain name. Use of DHCP for acquiring the domain name is specified, but alternative methods of acquisition are permitted.

This document specifies a means for a device to discover several alternative domain names that can be used as input to the DDDS process. These domain names are based on the IP address of the Device. Specifically, the domain names are a portion of the reverse DNS trees - either the ".in-addr.arpa." or ".ip6.arpa." tree.

A Device might be reachable at one of a number of IP addresses. In the process described, a Device first identifies each IP address that it is potentially reachable from. From each of these addresses, the Device then selects up to three domain names for use in discovery. These domain names are then used as input to the DDDS process.

4.1. Identification of IP Addresses

A Device identifies a set of potential IP addresses that currently result in packets being routed to it. These are ordered by proximity, with those addresses that are used in adjacent network segments being favoured over those used in public or remote networks. The first addresses in the set are those that are assigned to local network interfaces.

A Device can use the Session Traversal Utilities for NAT (STUN) [RFC5389] to determine its public reflexive transport address. The host uses the "Binding Request" message and the resulting "XOR-MAPPED-ADDRESS" parameter that is returned in the response.

Alternative methods for determining other IP addresses MAY be used by the host. Universal Plug and Play (UPnP) [UPnP-IGD-WANIPConnection1] and NAT Port Mapping Protocol (NAT-PMP) [I-D.cheshire-nat-pmp] are both able to provide the external address of a residential gateway device when enabled. These as well as proprietary methods for determining other addresses might also be available. Because there is no assurance that these methods will be supported by any access network these methods are not mandated. Note also that in some cases, methods that rely on the view of the network from the residential gateway device could reveal an address in a private address range (see Section 4.4).

In many instances, the IP address produced might be from a private address range. For instance, the address on a local network

interface could be from a private range allocated by the residential gateway. In other cases, methods that rely on the view of the network (UPnP, NAT-PMP) from the residential gateway device could reveal an address in a private address range if the access network also uses NAT. For a private IP address, the derived domain name is only usable where the DNS server used contains data for the corresponding private IP address range.

4.2. Domain Name Selection

The domain name selected for each resulting IP address is the name that would be used for a reverse DNS lookup. The domain name derived from an IP version 4 address is in the ".in-addr.arpa." tree and follows the construction rules in Section 3.5 of [RFC1035]. The domain name derived from an IP version 6 address is in the ".ip6.arpa." tree and follows the construction rules in Section 2.5 of [RFC3596].

Additional domain names are added to allow for a single record to cover a larger set of addresses. If the search on the domain derived from the full IP address does not produce a NAPTR record with the desired service tag (e.g., "LIS:HELD"), a similar search is repeated based on a shorter domain name, using a part of the IP address:

- o For IP version 4, the resulting domain name SHOULD be shortened successively by one and two labels and the query repeated. This corresponds to a search on a /24 or /16 network prefix. This allows for fewer DNS records in the case where a single access network covering an entire /24 or /16 network is served by the same LIS.
- o For IP version 6, the resulting domain SHOULD be shortened successively by 16, 20 and 24 labels and the query repeated. This corresponds to a search on a /64, /48 or /32 network prefix.

DNS queries on other prefixes than those listed above SHOULD NOT be performed to limit the number of DNS queries performed by Devices. If no LIS is discovered by this method, no more than four U-NAPTR resolutions are invoked for each IP address.

4.3. When To Use This Method

The DHCP method described in [I-D.ietf-geopriv-lis-discovery] SHOULD be attempted on all local network interfaces before attempting this method. This method is employed either because DHCP is unavailable, when the DHCP server does not provide a value for the access network domain name option, or if a request to the resulting LIS results in a HELD "notLocatable" error or equivalent.

This method can also be used to facilitate third party queries, as described in Section 3.2. Based on a known IP address, the LIS that serves that address can be identified as long as the corresponding NAPTR records are provided.

4.4. Necessary Assumptions and Restrictions

When used by a Device for LIS discovery this is an UNSAF application and is subject to the limitations described in Section 7.

It is not necessary that the IP address used is unique to the Device, only that the address can be somehow related to the Device or the access network that serves the Device. This allows a degree of flexibility in determining this value, although security considerations (Section 6) might require that the address be verified to prevent falsification.

Addresses from private address space [RFC1918] MAY be used as input to this method. However, it is assumed that a DNS server with a view of the same address space is used in order to provide the corresponding DNS mappings; the public DNS does not contain useful records for all possible address spaces.

This does not preclude the use of private address spaces; use of a private address space in discovery can provide an access network operator more granular control over discovery. This assumes that the DNS server used in the U-NAPTR resolution is able to view the address realm. Addresses from the public address space are more likely to be able to be resolved by any DNS server. Thus, use of the public reflexive transport addresses acquired from a STUN server provide better chance of the DNS server being able to produce a usable result. Therefore, access to a STUN server that is able to view addresses from the public Internet is necessary.

This solution assumes that the public reflexive transport address used by a Device is in some way controlled by their ISP, or some other related party. This implies that the corresponding ".in-addr.arpa." or ".ip6.arpa." record can be updated by that entity to include a useful value for the LIS address.

4.5. Failure Modes

Successful use of private addresses relies on a DNS server that is able to see the private address space; therefore, a means to determine a public IP address is necessary. This document relies on STUN to provide the Device with a public reflexive transport address. Configuration of STUN server is necessary to ensure that this is successful.

Alternative methods for discovering external IP addresses are possible, including UPnP and NAT-PMP. However, these methods might not be enabled on the residential gateway; thus, these methods cannot be relied upon.

In cases where a virtual private network (VPN) or other tunnel is used, the entity providing a public IP address might not be able to provide the Device with location information. It is assumed that this entity is able to identify this problem and indicate this to the Device (using the "notLocatable" HELD error, or similar). This problem is described in more detail in [I-D.ietf-geopriv-http-location-delivery].

4.6. Deployment Considerations

An access network provider SHOULD provide NAPTR records for each public IP address that is used for Devices within the access network. If the access network provider uses NAT, any DNS internal to that NAT SHOULD also include records for the private address range.

NAPTR records can be provided for individual IP addresses. To limit the proliferation of identical records, a single record can be placed at a the higher nodes of the tree (corresponding to /24 and /16 for IPv4; /64, /48 and /32 for IPv6). A record at a higher point in the tree (those with a shorter prefix) applies to all addresses lower in the tree (those with a longer prefix); records at the lower point override those at higher points, allowing for exceptions to be provided for at the lower point.

5. IANA Considerations

[RFC Editor: please remove this section prior to publication.]

This document has no IANA actions.

6. Security Considerations

The security considerations described in [I-D.ietf-geopriv-lis-discovery] apply to the discovery process as a whole. The primary security concern is with the potential for an attacker to impersonate a LIS.

The added ability for a third party to discover the identity of a LIS does not add any concerns, since the identity of a LIS is considered public information.

In addition to existing considerations, this document introduces further security considerations relating to the identification of the IP address. It is possible that an attacker could attempt to provide a falsified IP addresses in an attempt to subvert the rest of the process.

[RFC5389] describes attacks where an attacker is able to ensure that a Device receives a falsified reflexive address. Even if the STUN server is trusted, an attacker might be able to ensure that a falsified address is provided to the Device.

This attack is an effective means of denial of service, or a means to provide a deliberately misleading service. Notably, any LIS that is identified based on a falsified IP address could still be a valid LIS for the given IP address, just not one that is useful for providing the Device with location information. In this case, the LIS provides a HELD "notLocatable" error, or an equivalent. If the falsified IP address is under the control of the attacker, it is possible that misleading (but verifiable) DNS records could indicate a malicious LIS that provides false location information.

In all cases of falsification, the best remedy is to perform some form of independent verification of the result. No specific mechanism is currently available to prevent attacks based on falsification of reflexive addresses; it is suggested that Devices attempt to independently verify that the reflexive transport address provided is accurate.

Use of private address space effectively prevents use of the usual set of trust anchors for DNSSEC. Only a DNS server that is able to see the same private address space can provide useful records. A Device that relies on DNS records in the private address space portion of the ".in-addr.arpa." or ".ip6.arpa." trees MUST either use an alternative trust anchor for these records or rely on other means of ensuring the veracity of the DNS records.

7. IAB Considerations

The IAB has studied the problem of Unilateral Self-Address Fixing (UNSAF) [RFC3424], which is the general process by which a client attempts to determine its address in another realm on the other side of a NAT through a collaborative protocol reflection mechanism, such as STUN.

This section only applies to the use of this method of LIS discovery by Devices and does not apply to its use for third-party LIS discovery.

The IAB requires that protocol specifications that define UNSAF mechanisms document a set of considerations.

1. Precise definition of a specific, limited-scope problem that is to be solved with the UNSAF proposal.

Section 3 describes the limited scope of the problem addressed in this document.

2. Description of an exit strategy/transition plan.

[I-D.ietf-geopriv-lis-discovery] describes behaviour that residential gateways require in order for this short term solution to be rendered unnecessary. When implementations of the recommendations in LIS discovery are widely available, this UNSAF mechanism can be made obsolete.

3. Discussion of specific issues that may render systems more "brittle".

A description of the necessary assumptions and limitations of this solution are included in Section 4.4.

Use of STUN for discovery of a reflexive transport address is inherently brittle in the presence of multiple NATs or address realms. In particular, brittleness is added by the requirement of using a DNS server that is able to view the address realm that contains the IP address in question. If address realms use overlapping addressing space, then there is a risk that the DNS server provides information that is not useful to the Device.

4. Identify requirements for longer term, sound technical solutions; contribute to the process of finding the right longer term solution.

A longer term solution is already provided in

[I-D.ietf-geopriv-lis-discovery]. However, that solution relies on widespread deployment. The UNSAF solution provided here is provided as an interim solution that enables LIS access for Devices that are not able to benefit from deployment of the recommendations in [I-D.ietf-geopriv-lis-discovery].

5. Discussion of the impact of the noted practical issues with existing deployed NATs and experience reports.

The UNSAF mechanism depends on the experience in deployment of STUN [RFC5389]. On the whole, existing residential gateway devices are able to provide access to STUN and DNS service reliably, although regard should be given to the size of the DNS response (see [RFC5625]).

8. References

8.1. Normative References

- [RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, November 1987.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3424] Daigle, L. and IAB, "IAB Considerations for UNilateral Self-Address Fixing (UNSAF) Across Network Address Translation", RFC 3424, November 2002.
- [RFC3596] Thomson, S., Huitema, C., Ksinant, V., and M. Souissi, "DNS Extensions to Support IP Version 6", RFC 3596, October 2003.
- [I-D.ietf-geopriv-http-location-delivery]
Barnes, M., Winterbottom, J., Thomson, M., and B. Stark, "HTTP Enabled Location Delivery (HELD)", draft-ietf-geopriv-http-location-delivery-16 (work in progress), August 2009.
- [I-D.ietf-geopriv-lis-discovery]
Thomson, M. and J. Winterbottom, "Discovering the Local Location Information Server (LIS)", draft-ietf-geopriv-lis-discovery-15 (work in progress), March 2010.

8.2. Informative References

- [RFC1918] Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996.
- [RFC2131] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, March 1997.
- [RFC3315] Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, July 2003.
- [RFC3693] Cuellar, J., Morris, J., Mulligan, D., Peterson, J., and J. Polk, "Geopriv Requirements", RFC 3693, February 2004.
- [RFC4848] Daigle, L., "Domain-Based Application Service Location Using URIs and the Dynamic Delegation Discovery Service

- (DDDS)", RFC 4848, April 2007.
- [RFC5012] Schulzrinne, H. and R. Marshall, "Requirements for Emergency Context Resolution with Internet Technologies", RFC 5012, January 2008.
- [RFC5389] Rosenberg, J., Mahy, R., Matthews, P., and D. Wing, "Session Traversal Utilities for NAT (STUN)", RFC 5389, October 2008.
- [RFC5687] Tschofenig, H. and H. Schulzrinne, "GEOPRIV Layer 7 Location Configuration Protocol: Problem Statement and Requirements", RFC 5687, March 2010.
- [I-D.ietf-sipcore-location-conveyance]
Polk, J., Rosen, B., and J. Peterson, "Location Conveyance for the Session Initiation Protocol", draft-ietf-sipcore-location-conveyance-03 (work in progress), July 2010.
- [UPnP-IGD-WANIPConnection1]
UPnP Forum, "Internet Gateway Device (IGD) Standardized Device Control Protocol V 1.0: WANIPConnection:1 Service Template Version 1.01 For UPnP Version 1.0", DCP 05-001, Nov 2001.
- [I-D.cheshire-nat-pmp]
Cheshire, S., "NAT Port Mapping Protocol (NAT-PMP)", draft-cheshire-nat-pmp-03 (work in progress), April 2008.
- [I-D.ietf-geopriv-held-identity-extensions]
Winterbottom, J., Thomson, M., Tschofenig, H., and R. Barnes, "Use of Device Identity in HTTP-Enabled Location Delivery (HELD)", draft-ietf-geopriv-held-identity-extensions-04 (work in progress), June 2010.
- [RFC5625] Bellis, R., "DNS Proxy Implementation Guidelines", BCP 152, RFC 5625, August 2009.

Authors' Addresses

Martin Thomson
Andrew Corporation
PO Box U40
Wollongong University Campus, NSW 2500
AU

Phone: +61 2 4221 2915
Email: martin.thomson@andrew.com
URI: <http://www.andrew.com/>

Ray Bellis
Nominet UK
Edmund Halley Road
Oxford OX4 4DQ
United Kingdom

Phone: +44 1865 332211
Email: ray.bellis@nominet.org.uk
URI: <http://www.nominet.org.uk/>

GEOPRIV
Internet-Draft
Intended status: Standards Track
Expires: June 20, 2011

J. Winterbottom
M. Thomson
Andrew Corporation
R. Barnes
BBN Technologies
December 17, 2010

Specifying Civic Address Extensions in PIDF-LO
draft-winterbottom-geopriv-local-civic-04

Abstract

New fields are occasionally added to civic addresses. A backwardly-compatible mechanism for adding civic address elements to the Geopriv civic address format is described. A formal mechanism for handling unsupported extensions when translating between XML and DHCP civic address forms is defined for entities that need to perform this translation.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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

This Internet-Draft will expire on June 20, 2011.

Copyright Notice

Copyright (c) 2010 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	3
1.1. Motivating Example	3
1.2. Terminology	4
2. Specifying Civic Address Extensions	4
3. Translating Unsupported Elements	5
3.1. XML to DHCP Format Translation	6
3.2. Extension Civic Address Type (CAtype)	6
3.3. DHCP to XML Format Translation	6
3.4. Conversion Example	7
4. Security Considerations	8
5. IANA Considerations	8
5.1. CAtype Registration for Extensions	8
5.2. End of Numeric CAtype Registration	8
5.3. Registration Template	8
5.4. Registration Policy and Expert Guidance	9
6. Acknowledgements	10
7. References	10
7.1. Normative References	10
7.2. Informative References	10
Authors' Addresses	11

1. Introduction

The Geopriv civic location specifications ([RFC4776], [RFC5139]) define an XML and binary representations for civic addresses that allow for the expression of civic addresses. Guidance for the use of these formats for the civic addresses in different countries is included in [RFC5774].

Subsequent to these specifications being produced, use cases for extending the civic address format with new elements have emerged. Extension elements do not readily fit existing elements, as recommended in [RFC5774].

The XML format for civic addresses [RFC5139] provides a mechanism that allows for the addition of standardized or privately understood elements. A similar facility for private extension is not provided for the DHCP format [RFC4776], though new specifications are able to define new CATypes (civic address types).

A recipient of a civic address in either format currently has no option other than to ignore elements that it does not understand. This results in any elements that are unknown to that recipient being discarded if a recipient performs a translation between the two formats. In order for a new extension to be preserved through translation by any recipient, the recipient has to understand the extension and know how to correlate an XML element with a CAType.

This document describes how new civic address elements are added. Extension always starts with the definition of XML elements. A mechanism for carrying the extension in the DHCP format is described.

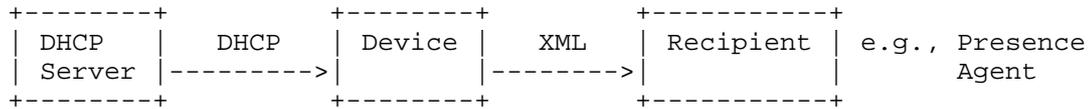
These mechanisms ensure that any translation between formats can be performed consistently and without loss of information. Translation between formats can occur without knowledge of every extension that is present.

These additions described in this document are backwardly compatible. Existing implementations may cause extension information to be lost, but the presence of extensions does not affect an implementation that conforms to either [RFC4776] or [RFC5139].

1.1. Motivating Example

One instance where translation might be necessary is where a device receives location configuration using DHCP [RFC4776]. Conversion of DHCP information to an XML form is necessary if the device wishes to use the DHCP-provided information in a range of applications, including location-based presence services [RFC4079], and emergency

calling [RFC5012].



Conversion Scenario

The Device that performs the translation between the DHCP and XML formats might not be aware of some of the extensions that are in use. Without knowledge of these extensions and how they are represented in XML, the Device is forced to discard them.

These extensions could be useful - or critical - to the ultimate consumers of this information. For instance, an extension element might provide a presence watcher with important information in locating the Device or an extension might be significant in choosing a particular call route.

1.2. Terminology

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

2. Specifying Civic Address Extensions

The civic schema in [RFC5139] defines an ordered structure of elements that can be combined to describe a civic address. The XML extension point at the end of this sequence is used to extend the address.

New elements are defined in a new XML namespace [XMLNS]. This is true of address elements with significance within private or localized domains, as well as those that are intended for global applicability.

New elements SHOULD use the basic "caType" schema type defined in [RFC5139]. This type provides an optional "xml:lang" attribute.

For example, suppose the (fictitious) Central Devon Canals Authority wishes to introduce a new civic element called "bridge". The authority defines an XML namespace that includes a "bridge" element. The namespace needs to be a unique URI, for example "http://devon.canals.org.uk/civic".

A civic address that includes the new "bridge" element is shown in Figure 1.

```
<civicAddress xml:lang="en-GB"
  xmlns="urn:ietf:params:xml:ns:pidf:geopriv10:civicAddr"
  xmlns:cdc="http://devon.canals.org.uk/civic">
  <country>UK</country>
  <A1>Devon</A1>
  <A3>Monkokehampton</A3>
  <RD>Deckport</RD>
  <STS>Cross</STS>

  <cdc:bridge>21451338</cdc:bridge>

</civicAddress>
```

Figure 1: Extended Civic Address Example

An entity that receives this location information might not understand the extension address element. As long as the added element is able to be safely ignored, the remainder of the civic address can be used. The result is that the information is not as useful as it could be, but the added element does not prevent the use of the remainder of the address.

The address can be passed to other applications, such as a LoST server [RFC5222], without modification. If the application understands the added elements, it is able to make use of that information. For example, if this civic address is acquired using HELD [RFC5985], it can be included in a LoST request directly.

3. Translating Unsupported Elements

Unsupported civic address elements can be carried without consequence only as long as the format of the address does not change. When converting between the XML and DHCP formats, these unsupported elements are necessarily discarded: the entity performing the translation has no way to know the correct element to use in the target format.

All extensions **MUST** be defined using the mechanism described in this document. Extensions that use numeric CATypes or other mechanisms cannot be safely translated between XML and DHCP representations.

An entity that does not support these extension mechanisms is expected to remove elements it doesn't understand when performing conversions.

numeric CAtypes. No new registrations of numeric CAtypes can be made.

Extension for the DHCP civic address format is performed by first describing an XML extension. This extension is then carried in the DHCP form in an extension CAtype.

When converting to XML, the namespace prefix used for the extension element is selected by the entity that performs the conversion.

3.4. Conversion Example

The following example civic address contains two extensions:

```
<civicAddress xml:lang="en-US"
  xmlns="urn:ietf:params:xml:ns:pidf:geopriv10:civicAddr"
  xmlns:post="http://postsoftheworld.net/ns"
  xmlns:ap="http://example.com/airport/5.0">
  <country>US</country>
  <A1>CA</A1>

  <post:lamp>2471</post:lamp>
  <post:pylon>AQ-374-4(c)</post:pylon>

  <ap:airport>LAX</ap:airport>
  <ap:terminal>Tom Bradley</ap:terminal>
  <ap:concourse>G</ap:concourse>
  <ap:gate>36B</ap:gate>
</civicAddress>
```

Figure 3: XML Example with Multiple Extensions

This is converted to a DHCP form as follows:

```
country      = US
CAtype[0]    = en-US
CAtype[1]    = CA
CAtype[XX]   = http://postsoftheworld.net/ns lamp 2471
CAtype[XX]   = http://postsoftheworld.net/ns lamp AQ-374-4(c)
CAtype[XX]   = http://example.com/airport/5.0 airport LAX
CAtype[XX]   = http://example.com/airport/5.0 terminal Tom Bradley
CAtype[XX]   = http://example.com/airport/5.0 concourse G
CAtype[XX]   = http://example.com/airport/5.0 gate 36B
```

Figure 4: Converted DHCP Example with Multiple Extensions

4. Security Considerations

This document defines a formal way to extend the existing Geopriv civic address schema. No security threats are introduced by this document.

Security threats applicable to the civic address formats are described in [RFC4776] (DHCP) and [RFC5139] (XML).

5. IANA Considerations

This document alters the "CATypes" registry established by [RFC4776].

5.1. CAtype Registration for Extensions

IANA has allocated a CAtype code of XX for the extension CAtype.

[[IANA/RFC-EDITOR: Please replace XX with the allocated CAtype]]

5.2. End of Numeric CAtype Registration

No further registration of numeric CATypes is permitted. New registrations in this registry use the registration template in Section 5.3.

5.3. Registration Template

New registrations in the "CATypes" registry require the following information:

CAtype: The assigned numeric CAtype. All new registrations use the value XX. [[IANA/RFC-Editor: update XX] Existing registrations use their assigned value.

Namespace URI: A unique identifier for the XML namespace used for the extension element.

Local Name: The local name of an XML element that carries the civic address element.

Description: A brief description of the semantics of the civic address element.

(Optional) Example: One or more simple examples of the element.

Contact: Contact details for the person providing the extension.

(Optional) Specification: A reference to a specification for the civic address element.

(Optional) Schema: A reference to a formal schema (XML schema, RelaxNG, or other form) that defines the extension.

Registrations from [RFC4776] and [RFC5139] are registered with the following form:

CAtype: (The existing CAtype.)

Namespace URI: urn:ietf:params:xml:ns:pidf:geopriv10:civicAddr

Local Name: (The contents of the PIDF column.)

Description: (The existing description for the element, including a note about the equivalent NENA field, if present.)

Contact: The IESG (iesg@ietf.org); the GEOPRIV working group (geopriv@ietf.org).

Specification: RFC4776 and RFC5139

Schema: urn:ietf:params:xml:schema:pidf:geopriv10:civicAddr

5.4. Registration Policy and Expert Guidance

The "CAtypes" registry is altered to operate on a registration policy of "Expert Review", and optionally "Specification Required" [RFC5226].

The registration rules for "Specification Required" are followed only if a registration includes a reference to a specification. Registrations can be made without a specification reference.

All registrations are reviewed to identify potential duplication between registered elements. Duplicated semantics are not prohibited in the registry, though it is preferred if existing elements are used. The expert review is advised to recommend the use of existing elements following the guidance in [RFC5774]. Any registration that is a duplicate or could be considered a close match for the semantics of an existing element SHOULD include a discussion of the reasons that the existing element was not reused.

6. Acknowledgements

Thanks to Brian Rosen, Delaine Arnold, Robins George, and anyone else who has tried to extend the civic schema and found it a little unintuitive.

7. References

7.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC4776] Schulzrinne, H., "Dynamic Host Configuration Protocol (DHCPv4 and DHCPv6) Option for Civic Addresses Configuration Information", RFC 4776, November 2006.
- [RFC5139] Thomson, M. and J. Winterbottom, "Revised Civic Location Format for Presence Information Data Format Location Object (PIDF-LO)", RFC 5139, February 2008.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.
- [XMLNS] Hollander, D., Layman, A., Tobin, R., and T. Bray, "Namespaces in XML 1.1 (Second Edition)", World Wide Web Consortium Recommendation REC-xml-names11-20060816, August 2006, <<http://www.w3.org/TR/2006/REC-xml-names11-20060816>>.

7.2. Informative References

- [RFC3629] Yergeau, F., "UTF-8, a transformation format of ISO 10646", STD 63, RFC 3629, November 2003.
- [RFC4079] Peterson, J., "A Presence Architecture for the Distribution of GEOPRIV Location Objects", RFC 4079, July 2005.
- [RFC5012] Schulzrinne, H. and R. Marshall, "Requirements for Emergency Context Resolution with Internet Technologies", RFC 5012, January 2008.
- [RFC5222] Hardie, T., Newton, A., Schulzrinne, H., and H. Tschofenig, "LoST: A Location-to-Service Translation Protocol", RFC 5222, August 2008.

[RFC5774] Wolf, K. and A. Mayrhofer, "Considerations for Civic Addresses in the Presence Information Data Format Location Object (PIDF-LO): Guidelines and IANA Registry Definition", BCP 154, RFC 5774, March 2010.

[RFC5985] Barnes, M., "HTTP-Enabled Location Delivery (HELD)", RFC 5985, September 2010.

Authors' Addresses

James Winterbottom
Andrew Corporation
Andrew Building (39)
Wollongong University Campus
Northfields Avenue
Wollongong, NSW 2522
AU

Phone: +61 242 212938
Email: james.winterbottom@andrew.com

Martin Thomson
Andrew Corporation
Andrew Building (39)
Wollongong University Campus
Northfields Avenue
Wollongong, NSW 2522
AU

Phone: +61 2 4221 2915
Email: martin.thomson@andrew.com

Richard Barnes
BBN Technologies
9861 Broken Land Parkway
Columbia, MD 21046
US

Phone: +1 410 290 6169
Email: rbarnes@bbn.com

