

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																		
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>

GEOPRIV
Internet-Draft
Updates: 3693, 3694
(if approved)
Intended status: BCP
Expires: April 14, 2011

R. Barnes
M. Lepinski
BBN Technologies
A. Cooper
J. Morris
Center for Democracy &
Technology
H. Tschofenig
Nokia Siemens Networks
H. Schulzrinne
Columbia University
October 11, 2010

An Architecture for Location and Location Privacy in Internet
Applications
draft-ietf-geopriv-arch-03

Abstract

Location-based services (such as navigation applications, emergency services, management of equipment in the field) need geographic location information about Internet hosts, their users, and other related entities. These applications need to securely gather and transfer location information for location services, and at the same time protect the privacy of the individuals involved. This document describes an architecture for privacy-preserving location-based services in the Internet, focusing on authorization, security, and privacy requirements for the data formats and protocols used by these services.

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 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. 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
1.1. Binding Rules to Data	4
1.2. Location-Specific Privacy Risks	5
1.3. Privacy Paradigms	6
2. Terminology Conventions	7
3. Overview of the Architecture	7
3.1. Basic Geopriv Scenario	8
3.2. Roles and Data Formats	10
4. The Location Life-Cycle	13
4.1. Positioning	14
4.1.1. Determination Mechanisms and Protocols	14
4.1.2. Privacy Considerations for Positioning	16
4.1.3. Security Considerations for Positioning	17
4.2. Location Distribution	17
4.2.1. Privacy Rules	18
4.2.2. Location Configuration	20
4.2.3. Location References	20
4.2.4. Privacy Considerations for Distribution	21
4.2.5. Security Considerations for Distribution	23
4.3. Location Use	24
4.3.1. Privacy Considerations for Use	24
4.3.2. Security Considerations for Use	24
5. Security Considerations	25
6. Example Scenarios	27
6.1. Minimal Scenario	27
6.2. Location-based Web Services	28
6.3. Emergency Calling	30
6.4. Combination of Services	32
7. Glossary	34
8. Acknowledgements	37
9. IANA Considerations	37
10. References	37
10.1. Normative References	37
10.2. Informative References	37
Authors' Addresses	39

1. Introduction

Location-based services (applications that require information about the geographic location of an individual or device) are becoming increasingly common on the Internet. Navigation and direction services, emergency services, friend finders, management of equipment in the field and many other applications require geographic location information about Internet hosts, their users, and other related entities. As the accuracy of location information improves and the expense of calculating and obtaining it declines, the distribution and use of location information in Internet-based services will likely become increasingly pervasive. Ensuring that location information is transmitted and accessed in a secure and privacy-protective way is essential to the future success of these services, as well as the minimization of the privacy harms that could flow from their wide deployment and use.

Standards for communicating location information over the Internet have an important role to play in providing a technical basis for privacy and security protection. This document describes a standardized privacy- and security-focused architecture for location-based services in the Internet: the Geopriv architecture. The central component of the Geopriv architecture is the location object, which is used to convey both location information about an individual or device and user-specified privacy rules governing that location information. As location information moves through its life cycle -- positioning, distribution, and use by its ultimate recipient(s) -- Geopriv provides mechanisms to secure the integrity and confidentiality of location objects and to ensure that location information is only transmitted in compliance with the user's privacy rules.

The goals of this document are two-fold: First, the architecture described revises and expands on the basic Geopriv Requirements [2][3], in order to clarify how these privacy concerns and the Geopriv architecture apply to use cases that have arisen since the publication of those documents. Second, this document provides a general introduction to Geopriv and Internet location-based services, and is useful as a good first document for readers new to Geopriv.

1.1. Binding Rules to Data

A central feature of the Geopriv architecture is that location information is always bound to privacy rules to ensure that entities that receive location are informed of how they may use it. These rules can convey simple directives ("do not share my location with others"), or more robust preferences ("allow my spouse to know my exact location all of the time, but only allow my boss to know it

during work hours"). By creating a structure to convey the user's preferences along with location information, the likelihood that those preferences will be honored necessarily increases. In particular, no recipient of the location information can disavow knowledge of users' preferences for how their location may be used. The binding of privacy rules to location information can convey users' desire for and expectations of privacy, which in turn helps to bolster social and legal systems' protection of those expectations.

Binding of usage rules to sensitive information is a common way of protecting information. Several emerging schemes for expressing copyright information provide for rules to be transmitted together with copyrighted works. The Creative Commons [28] model is the most prominent example, allowing an owner of a work to set four types of rules ("Attribution," "Noncommercial," "No Derivative Works" and "ShareAlike") governing the subsequent use of the work. After the author sets these rules, the rules are conveyed together with the work itself, so that every recipient is aware of the copyright terms.

Classification systems for controlling sensitive documents within an organization are another example. In these systems, when a document is created, it is marked with a classification such as "SECRET" or "PROPRIETARY." Each recipient of the document knows from this marking that the document should only be shared with other people who are authorized to access documents with that marking. Classification markings can also convey other sorts of rules, such as a specification for how long the marking is valid (a declassification date). The United States Department of Defense guidelines for classification [4] provide one example.

1.2. Location-Specific Privacy Risks

While location-based services raise some privacy concerns that are common to all forms of personal information, many of them are heightened and others are uniquely applicable in the context of location information.

Location information is frequently generated on or by mobile devices. Because individuals often carry their mobile devices with them, location data may be collected everywhere and at any time, often without user interaction, and it may potentially describe both what a person is doing and where he or she is doing it. For example, location data can reveal the fact that an individual was at a particular medical clinic at a particular time. The ubiquity of location information may also increase the risks of stalking and domestic violence if perpetrators are able to use (or abuse) location-based services to gain access to location information about their victims.

Location information is also of particular interest to governments and law enforcers around the world. The existence of detailed records of individuals' movements should not automatically facilitate the ability for governments to track their citizens, but in some jurisdictions, laws dictating what government agents must do to obtain location data are either non-existent or out-of-date.

1.3. Privacy Paradigms

Traditionally, the extent to which data about individuals enjoys privacy protections on the Internet has largely been decided by the recipients of the data. Internet users may or may not be aware of the privacy practices of the entities with whom they share data. Even if they are aware, they have generally been limited to making a binary choice between sharing data with a particular entity or not sharing it. Internet users have not historically been granted the opportunity to express their own privacy preferences to the recipients of their data and to have those preferences honored.

This paradigm is problematic because the interests of data recipients are often not aligned with the interests of data subjects. While both parties may agree that data should be collected, used, disclosed and retained as necessary to deliver a particular service to the data subject, they may not agree about how the data should otherwise be used. For example, an Internet user may gladly provide his email address on a Web site to receive a newsletter, but he may not want the Web site to share his email address with marketers, whereas the Web site may profit from such sharing. Neither providing the address for both purposes nor deciding not to provide it is an optimal option from the Internet user's perspective.

The Geopriv model departs from this paradigm for privacy protection. As explained above, location information can be uniquely sensitive. And as siloed location-based services emerge and proliferate, they increasingly require standardized protocols for communicating location information between services and entities. Recognizing both of these dynamics, Geopriv gives data subjects the ability to express their choices with respect to their own location information, rather than allowing the recipients of the information to define how it will be used. The combination of heightened privacy risk and the need for standardization compelled the Geopriv designers to shift away from the prevailing Internet privacy model, instead empowering users to express their privacy preferences about the use of their location information.

Geopriv does not, by itself, provide technical means through which it can be guaranteed that users' location privacy rules will be honored by recipients. The privacy protections in the Geopriv architecture

are largely provided by virtue of the fact that recipients of location are informed of relevant privacy rules, and are expected to only use location in accordance with those rules. The distributed nature of the architecture inherently limits the degree to which compliance can be guaranteed and verified by technical means. Section 5 describes how some security mechanisms can address this to a limited extent.

By binding privacy rules to location information, however, Geopriv provides valuable information about users' privacy preferences, so that non-technical forces such as legal contracts, governmental consumer protection authorities, and marketplace feedback can better enforce those privacy preferences. If a commercial recipient of location information, for example, violates the location rules bound to the information, the recipient can in a growing number of countries be charged with violating consumer or data protection laws. In the absence of a binding of rules with location information, consumer protection authorities would be less able to protect individuals whose location information has been abused.

2. Terminology Conventions

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 [1].

3. Overview of the Architecture

This section provides an overview of the Geopriv architecture for the secure and private distribution of location information on the Internet. We describe the three phases of the "location life cycle" -- positioning, distribution and use -- and discuss how the components of the architecture fit within each phase. The next section provides additional detail about how each phase can be achieved in a private and secure manner.

The risks discussed in the previous section all arise from unauthorized disclosure or usage of location information. Thus, the Geopriv architecture has two fundamental privacy goals:

1. Ensure that location information is distributed only to authorized entities, and
2. Provide information to those entities about how they are authorized to use the location information.

If these two goals are met, all parties that receive location information will also receive directives about how they can use that information. Privacy-preserving entities will only engage in authorized uses, and entities that violate privacy will do so knowingly, since they have been informed of what is authorized (and thus, implicitly, of what is not).

Privacy rules and their distribution are thus the central technical components of the privacy system, since they inform location recipients about how they are authorized to use that information. The two goals in the preceding paragraph are enabled by two classes of rules:

1. Access control rules: Rules that describe which entities may receive location information and in what form
2. Usage rules: Rules that describe what uses of location information are authorized

Within this framework for privacy, security mechanisms provide support for the application of privacy rules. For example, authentication mechanisms validate the identities of entities requesting location (so that authorization and access-control policies can be applied), and confidentiality mechanisms protect location information en route between privacy-preserving entities. Security mechanisms can also provide assurances that are outside the purview of privacy by, for example, assuring location recipients that location information has been faithfully transmitted to them by its creator.

3.1. Basic Geopriv Scenario

As location information is transmitted among Internet hosts, it goes through a "location life-cycle": first, the location is computed based on some external information (positioning), then it is transmitted from one host to another (distribution) until finally it is used by a recipient (use).

For example, suppose Alice is using a mobile device, she learns of her location from a wireless location service, and she wishes to share her location privately with her friends by way of a presence service. Alice clearly needs to provide the presence server with her location and rules about which friends can be provided with her location. To enable Alice's friends to preserve her privacy, they need to be provided with privacy rules. Alice may tell some of her friends the rules directly, or she can have the presence server provide the rules to her friends when it provides them with her location. In this way, every friend who receives Alice's location is

authorized by Alice to receive it, and every friend who receives it knows the rules. Good friends will obey the rules. If a bad friend breaks them and Alice finds out, the bad friend cannot claim that he was unaware of the rules.

Some of Alice's friends will be interested in using Alice's location only for their own purposes (to meet up with her or plot her location over time, for example). The usage rules that they receive direct them as to what they can or cannot do (for example, Alice might not want them keeping her location for more than, say, two weeks).

Consider one friend, Bob, who wants to send Alice's location to some of his friends. To operate in a privacy-protective way, Bob needs not only usage rules for himself, but also access control rules that describe who he can send information to and rules to give to the recipients. If the rules he received from the presence server authorize him to give Alice's location to others, he may do so; otherwise, he will require additional rules from Alice before he is authorized to distribute her location. If recipients who receive Alice's location from Bob want to distribute the location on further, they must go through the same process as Bob.

The whole example is illustrated in the following figure:

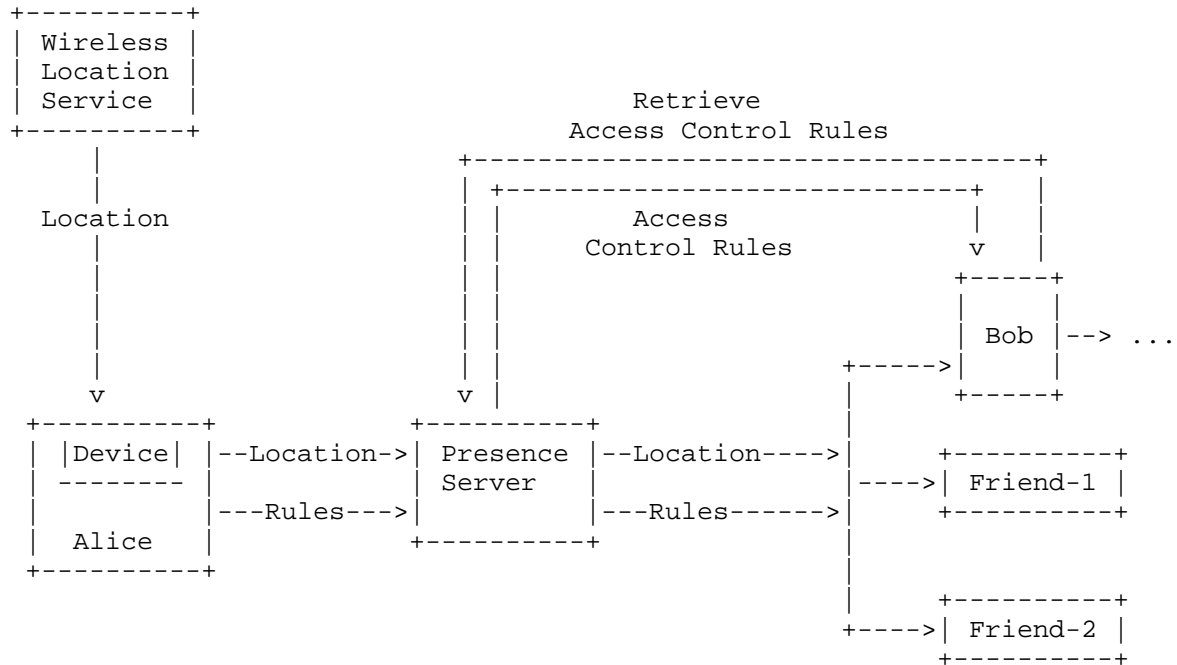


Figure 1: Basic Geopriv Scenario

3.2. Roles and Data Formats

The above example illustrates the six basic roles in the Geopriv architecture:

Target: An individual or other entity whose location is sought in the Geopriv architecture. In many cases the Target will be the human user of a Device, but it can also be an object such as a vehicle or shipping container to which a Device is attached. In some instances the Target will be the Device itself. The Target is the entity whose privacy Geopriv seeks to protect. Alice is the Target in Figure 1.

Device: The technical device whose location is tracked as a proxy for the location of a Target. Alice's device is the Device in Figure 1.

Rule Maker (RM): Performs the role of creating rules governing access to location information for a Target. In some cases the Target performs the Rule Maker role (as is the case with Alice), and in other cases they are separate. For example, a parent may serve as the Rule Maker when the Target is his child, or a corporate security officer may serve as the Rule Maker for devices owned by the corporation but used by employees. The Rule Maker is also not necessarily the owner of the Device. For example, a corporation may provide a Device to an employee but permit the employee to serve as the Rule Maker and set her own privacy rules.

Location Generator (LG): Performs the roles of initially determining or gathering the location of the Device and providing it to Location Servers. Location Generators may be any sort of software or hardware used to obtain the Device's location (examples include GPS chips and cellular networks). A Device may even perform the Location Generator role for itself; Devices capable of unassisted satellite-based positioning and Devices that accept manually entered location information are two examples. The wireless location service plays the Location Generator role in Figure 1.

Location Server (LS): Performs the roles of receiving location information and rules, applying the rules to the location information to determine what other entities, if any, can receive location information, and providing the location to Location Recipients. Location Servers receive location information from Location Generators and rules from Rule Makers, and then apply the rules to the location information. Location Servers may not necessarily be "servers" in the colloquial sense of hosts in remote data centers servicing requests. Rather, a Location Server can be any software or hardware component that distributes location information. Examples include a server in an access network, a presence server, or a Web browser or other software running on a Device. The above example includes three Location Servers: Alice, the presence service and Bob.

Location Recipient (LR): Performs the role of receiving location information. A Location Recipient may ask for location explicitly (by sending a query to a Location Server), or it may receive location asynchronously. The presence service, Bob, Friend-1 and Friend-2 are Location Recipients in Figure 1.

In general, these roles may or may not be performed by physically separate entities, as demonstrated by the entities in Figure 1, many of which perform multiple roles. It is not uncommon for the same entity to perform both the Location Generator and Location Server roles, or both the Location Recipient and Location Server roles. A

single entity may take on multiple roles simply by virtue of its own capabilities and the permissions provided to it.

Although in the above example there is only a single Location Generator and a single Rule Maker, in some cases a Location Server may receive Location Objects from multiple Location Generators or Rules from multiple Rule Makers. Likewise, a single Location Generator may publish location information to multiple Location Servers, and a single Location Recipient may receive Location Objects from multiple Location Servers.

There is a close relationship between a Target and its Device. The term "Device" is used when discussing protocol interactions, whereas the term "Target" is used when discussing generically the person or object being located and its privacy. While in the example above there is a one-to-one relationship between the Target and the Device, Geopriv can also be used to convey location information about a device that is not directly linked to a single individual or object, such as a Device shared by multiple individuals.

Two data formats are necessary within this architecture:

Location Object (LO): An object used to convey location information together with Privacy Rules. Geopriv supports both geodetic location data (latitude/longitude/altitude/etc.) and civic location data (street/city/state/etc.). Either or both types of location information may be present in a single LO (see the considerations in [5] for LOs containing multiple locations). Location Objects typically include some sort of identifier associated with the Target.

Privacy Rule: A directive that regulates an entity's activities with respect to location information, including the collection, use, disclosure, and retention of the location information. Privacy Rules describe which entities may obtain location information in what form (access control rules) and how location information may be used by an entity (usage rules).

The whole example, using Geopriv roles and formats, is illustrated in the following figure:

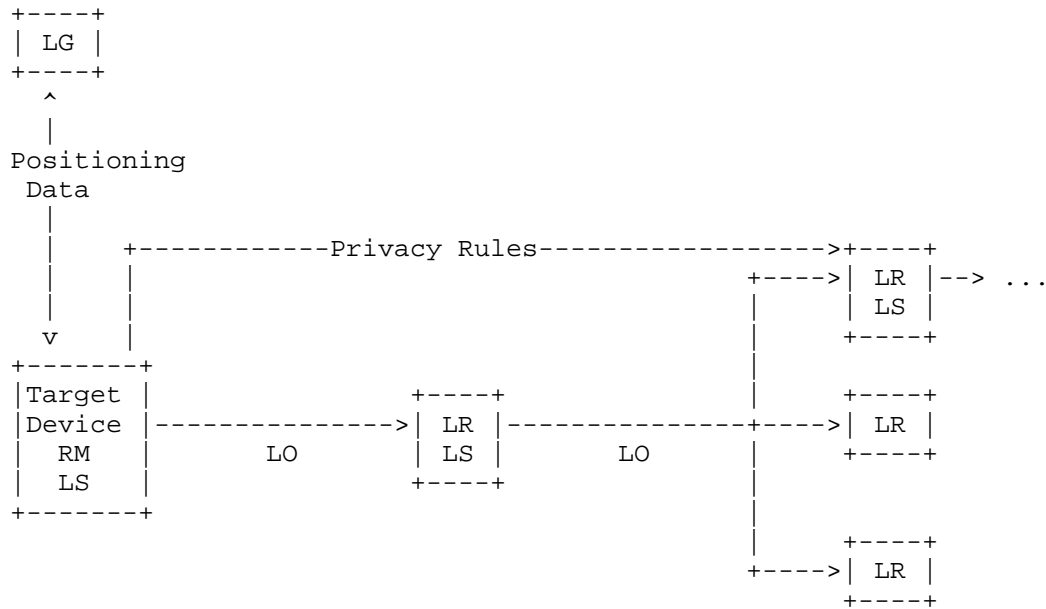


Figure 2: Basic Geopriv Scenario

4. The Location Life-Cycle

The previous section gave an example of how an individual's location can be distributed through the Internet. In general, the location life-cycle breaks down into three phases:

1. Positioning: A Location Generator determines the Device's location.
2. Distribution: Location Servers send location to Location Recipients, which may in turn act as Location Servers and further distribute location to other Location Recipients (possibly several times).
3. Use: A Location Recipient receives the location and uses it.

Each of these phases involves a different set of Geopriv roles and each has a different set of privacy and security implications. The Geopriv roles are mapped onto the location life-cycle in the figure below.

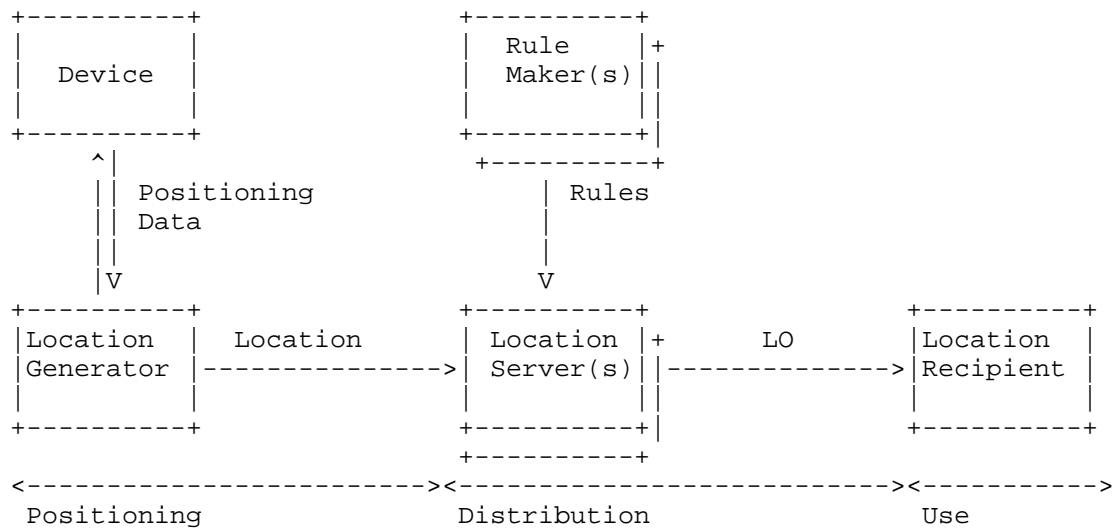


Figure 3: Location Life-Cycle

4.1. Positioning

Positioning is the process by which the physical location of the Device is computed, based on some observations about the Device's situation in the physical world. (This process goes by several other names, including Location Determination or Sighting.) The input to the positioning process is some information about the Device, and the outcome is that the LG knows the location of the Device.

In this section, we give a brief taxonomy of current positioning systems, their requirements for protocol support, and the privacy and security requirements for positioning.

4.1.1. Determination Mechanisms and Protocols

While the specific positioning mechanisms that can be applied for a given Device are strongly dependent on the physical situation and capabilities of the Device, these mechanisms generally fall into the three categories described in detail below:

- o Device-based
- o Network-based
- o Network-assisted

As suggested by the above names, a positioning scheme can rely on the Device, an Internet-accessible resource (not necessarily a network operator), or a combination of the two. For a given scheme, the nature of this reliance will dictate the protocol mechanisms needed to support it.

With Device-based positioning mechanisms, the Device is capable of determining its location by itself. This is the case for manually-entered location or for (unassisted) satellite-based positioning (using a Global Navigation Satellite System, or GNSS). In these cases, the Device acts as its own LG, and there are no protocols required to support positioning (since no information needs to be communicated).

In network-based positioning schemes, an external LG (an Internet host other than the Device) has access to sufficient information about the Device, through out-of-band channels, to establish the position of the Device. The most common examples of this type of LG are entities that have a physical relationship to the Device (such as ISPs). In wired networks, wiremap-based location is a network-based technique; in wireless networks, timing and signal-strength based techniques that use measurements from base stations are considered to be network-based. Large-scale IP-to-geo databases (for example, those based on WHOIS data or latency measurements) are also considered to be network-based positioning mechanisms.

For network-based positioning as for Device-based, no protocols are strictly necessary to support positioning, since positioning information is collected outside of the location distribution system (at lower layers of the network stack, for example). This does not rule out the use of other Internet protocols (like SNMP) to collect inputs to the positioning process. Rather, since these inputs can only be used by certain LGs to determine location, they are not controlled as private information. Network-based positioning often provides location to protocols by which the network informs a Device of its own location (these are known as Location Configuration Protocols, see Section 4.2.2 for further discussion).

Network-assisted systems account for the greatest number and diversity of positioning schemes. In these systems, the work of positioning is divided between the Device and an external LG via some communication (possibly over the Internet), typically in one of two ways:

- o The Device provides measurements to the LG
- o The LG provides assistance data to the Device

"Measurements" are understood to be observations about the Device's environment, ranging from wireless signal strengths to the MAC address of a first-hop router. "Assistance" is the complement to measurement, namely the positioning information that enables the computation of location based on measurements. A set of wireless base station locations (or wireless calibration information) would be an assistance datum, as would be a table that maps routers to buildings in a corporate campus.

For example, wireless and wired networks can serve as the basis for network-assisted positioning. In several current 802.11 positioning systems, the Device sends measurements (e.g., MAC addresses and signal strengths) to an LG, and the LG returns a location to the client. In wired networks, the Device can send its MAC address to the LG, which can query the MAC-layer infrastructure to determine the switch and port to which that MAC address is connected, then query a wire map to determine the location at which the wire connected to that port terminates.

As an aside, the common phrase "assisted GPS" ("assisted GNSS" more broadly) actually encompasses techniques that transmit both measurements and assistance data. Systems in which the Device provides the LG with GNSS measurements are measurement-based, while those in which the assistance server provide ephemeris or almanac data are assistance-based in the above terminology. (Those familiar with GNSS positioning will note that there are of course cases in which both of these interactions occur within a single location determination protocol, so the categories are not mutually exclusive.)

Naturally, the exchange of measurement or positioning data between the Device and the LG requires a protocol over which the information is carried. The structure of this protocol will depend on which of the two patterns a network-assisted scheme follows. Conversely, the structure of the protocol will determine which of the two parties (the Device, the LG, or both) is aware of the Device's location at the end of the protocol interaction.

4.1.2. Privacy Considerations for Positioning

Positioning is the first point at which location may be associated with a particular Target's identity. Local identifiers, unlinked pseudonyms, or private identifiers that are not linked to the real identity of the Target should be used as forms of identity whenever possible. This provides privacy protection by disassociating the location from the Target's identity before it is distributed.

At the conclusion of the positioning process, the entity acting as

the LG has the Device's location (if the Device is performing the LG role, then they both have it). If the entity acting as the LG also performs the role of LS, the privacy considerations in Section 4.2.4 apply.

In some deployment scenarios, positioning functions and distribution functions may need to be provided by separate entities, in which case the LG and LS roles will not be performed by the same entity. In this situation, the LG acts as a "dumb," non-privacy-aware positioning resource, and the LS provides the privacy logic necessary to support distribution (possibly with multiple LSes using the same LG). In order to allow the privacy-unaware LG to distribute location to these LSes while maintaining privacy, the relationship between the LG and its set of LSes MUST be tightly constrained, effectively "hard-wired." That is, the LG MUST only provide location to a small fixed set of LSes, and each of these LSes MUST comply with the requirements of Section 4.2.4.

4.1.3. Security Considerations for Positioning

Manipulation of the positioning process can expose location through two mechanisms:

1) A third party could guess or derive measurements about a specific device and use them to get the location of that Device. To mitigate this risk, the LG SHOULD be able to authenticate and authorize devices providing measurements and, if possible, verify that the presented measurements are likely to be the actual physical values measured by that client. These security procedures rely on the type of positioning being done, and may not be technically feasible in all cases.

2) By eavesdropping, a third party may be able to obtain measurements sent by the Device itself that indicate the rough position of the Device. To mitigate this risk, protocols used for positioning MUST provide confidentiality and integrity protections in order to prevent observation and modification of transmitted positioning data while en route between the Target and the LG.

If an LG or a Target chooses to act as an LS, it inherits the security requirements for an LS, described in Section 4.2.5.

4.2. Location Distribution

When an entity receives location (from an LG or an LS) and redistributes it to other entities, it acts as an LS. Location Distribution is the process by which one or more LSes provide LOs to LRs in a privacy-preserving manner.

The role of an LS is thus two-fold: First, it must collect location information and Rules that control access to that information. Rules can be communicated within an LO, within a protocol that carries LOs, or through a separate protocol that carries Rules. Second, the LS must process requests for location and apply the Rules to these requests in order to determine whether it is authorized to fulfill them by returning location.

An LS thus has at least two types of interactions with other hosts, namely receiving and sending LOs. An LS may optionally implement a third interaction, allowing Rule Makers to provision it with Rules. The distinction between these two cases is important in practice, because it determines whether the LS has a direct relationship with a Rule Maker: An LS that accepts Rules directly from a Rule Maker has such a relationship, while an LS that acquires all its Rules through LOs does not.

4.2.1. Privacy Rules

Privacy Rules are the central mechanism in Geopriv for maintaining a Target's privacy, because they provide a recipient of an LO (an LS or LR) with information on how the LO may be used.

Throughout the Geopriv architecture, Privacy Rules are communicated in rules languages with a defined syntax and semantics. For example, the Common Policy rules language has been defined [6] to provide a framework for broad-based rule specifications. Geopriv Policy [7] defines a language for creating location-specific rules. XCAP [8] can be used as a protocol to install rules in both of these formats.

Privacy Rules follow a default-deny pattern: an empty set of Rules implies that all requests for location should be denied (other than requests made by the Target itself), with each Rule added to the set granting a specific permission. Adding a Rule can only augment privacy protections because all Rules are positive grants of permission.

The following are examples of Privacy Rules governing location distribution:

- o Retransmit location when requested from example.com
- o Retransmit only city and country
- o Retransmit location with no less than a 100 meter radius of uncertainty

- o Retransmit location only for the next two weeks

LSes enforce Privacy Rules in two ways: by denying requests for location, or by transforming the location information before retransmitting it.

LSes may also receive Rules governing location retention, such as "Retain location only for 48 hours." Such Rules are simply directives about how long the Target's location information can be retained.

Privacy Rules can govern the behavior of both LSes and LRs. Rules that direct LSes about how to treat a Target's location information are known as Local Rules. Local Rules are used internally by the LS to handle requests from LRs. They are not distributed to LRs.

Forwarded Rules, on the other hand, travel inside LOs and direct LSes and LRs about how to handle the location information they receive. Because the Rules themselves may reveal potentially sensitive information about the Target, only the minimal subset of Forwarded Rules necessary to handle the LO is distributed.

An example can illustrate the interaction between Local Rules and Forwarded Rules. Suppose Alice provides the following Local Rules to an LS:

- o The LS may retransmit Alice's precise location to Bob, who in turn is permitted to retain the location information for one month
- o The LS may retransmit Alice's city, state, and country to Steve, who in turn is permitted to retain the location information for one hour
- o The LS may retransmit Alice's country to a photo-sharing website, which in turn is permitted to retain the location information for one year and retransmit it to any requesters

When Steve asks for Alice's location, the LS can transmit to Steve the limited location information (city, state, and country) along with Forwarded Rules instructing Steve to (a) not further retransmit Alice's location information, and (b) only retain the location information for one hour. By only sending these specifically applicable Forwarded Rules to Steve (as opposed to the full set of Local Rules), the LS is protecting Alice's privacy by not disclosing to Steve that (for example) Alice allows Bob to obtain more precise location information than Alice allows Steve to receive.

Geopriv is designed to be usable even by devices with constrained

processing capabilities. To ensure that Forwarded Rules can be processed on constrained devices, LOs are required to carry only a limited set of Forwarded Rules, with an option to reference a more robust set of external Rules. The limited Rule set covers two privacy aspects: how long the Target's location may be retained ("Retention"), and whether or not the Target's location may be retransmitted ("Retransmission"). A LO may contain a pointer to more robust Rules, such as those shown in the set of four Rules at the beginning of this section.

4.2.2. Location Configuration

Some entities performing the LG role are designed only to provide Targets with their own locations (as opposed to distributing a Target's location to others). The process of providing a Target with its own location is known within Geopriv as Location Configuration. The term Location Information Server (LIS) is often used to describe the entity that performs this function (although a LIS may also perform other functions, such as providing a Target's location to other entities).

A Location Configuration Protocol (LCP) [9] is one mechanism that can be used by a Device to discover its own location from a LIS. LCPs provide functions in the way they obtain, transport and deliver location requests and responses between a LIS and a Device such that the LIS can trust that the location requests and responses handled via the LCP are in fact from/to the Target. Several LCPs have been developed within Geopriv [10][11][12][13].

A LIS whose sole purpose is to perform Location Configuration need only follow a simple privacy-preserving policy: transmit a Target's location only to the Target itself. This is known as the "LCP policy."

Importantly, if an LS is also serving in the role of LG and it has not been provisioned with Privacy Rules for a particular Target, it MUST follow the LCP policy, whether it is a LIS or not. In the positioning phase, an entity serving the roles of both LG and LS that has not received Privacy Rules must follow this policy. The same is true for any LS in the distribution phase.

4.2.3. Location References

The location distribution process occurs through a series of transmissions of LOs: transmissions of location "by value." Location "by value" can be expressed in terms of geodetic location data (latitude/longitude/altitude/etc.) and civic location data (street/city/state/etc.).

Location can also be distributed "by reference," where a reference is represented by a URI that can be dereferenced to obtain the LO. This document summarizes the properties of location-by-reference that are discussed at length in [14].

Distribution of location by reference (distribution of location URIs) offer several benefits. Location URIs are a more compact way of transmitting location, since URIs are usually smaller than LOs. A recipient of location can make multiple requests to a URI over time to receive updated location (if the URI is configured to provide fresh location rather than a single "snapshot").

From a positioning perspective, location by reference can offer the additional benefit of "just in time" positioning. If location is distributed by reference, an entity acting as a combined LG/LS only needs to perform positioning operations when a recipient dereferences a previously distributed URI.

From a privacy perspective, distributing location as a URI instead of as an LO can help protect privacy by forcing each recipient of the location to request location from the referenced LS, which can then apply access controls individually to each recipient. But the benefit provided here is contingent on the LS applying access controls. If the LS does not apply an access control policy to requests for a location URI (in other words, if it enforces the "possession model" defined in [14]), then transmitting a location URI presents the same privacy risks as transmitting the LO itself. Moreover, the use of location URIs without access controls can introduce additional privacy risks: If URIs are predictable, an attacker to whom the URI has not been sent may be able to guess the URI and use it to obtain the referenced LO. To mitigate this, location URIs without access controls need to be constructed so that they contain a random component with sufficient entropy to make guessing infeasible.

4.2.4. Privacy Considerations for Distribution

Location information MUST be accompanied by Rules throughout the distribution process. Otherwise, a recipient will not know what uses are authorized, and will not be able to use the LO. Consequently, LOs MUST be able to express Rules that convey appropriate authorizations.

An LS MUST only accept Rules from authorized Rule Makers. For an LS that receives Rules exclusively in LOs and has no direct relationship with a Rule Maker, this requirement is met by applying the Rules provided in an LO to the distribution of that LO. For an LS with a direct relationship to a Rule Maker, this requirement means that the LS MUST be configurable with an RM authorization policy. An LS

SHOULD define a prescribed set of RMs that may provide Rules for a given Target or LO. For example, an LS may only allow the Target to set Rules for itself, or it might allow an RM to set Rules for several Targets (e.g., a parent for children, or a corporate security officer for employees).

No matter how Rules are provided to an LS, for each LO it receives, it MUST combine all Rules that apply to the LO into a Rule set that defines which transmissions are authorized, and it MUST transmit location only in ways that are authorized by these Rules.

An LS that receives Rules exclusively through LOs MUST examine the Rules that accompany a given LO in order to determine how the LS may use the LO (if any Rules are included by reference, the LS SHOULD attempt to download them). If the LO includes no Rules that allow the LS to transmit the LO to another entity, then the LS MUST NOT transmit the LO. If the LO contains no Rules at all (if it is in a format with no Rules syntax, for example), then the LS MUST delete it (emergency services provide an exception in that Rules can be implicit, see [15]). If the LO included Rules by reference, but these Rules were not obtained for any reason, the LS MUST NOT transmit the LO and MUST delete it.

An LS that receives Rules both directly from one or more Rule Makers and through LOs MUST combine the Rules in a given LO with Rules it has received from the RMs. The strategy the LS uses to combine these sets of Rules is a matter for local policy, depending on the relative priority that the LS grants to each source of Rules. Some example policies:

Union: A transmission of location is authorized if it is authorized by either a rule in the LO or an RM-provided rule.

Intersection: A transmission of location is authorized if it is authorized by both a rule in the LO and an RM-provided rule.

RM Override: A transmission of location is authorized if it is authorized by an RM-provided rule (regardless of the LO Rules).

LO Override: A transmission of location is authorized if it is authorized by an LO-provided rule (regardless of the RM Rules).

Different policies may be applicable in different scenarios. In cases where an external RM is more trusted than the source of the LO, the "RM Override" policy may be suitable (for example, if the external RM is the Target, and the LO is provided by a third party). Conversely, the "LO Override" policy is better suited to cases where the LO provider is more trusted than the RM (for example, if the RM is

the user of a mobile device LS and the LO contains Rules from the RM's parents or corporate security office). The "Intersection" policy takes the strictest view of the permission grants, giving equal weight to all RMs (including the LO creator).

Each of these policies will also have different privacy consequences. Following the "Intersection" policy ensures that the most privacy-protective subset of all RMs' rules will be followed. The "Union" policy and both "Override" policies may defy the expectations of any RM (including, potentially, the Target) whose policy is not followed. For example, if a Target acting as an RM sets Rules and those Rules are overridden by the application of a more permissive LO Override policy that has been set by the Target's parent or employer acting as an RM, the retransmission or retention of the Target's data may come as a surprise to the Target. For this reason, it is RECOMMENDED that LSes provide a way for RMs to be able to find out which policy will be applied to the distribution of a given LO.

4.2.5. Security Considerations for Distribution

An LS's decisions about how to transmit location are based on the identities of entities requesting information and other aspects of requests for location. In order to ensure that these decisions are made properly, the LS needs assurance of the reliability of information on the identities of the entities with which the LS interacts (including LRs, LSes, and RMs) and other information in the request.

Protocols to convey LOs and protocols to convey Rules MUST provide information on the identity of the recipient of location and the identity of the RM, respectively. In order to ensure the validity of this information, these protocols MUST allow for mutual authentication of both parties, and MUST provide integrity protection for protocol messages. These security features ensure that the LG has sufficient information (and sufficiently reliable information) to make privacy decisions.

As they travel through the Internet, LOs necessarily pass through a sequence of intermediaries, ranging from layer-2 switches to IP routers to application-layer proxies and gateways. The ability of an LS to protect privacy by making access control decisions is reduced if these intermediaries have access to an LO as it travels between privacy-preserving entities.

Ideally, LOs SHOULD be transmitted with confidentiality protection end-to-end between an LS that transmits location and the LR that receives it. In some cases, the protocol conveying an LO provides confidentiality protection as a built-in security solution for its

signaling (and potentially its data traffic). In this case, carrying an unprotected LOs within such an encrypted channel is sufficient. Many protocols, however, are offering communication modes where messages are either unprotected or protected on a hop-by-hop basis (for example, between intermediaries in a store-and-forward protocol). In such a case it is RECOMMENDED that the protocol allows for the use of encrypted LOs, or for the transmission of a reference to location in place of an LO [14].

4.3. Location Use

The primary privacy requirement of an LR is to constrain its usage of location to the set of uses authorized by the Rules in an LO. If an LR only uses an LO in ways that have minimal privacy impact -- specifically, if it does not transmit the LO to any other entity, and does not retain the LO for longer than is required to complete its interaction with the LS -- then no further action is necessary for the LR to comply with Geopriv requirements.

As an example of this simplest case, if an LR (a) receives a location, (b) immediately provides to the Target information or a service based on the location, (c) does not retain the information, and (d) does not retransmit the location to any other entity, then the LR will comply with any set of Rules that are permissible under Geopriv. Thus, a service that, for example, only provides directions to the closest bookstore in response to an input of location, and promptly then discards the input location, will be in compliance with any Geopriv Rule set.

LRs that make other uses of an LO (e.g., those that store LOs, or send them to other service providers to obtain location-based services) MUST meet the requirements below to assure that these uses are authorized.

4.3.1. Privacy Considerations for Use

The principal privacy requirement for LR is to follow usage rules. Any LR that wants to retransmit or retain the LO is REQUIRED to examine the rules included with that LO. Any usage the LR makes of the LO MUST be explicitly authorized by these Rules. Since Rules are positive grants of permission, any action not explicitly authorized is denied by default.

4.3.2. Security Considerations for Use

Since the LR role does not involve transmission of location, there are no protocol security considerations required to support privacy (other than ensuring that data does not leak unintentionally caused

by security breaches).

Aside from privacy, LRs often require some assurance that an LO is reliable (assurance of the integrity, authenticity, and validity of an LO), since LRs use LOs in order to deliver location-based services. Threats against this reliability and corresponding mitigations are discussed in the Security Considerations below.

5. Security Considerations

Security considerations related to the privacy of LOs are discussed throughout this document. In this section we summarize those concerns and consider security risks not related to privacy.

The life-cycle of an LO often consists of a series of location transmissions. Protocols that carry location can provide strong assurances, but only for a single segment of the LO's life cycle. In particular, a protocol can provide integrity protection and confidentiality for the data exchanged, and mutual authentication of the parties involved in the protocol, by using a secure transport such as IPSec [16] or TLS [17].

Additionally, if (1) the protocol provides mutual authentication for every segment, and (2) every entity in the location distribution chain exchanges information only with entities with whom it has a trust relationship, entities can transitively obtain assurances regarding the origin and ultimate destination of the LO. Of course, direct assurances are always preferred over assurances requiring transitive trust, since they require fewer assumptions.

Using protocol mechanisms alone, the entities can receive assurances only about a single hop in the distribution chain. For example, suppose that an LR receives location from an LS over an integrity- and confidentiality-protected channel. The LR knows that the transmitted LO has not been modified or observed en route. However, the assurances provided by the protocol do not guarantee that the transmitted LO was not corrupted before it was sent to the LS (by a previous LS, for example). Likewise, the LR can verify that the LO was transmitted by the LS, but cannot verify the origin of the LO if it did not originate with the LS.

Security mechanisms in protocols are thus unable to provide direct assurances over multiple transmissions of an LO. However, the transmission of location "by reference" can be used to effectively turn multi-hop paths into single-hop paths. If the multiple transmissions of an LO are replaced by multiple transmissions of a URI (a multi-hop dissemination channel), the LO need only traverse a

single hop, namely the dereference transaction between the LR and the dereference server. The requirements for securing location passed by reference [14] are applicable in this case.

The major threats to the security of LOs can be grouped into two categories. First, threats against the integrity and authenticity of LOs can expose entities that rely on LOs. Second, threats against the confidentiality of LOs can allow unauthorized access to location information.

An LO contains four essential types of information: identifiers for the described Target, location information, time-stamps, and Rules. By grouping values of these various types together within a single structure, an LO encodes a set of bindings among them. That is, the LO asserts that the identified Target was present at the given location at the given time and that the given Rules express the Target's desired policy at that time for the distribution of his location. Below, we provide a description of the assurances required by each party involved in the location distribution in order to mitigate the possible attacks on these bindings.

Rule Maker: The Rule Maker is responsible for creating the Target's Privacy Rules and for uploading them to the LSes. The primary assurance required by the Rule Maker is that the Target's Privacy Rules are correctly associated with the Target's identity when they are conveyed to each LS that handles the LO. Ensuring the integrity of the Privacy Rules distributed to the LSes prevents rule-tampering attacks. In many circumstances, the privacy policy of the Target may itself be sensitive information; in these cases, the Rule Maker also requires the assurance that the binding between the Target's identity and the Target's Privacy Rules are not deducible by anyone other than an authorized LS.

Location Server: The Location Server is responsible for enforcing the Target's Privacy Rules. The first assurance required by the LS is that the binding between the Target's Privacy Rules and the Target's identity is authentic. Authenticating and authorizing the Rule Maker who creates, updates and deletes the Privacy Rules prevents rule-tampering attacks. The LS has to ensure that the authorization policies are not exposed to third parties, if so desired by the Rule Maker (when the rules themselves are privacy-sensitive).

Location Recipient: The Location Recipient is the consumer of the LO. The LR thus requires assurances about the authenticity of the bindings between the Target's location, the Target's identity and the time. Ensuring the authenticity of these bindings helps to prevent various attacks, such as falsifying the location, modifying

the time-stamp, faking the identity, replaying LOss.

Location Generator: The primary assurance required by the Location Generator is that the LS to which the LO is initially published is one that is trusted to enforce the Target's Privacy Rules. Authenticating the trusted LS mitigates the risk of server impersonation attacks. Additionally, the LG is responsible for the location determination process, which is also sensible from a security perspective because wrong input provided by external entities can lead to undesirable disclosure or access to location information.

Assurances as to the integrity and confidentiality of a Location Object can be provided directly through the LO format. RFC 4119 [18] provides a description for usage of S/MIME to integrity and confidentiality protection. Although such direct, end-to-end assurances are desirable, and these mechanisms should be used whenever possible, there are many deployment scenarios where directly securing an LO is impractical. For example, in some deployment scenarios a direct trust relationship may not exist between the creator of the Location Object and the recipient. Additionally, in a scenario where many recipients are authorized to receive a given LO, the creator of the LO cannot guarantee end-to-end confidentiality without knowing precisely which recipient will receive the LO. Many of these cases can, however, be addressed by the usage of a Location-by-Reference (possibly combined with an LO).

6. Example Scenarios

This section contains a set of example of how the Geopriv architecture can be deployed in practice. These examples are meant to illustrate key points of the architecture, rather than to form an exhaustive set of use cases.

For convenience and clarity in these examples, we assume that the Privacy Rules that an LO carries are equivalent to those in a PIDF-LO (namely, that the principal Rules that can be set are limits on the retransmission and retention of the LO). While these two Rules are the most well-known and important examples, the specific types of Rules an LS or LR must consider will in general depend on the types of LO it processes.

6.1. Minimal Scenario

One of the simplest scenarios in the Geopriv architecture is when a Device determines its own location and uses that LO to request a service (e.g., by including the LO in an HTTP POST request [19] or

SIP INVITE message [20]), and the server delivers that service immediately (e.g., in a 200 OK response in HTTP or SIP), without retaining or retransmitting the Device's location. The Device acts as an LG by using a Device-based positioning algorithm (e.g., manual entry) and as an LS by interpreting the rule and transmitting the LO. The Target acts as a Rule Maker by specifying that the location should be sent to the server. The server acts as an LR by receiving and using the LO.

In this case, the privacy of location information is maintained in two steps: The first step is that location is only transmitted as directed by the single Rule Maker, namely the Target. The second step is simply the fact that the server, as LR, does not do anything that creates a privacy risk -- it does not retain or retransmit location. Because the server limits its behavior in this way, it does not need to read the Rules in the LO (even though they were provided) -- no Rule would prevent it from using location in this safe manner.

The following outline summarizes this scenario:

- o Positioning: Device-based, Device=LG
- o Distribution hop 1: HTTP UA --> Ephemeral web service, privacy via user indication
- o Use: Ephemeral web service delivers response without retaining or retransmitting location
- o Key points:
 - * LRs that do not behave in ways that risk privacy are Geopriv-compliant by default. No further action is necessary.

6.2. Location-based Web Services

Many location-based services are delivered over the Web, using Javascript code to orchestrate a series of HTTP requests for location specific information. To support these applications, browser extensions have been developed that support Device-based positioning (manual entry and Global Positioning System (GPS)) and network-assisted positioning (via Assisted GPS (AGPS), and multilateration with 802.11 and cellular signals), exposing location to web pages through Javascript APIs.

In this scenario, we consider a Target that uses a browser with a network-assisted positioning extension. When the Target uses this browser to request location-based services from a web page, the

browser prompts the user to grant the page permission to access the user's location. If the user grants permission, the browser extension sends 802.11 signal strength measurements to a positioning server, which then returns the position of the host. The extension constructs an LO with this location and Rules set by the user, then passes the LO to the page through its Javascript API. The page then obtains location-relevant information using an XMLHttpRequest [21] to a server in the same domain as the page and renders this information to the user.

At first blush, this scenario seems much more complicated than the minimal scenario above. However, most of the privacy considerations are actually the same.

The positioning phase in this scenario begins when the browser extension contacts the positioning server. The positioning server acts as an LG.

The distribution phase actually occurs entirely within the Target host. This phase begins when the positioning server, now acting as LS, follows the LCP policy by providing location only to the Target. The next hop in distribution occurs when the browser extension (an entity under the control of the Target) passes an LO to the web page (an entity under the control of its author). In this phase, the browser extension acts as an LS, with the Target as the sole Rule Maker; the user interface for rule-making is effectively a protocol for conveying Rules, and the extension's API effectively defines a way to communicate LOs and an LO Format. The web site acts as an LR when the web page accepts the LO.

The use phase encompasses the web site's use of the LO. In this context, the phrase "web site" encompasses not only the web page, but also the dedicated supporting logic behind it. Considering the entire web site as a recipient, rather than a single page, it becomes clear that sending the LO in an XMLHttpRequest to a back-end server is like passing it to a separate component of the LR (as opposed to retransmitting it to another entity). Thus, even in this case, where location-relevant information is obtained from a back-end server, the LR does not retain or retransmit location, so its behavior is "privacy-safe" -- it doesn't need to interpret the Rules in the LO.

However, consider a variation on this scenario where the web page requests additional information (a map, for instance) from a third-party site. In this case, since location is being transmitted to a third party, the web site (either in the web page or in a back-end server) would need to verify that this transmission is allowed by the LO's Privacy Rules. Similarly, if the site wanted to log the user's location information, then it would need to examine the LO to

determine how long this information can be retained. In such a case, if the LR needs to do something that is not allowed by the Rules, it may have to deny service to the user (hopefully providing a message with the reason). Nonetheless, if the Rules permit retention or retransmission (even if this retransmission is limited by access control rules), then the LR may do so to the extent the Rules allow.

The following outline summarizes this scenario:

- o Positioning: Network-assisted, positioning server=LG
- o Rule installation: RM (=Target) gives permission to sites and sets LO Rules
- o Distribution hop 1: positioning server=LS --> Target, privacy via LCP policy
- o Distribution hop 2: Browser=LS --> Web site=LR, privacy via user confirmation
- o Use: Back-end server delivers location-relevant information without further retransmission, then deletes location; privacy via safe behavior
- o Key points:
 - * Privacy in this scenario is provided by a combination of explicit user direction and Rules in an LO
 - * Distribution can occur within a host, between mutually untrusting components
 - * Some transmissions of location are actually internal to an LR
 - * LRs that do things that might be constrained by Rules need to verify that these actions are allowed for a particular LO

6.3. Emergency Calling

Support for emergency calls by Voice-over-IP devices is a critical use case for location information about Internet hosts. The details of the Internet architecture for emergency calling are described in [22][23]. In this architecture, there are three critical steps in the placement of an emergency call, each involving location information:

1. Determine the location of the caller

2. Determine the proper Public Safety Answering Point (PSAP) for the caller's location
3. Send a SIP INVITE message (including the caller's location) to the PSAP

The first step in an emergency call is to determine the location of the caller. This step is the positioning phase of the location life-cycle. Location is determined by whatever means are available to the caller's device, or to the network, if this step is being done by a proxy. Whichever entity does the positioning (either the caller or a proxy) acts as an LS, preserving the privacy of location information by only including it in emergency calls.

The second step in an emergency call encompasses location distribution and use. The entity that is routing the emergency call sends location through the LoST protocol [15] to a mapping server. In this role, the routing entity acts as an LS and the LoST server acts as an LR. The LO format within LoST does not allow Rules to be sent along with location, but because LoST is an application-specific protocol, the sending of location within a LoST message authorizes the LoST server to use the location to complete the protocol, namely to route the message as necessary through the LoST mapping architecture [24]. That is, the LoST server is authorized to complete the LoST protocol, but to do nothing else.

The third step in an emergency call is again a combination of distribution and use. The caller (or another entity that inserts the caller's location) acts as an LS and the PSAP acts as an LR. In this specific example, the caller's location is transmitted either as a PIDF-LO object or as a reference that returns a PIDF-LO (or both); in the latter case, the reference should be appropriately protected so that only the PSAP has access. In any case, the receipt of an LO implies that the PSAP should obey the Rules in those LOs in order to preserve privacy. Depending on the regulatory environment, the PSAP may have the option to ignore those constraints in order to respond to an emergency, or it may be bound to respect these Rules (in spite of the emergency situation).

The following outline summarizes this scenario:

- o Positioning: Any
- o Distribution/use hop 1: Target=LS --> LoST infrastructure (no Rules), privacy via authorization implicit in protocol
- o Distribution/use hop 2: Target=LS --> PSAP, privacy via Rules in LO

- o Use: PSAP uses location to deliver emergency services
- o Key points:
 - * Privacy in this scenario is provided by a combination of explicit user direction, implicit authorization particular to a protocol, and Rules in an LO
 - * LRs may be constrained to respect or ignore Privacy Rules by local regulation

6.4. Combination of Services

In modern Internet applications, users frequently receive information via one channel and broadcast it via another. In this sense, both users and channels (e.g., web services) become LSess. Here we consider a more complex example that illustrates this pattern across multiple logical hops.

Suppose Alice (the Target) subscribes to a wireless ISP that determines her location using a network-based positioning technique (e.g., via the location of the base station serving the Target), and provides that information directly to a location-enhanced presence provider (which might use SIP, XMPP [25], or another protocol). The location-enhanced presence provider allows Alice to specify Rules for how this location is distributed: which friends should receive Alice's location and what Rules they should get with it. Alice uses a few other location-enhanced services as well, so she sends Rules that allow her location to be shared with those services, and allow those services to retain and retransmit her location.

Bob is one of Alice's friends, and he receives her location via this location-enhanced presence service. Noting that she's at their favorite coffee shop, Bob wants to upload a photo of the two of them at the coffee shop to a photo-sharing site, along with an LO that marks the location. Bob checks the Rules in Alice's LO and verifies that the photo sharing site is one of the services that Alice authorized. Seeing that Alice has authorized him to give the LO to the photo-sharing site, he attaches it to the photo and uploads it.

Once the geo-tagged photo is uploaded, the photo sharing site reads the Rules in the LO and verifies that the site is authorized to store the photo and to share it with others. Since Alice has allowed the site to retransmit and retain without any constraints, the site fulfills Bob's request to make the geo-tagged photo publicly accessible.

Eve, another user of the photo sharing site, downloads the photo of

Alice and Bob at the coffee shop and receives Alice's LO along with it. Eve posts the photo and location to her public page on a social networking site without checking the Rules, even though the LO doesn't allow Eve to send the location anywhere else. The social networking site, however, observes that no retransmission or retention are allowed (both of which it needs for a public posting), and rejects the upload.

In terms of the location life-cycle, this scenario consists of a positioning step, followed by four distribution hops and use. Positioning is the simplest step: An LG in Alice's ISP monitors her location and transmits it to the presence service, maintaining privacy by only transmitting location to a single entity (to which Alice has delegated privacy responsibilities).

The first distribution hop occurs when the presence server sends location to Bob. In this transaction, the presence server acts as an LS, Alice acts as an RM, and Bob acts as an LR. The privacy of this transaction is assured by the fact that Alice has installed Rules on the presence server that dictate who it may allow to access her location. The second distribution hop is when Bob uploads the LO to the photo-sharing site. Here Bob acts as an LS, preserving the privacy of location information by verifying that the Rules in the LO allow him to upload it. The third distribution hop is when the photo-sharing site sends the LO to Eve, likewise following the Rules -- but a different set of Rules than Bob, since an LO can specify different Rule sets for different LSes.

Eve is the fourth LS in the chain, and fails to comply with Geopriv by not checking the Rules in the LO prior to uploading the LO to the social networking site. The site, however, is a responsible LR -- it checks the Rules in the LO, sees that they don't allow it to use the location as it needs to, and discards the LO.

The following outline summarizes this scenario:

- o Positioning: Network-based, LG in network, privacy via exclusive relationship with presence service
- o Distribution/use hop 1: Presence server --> Bob, privacy via Alice's access control rules
- o Distribution/use hop 2: Bob --> photo sharing site, privacy via Rules for Bob in LO
- o Distribution/use hop 3: Photo sharing site --> Eve, privacy via Rules for site in LO

- o Distribution/use hop 4: Eve --> Social networking site, violates privacy by retransmitting
- o Use: Social networking site, privacy via checking Rules and discarding
- o Key points:
 - * Privacy can be preserved through multiple hops
 - * A LO can specify different Rules for different entities
 - * An LS can still disobey the Rules, but even then, the architecture still works in some cases

7. Glossary

Various security-related terms not defined here are to be understood in the sense defined in RFC 4949 [26].

\$ Access Control Rule

A rule that describe which entities may receive location information and in what form.

\$ civic location

The geographic position of an entity in terms of a postal address or civic landmark. Examples of such data are room number, street number, street name, city, ZIP code, county, state and country.

\$ Device

The physical device whose location is tracked as a proxy for the location of a Target.

\$ geodetic location

The geographic position of an entity in a particular coordinate system (for example, a latitude-longitude pair).

\$ Local Rule

A Privacy Rules that directs a Location Server about how to treat a Target's location information. Local Rules are used internally by a Location Server to handle requests from Location Recipients. They are not distributed to Location Recipients.

\$ Location Generator (LG)

Performs the role of initially determining or gathering the location of a Target. Location Generators may be any sort of software or hardware used to obtain a Target's location (examples include GPS chips and cellular networks).

\$ Location Information Server (LIS)

An entity responsible for providing devices within an access network with information about their own locations. A Location Information Server uses knowledge of the access network and its physical topology to generate and distribute location information to devices.

\$ Location Object (LO)

A data unit that conveys location information together with Privacy Rules within the Geopriv architecture. A Location Object may convey geodetic location data (latitude/longitude/altitude), civic location data (street/city/state/etc.), or both.

\$ Location Recipient (LR)

An ultimate end point entity to which a Location Object is distributed. Location Recipients request location information about a particular Target from a Location Server. If allowed by the appropriate Privacy Rules, a Location Recipient will receive Location Objects describing the Target's location from the Location Server.

\$ Location Server (LS)

An entity that receives Location Objects from Location Generators, Privacy Rules from Rule Makers, and location requests from Location Recipients. A Location Server applies the appropriate Privacy Rules to a Location Object received from a Location Generator and may disclose the Location Object, in compliance with the Rules, to Location Recipients.

Location Servers may not necessarily be "servers" in the colloquial sense of hosts in remote data centers servicing requests. Rather, a Location Server can be any software or hardware component that receives and distributes location information. Examples include a positioning server (with a location interface) in an access network, a presence server, or a Web browser or other software running on a Target's device.

\$ Privacy Rule

A directive that regulates an entity's activities with respect to a Target's location information, including the collection, use, disclosure, and retention of the location information. Privacy Rules describe how location information may be used by an entity, the level of detail with which location information may be described to an entity, and the conditions under which location information may be disclosed to an entity. Privacy Rules are communicated from Rule Makers to Location Servers and conveyed in Location Objects throughout the Geopriv architecture.

\$ Rule

See Privacy Rule.

\$ Rule Maker (RM)

An individual or entity that is authorized to set Privacy Rules for a Target. In some cases a Rule Maker and a Target will be the same individual or entity, and in other cases they will be separate. For example, a parent may serve as the Rule Maker when the Target is his child. The Rule Maker is also not necessarily the owner of a Target device. For example, a corporation may own a device that it provides to an employee but permit the employee to serve as the Rule Maker and set her own Privacy Rules. Rule Makers provide the Privacy Rules associated with a Target to Location Servers.

\$ Forwarded Rule

A Privacy Rule that travels inside a Location Object. Forwarded Rules direct Location Recipients about how to handle the location information they receive. Because the Forwarded Rules themselves may reveal potentially sensitive information about a Target, only the minimal subset of Forwarded Rules necessary for a Location Recipient to handle a Location Object is distributed to the Location Recipient.

\$ Target

An individual or other entity whose location is sought in the Geopriv architecture. In many cases the Target will be the human user of a Device, or it may be an object such as a vehicle or shipping container to which a Device is attached. In some instances the Target will be the Device itself. The Target is the entity whose privacy Geopriv seeks to protect.

\$ Usage Rule

A rule that describe what uses of location information are authorized.

8. Acknowledgements

Section 5 is largely based on the security investigations conducted as part of the Geopriv Layer-7 Location Configuration Protocol design team, which produced [9]. We would like to thank all the members of the design team.

We would also like to thank Marc Linsner and Martin Thomson for their contributions regarding terminology and LCPs.

9. IANA Considerations

This document makes no request of IANA.

10. References

10.1. Normative References

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

10.2. Informative References

- [2] Cuellar, J., Morris, J., Mulligan, D., Peterson, J., and J. Polk, "Geopriv Requirements", RFC 3693, February 2004.
- [3] Danley, M., Mulligan, D., Morris, J., and J. Peterson, "Threat Analysis of the Geopriv Protocol", RFC 3694, February 2004.
- [4] U.S. Department of Defense, "National Industrial Security Program Operating Manual", DoD 5220-22M, January 1995.
- [5] 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.
- [6] 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.

- [7] 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-21 (work in progress), January 2010.
- [8] Rosenberg, J., "The Extensible Markup Language (XML) Configuration Access Protocol (XCAP)", RFC 4825, May 2007.
- [9] Tschofenig, H. and H. Schulzrinne, "GEOPRIV Layer 7 Location Configuration Protocol: Problem Statement and Requirements", RFC 5687, March 2010.
- [10] Polk, J., Schnizlein, J., and M. Linsner, "Dynamic Host Configuration Protocol Option for Coordinate-based Location Configuration Information", RFC 3825, July 2004.
- [11] Schulzrinne, H., "Dynamic Host Configuration Protocol (DHCPv4 and DHCPv6) Option for Civic Addresses Configuration Information", RFC 4776, November 2006.
- [12] 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-08 (work in progress), July 2010.
- [13] Barnes, M., "HTTP-Enabled Location Delivery (HELD)", RFC 5985, September 2010.
- [14] Marshall, R., "Requirements for a Location-by-Reference Mechanism", RFC 5808, May 2010.
- [15] Hardie, T., Newton, A., Schulzrinne, H., and H. Tschofenig, "LoST: A Location-to-Service Translation Protocol", RFC 5222, August 2008.
- [16] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, December 2005.
- [17] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, August 2008.
- [18] Peterson, J., "A Presence-based GEOPRIV Location Object Format", RFC 4119, December 2005.
- [19] 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.

- [20] 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.
- [21] World Wide Web Consortium, "The XMLHttpRequest Object", W3C document <http://www.w3.org/TR/XMLHttpRequest/>, April 2008.
- [22] Rosen, B., Schulzrinne, H., Polk, J., and A. Newton, "Framework for Emergency Calling using Internet Multimedia", draft-ietf-ecrit-framework-11 (work in progress), July 2010.
- [23] Rosen, B. and J. Polk, "Best Current Practice for Communications Services in support of Emergency Calling", draft-ietf-ecrit-phonebcf-15 (work in progress), July 2010.
- [24] Schulzrinne, H., "Location-to-URL Mapping Architecture and Framework", draft-ietf-ecrit-mapping-arch-04 (work in progress), March 2009.
- [25] Saint-Andre, P., Ed., "Extensible Messaging and Presence Protocol (XMPP): Core", RFC 3920, October 2004.
- [26] Shirey, R., "Internet Security Glossary, Version 2", RFC 4949, August 2007.
- [27] Polk, J. and B. Rosen, "Location Conveyance for the Session Initiation Protocol", draft-ietf-sip-location-conveyance-13 (work in progress), March 2009.

URIs

- [28] <<http://creativecommons.org/>>

Authors' Addresses

Richard Barnes
BBN Technologies
9861 Broken Land Pkwy, Suite 400
Columbia, MD 21046
USA

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

Matt Lepinski
BBN Technologies
10 Moulton St
Cambridge, MA 02138
USA

Phone: +1 617 873 5939
Email: mlepinski@bbn.com

Alissa Cooper
Center for Democracy & Technology
1634 I Street NW, Suite 1100
Washington, DC
USA

Email: acooper@cdt.org

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

Email: jmorris@cdt.org

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>

GEOPRIV Working Group
INTERNET-DRAFT
Obsoletes: 3825 (if approved)
Category: Standards Track
Expires: August 26, 2011
26 February 2011

J. Polk
M. Linsner
Cisco Systems
M. Thomson
Andrew Corporation
B. Aboba (ed)
Microsoft Corporation

Dynamic Host Configuration Protocol Options for
Coordinate-based Location Configuration Information

draft-ietf-geopriv-rfc3825bis-17.txt

Abstract

This document specifies Dynamic Host Configuration Protocol Options (both DHCPv4 and DHCPv6) for the coordinate-based geographic location of the client. The Location Configuration Information (LCI) includes Latitude, Longitude, and Altitude, with resolution or uncertainty indicators for each. Separate parameters indicate the reference datum for each of these values. This document obsoletes RFC 3825.

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 August 26, 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 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.	Introduction	4
1.1.	Conventions	5
1.2.	Resolution and Uncertainty	5
2.	DHCP Option Formats	6
2.1.	DHCPv6 GeoLoc Option	6
2.2.	DHCPv4 Options	8
2.3.	Latitude and Longitude Fields	11
2.4.	Altitude	14
2.5.	Datum	16
3.	Security Considerations.	17
4.	IANA Considerations.	17
4.1.	DHCP Options	17
4.2.	Altitude Type Registry	18
4.3.	Datum Registry	18
4.4.	GeoLoc Option Version Registry	19
5.	Acknowledgments	20
6.	References	20
6.1.	Normative References	20
6.2.	Informational References	21
Appendix A.	GML Mapping	23
A.1.	GML Templates	23
Appendix B.	Calculations of Resolution	26
B.1.	LCI of "White House" (Example 1)	27
B.2.	LCI of "Sears Tower" (Example 2)	29
Appendix C.	Calculations of Uncertainty	30
C.1.	LCI of "Sydney Opera House" (Example 3)	30
Appendix D.	Changes from RFC 3825	34
Authors' Addresses	35

1. Introduction

The physical location of a network device has a range of applications. In particular, emergency telephony applications rely on knowing the location of a caller in order to determine the correct emergency center.

The location of a device can be represented either in terms of geospatial (or geodetic) coordinates, or as a civic address. Different applications may be more suited to one form of location information; therefore, both the geodetic and civic forms may be used simultaneously.

This document specifies Dynamic Host Configuration Protocol v4 (DHCPv4) [RFC2131] and DHCPv6 [RFC3315] options for the coordinate-based geographic location of the client, to be provided by the server. "Dynamic Host Configuration Protocol (DHCPv4 and DHCPv6) Option for Civic Addresses Configuration Information" [RFC4776] specifies DHCP options for civic addresses.

The geodetic coordinate options defined in this document and the civic address options defined in RFC 4776 [RFC4776] enable a DHCP client to obtain its location. For example, a wired Ethernet host might use these options for location determination. In this case, the location information could be derived from a wiremap by the DHCP server, using the Circuit-ID Relay Agent Information Option (RAIO) defined (as Sub-Option 1) in RFC 3046 [RFC3046]. The DHCP server could correlate the Circuit-ID with the geographic location where the identified circuit terminates (such as the location of the wall jack).

The mechanism defined here may also be utilized to provide location to wireless hosts. DHCP relay agent sub-options (RAIO) [RFC3046] is one method a DHCP server might use to perform host location determination. Currently, the relay agent sub-options do not include data sets required for device level location determination of wireless hosts. In cases where the DHC server uses RAIO for location determination, a wireless host can use this mechanism to discover location of the radio access point, or the area of coverage for the radio access point.

An important feature of this specification is that after the relevant DHCP exchanges have taken place, the location information is stored on the end device rather than somewhere else, where retrieving it might be difficult in practice.

1.1. 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].

1.2. Resolution and Uncertainty

The DHCP options defined in this document include fields quantifying the resolution or uncertainty associated with a target location. No inferences relating to privacy policies can be drawn from either uncertainty or resolution values.

As utilized in this document, resolution refers to the accuracy of a reported location, as expressed by the number of valid bits in each of the Latitude, Longitude and Altitude fields.

In the context of location technology, uncertainty is a quantification of errors. Any method for determining location is subject to some sources of error; uncertainty describes the amount of error that is present. Uncertainty might be the coverage area of a wireless transmitter, the extent of a building or a single room.

Uncertainty is usually represented as an area within which the target is located. In this document, each of the three axes can be assigned an uncertainty value. In effect, this describes a rectangular prism, which may be used as a coarse representation of a more complex shape that fits within it. See Section 2.3.2 for more detail on the correspondence between shapes and uncertainty.

When representing locations from sources that can quantify uncertainty, the goal is to find the smallest possible rectangular prism that this format can describe. This is achieved by taking the minimum and maximum values on each axis and ensuring that the final encoding covers these points. This increases the region of uncertainty, but ensures that the region that is described encompasses the target location.

The DHCPv4 option formats defined in this document support resolution and uncertainty parameters. The DHCPv4 GeoConf Option 123 includes a resolution parameter for each of the dimensions of location. Since this resolution parameter need not apply to all dimensions equally, a resolution value is included for each of the three location elements. The DHCPv4 GeoLoc Option TBD1 as well as the DHCPv6 GeoLoc Option TBD2 format utilize an uncertainty parameter.

Appendix A describes the mapping of DHCP option values to the Geography Markup Language (GML). Appendix B of this document

provides examples showing the calculation of resolution values. Appendix C provides an example demonstrating calculation of uncertainty values.

Since the Presence Information Data Format Location Object (PIDF-LO) [RFC4119][RFC5491] is used to conveying location and the associated uncertainty within an emergency call [Convey], a mechanism is needed to convert the information contained within the DHCPv4 and DHCPv6 options to PIDF-LO. This document describes the following conversions:

DHCPv4 GeoConf Option 123 to PIDF-LO
 DHCPv4 GeoLoc Option TBD1 and DHCPv6 GeoLoc Option TBD2 to PIDF-LO
 PIDF-LO to DHCP GeoLoc Option TBD1 and DHCPv6 GeoLoc Option TBD2

Conversion to PIDF-LO does not increase uncertainty; conversion from PIDF-LO to the DHCPv4 GeoLoc Option TBD1 and the DHCPv6 GeoLoc Option TBD2 increases uncertainty by less than a factor of 2 in each dimension. Since it is not possible to translate an arbitrary PIDF-LO to the DHCP GeoConf Option 123 with a bounded increase in uncertainty, the conversion is not specified.

2. DHCP Option Formats

This section defines the format for the DHCPv4 and DHCPv6 options. These options utilize a similar format, differing primarily in the option code.

2.1. DHCPv6 GeoLoc Option

The format of the DHCPv6 [RFC3315] GeoLoc Option is as follows:

```

      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
+-----+-----+-----+-----+-----+-----+-----+-----+
|          Option Code (TBD2)          |          OptLen          |
+-----+-----+-----+-----+-----+-----+-----+-----+
|  LatUnc  |          Latitude          |                         |
+-----+-----+-----+-----+-----+-----+-----+-----+
| Lat (cont'd) |  LongUnc  |          Longitude          |         +
+-----+-----+-----+-----+-----+-----+-----+-----+
|          Longitude (cont'd)          | AType |  AltUnc  |  Altitude +
+-----+-----+-----+-----+-----+-----+-----+-----+
|          Altitude (cont'd)          | Ver | Res | Datum |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

Code: DHCP Option Code TBD2 (16 bits).

OptLen: Option Length. For version 1, the option length is 16.

LatUnc: 6 bits. When the Ver field = 1, this field represents latitude uncertainty. The contents of this field is undefined for other values of the Ver field.

Latitude: a 34 bit fixed point value consisting of 9 bits of integer and 25 bits of fraction, interpreted as described in Section 2.3.

LongUnc: 6 bits. When the Ver field = 1, this field represents longitude uncertainty. The contents of this field is undefined for other values of the Ver field.

Longitude: a 34 bit fixed point value consisting of 9 bits of integer and 25 bits of fraction, interpreted as described in Section 2.3.

AType: Altitude Type (4 bits), defined in Section 2.4.

AltUnc: 6 bits. When the Ver field = 1, this field represents altitude uncertainty. The contents of this field is undefined for other values of the Ver field.

Altitude: A 30 bit value defined by the AType field, described in Section 2.4.

Ver: The Ver field is two bits, providing for four potential versions. This specification defines the behavior of version 1. The Ver field is always located at the same offset from the beginning of the option, regardless of the version in use. DHCPv6 clients implementing this specification MUST support receiving version 1 responses. DHCPv6 servers implementing this specification MUST send version 1 responses.

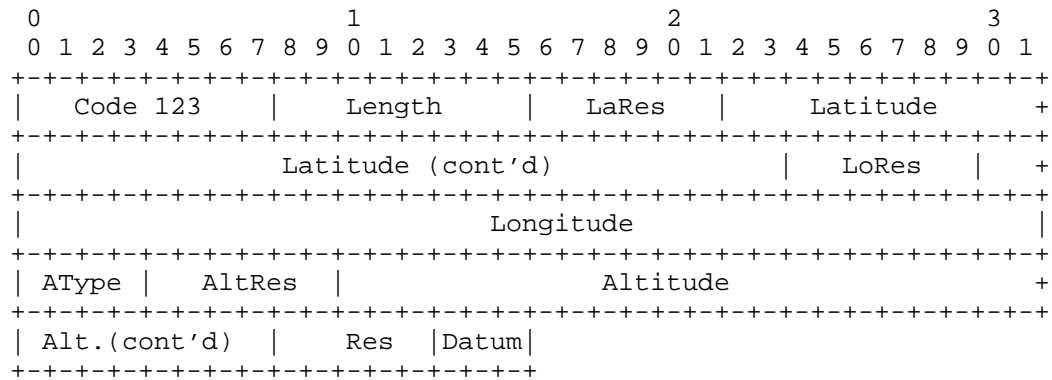
Res: The Res field which is 3 bits, is reserved. These bits have been used by [IEEE-802.11y], but are not defined within this specification.

Datum: 3 bits. The Map Datum used for the coordinates given in this Option.

2.2. DHCPv4 Options

2.2.1. DHCPv4 GeoConf Option

The format of the DHCPv4 GeoConf Option is as follows:



Code: 8 bits. The code for the DHCPv4 GeoConf Option (123).

Length: 8 bits. The length of the option, in octets.
The option length is 16.

LaRes: 6 bits. This field represents latitude resolution.

Latitude: a 34 bit fixed point value consisting of 9 bits of signed integer and 25 bits of fraction, interpreted as described in Section 2.3.

LoRes: 6 bits. This field represents longitude resolution.

Longitude: a 34 bit fixed point value consisting of 9 bits of signed integer and 25 bits of fraction, interpreted as described in Section 2.3.

AType: Altitude Type (4 bits), defined in Section 2.4.

AltRes: 6 bits. This field represents altitude resolution.

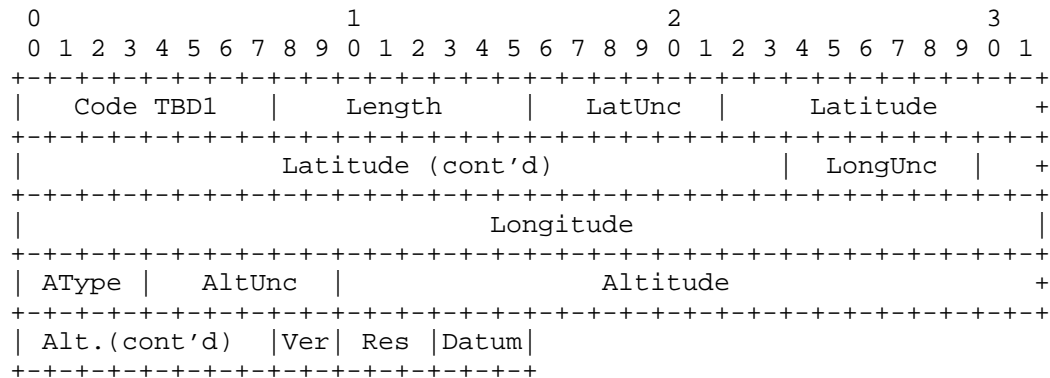
Altitude: A 30 bit value defined by the AType field, described in Section 2.4.

Res: The Res field which is 5 bits, is reserved. These bits have been used by IEEE 802.11y [IEEE-802.11y], but are not defined within this specification.

Datum: 3 bits. The Map Datum used for the coordinates given in this Option.

2.2.2. DHCPv4 GeoLoc Option

The format of DHCPv4 GeoLoc Option is as follows:



Code: 8 bits. The code for the DHCPv4 GeoLoc Option (TBD1).

Length: 8 bits. The length of the option, in octets.
For version 1, the option length is 16.

LatUnc: 6 bits. When the Ver field = 1, this field represents latitude uncertainty. The contents of this field is undefined for other values of the Ver field.

Latitude: a 34 bit fixed point value consisting of 9 bits of integer and 25 bits of fraction, interpreted as described in Section 2.3.

LongUnc: 6 bits. When the Ver field = 1, this field represents longitude uncertainty. The contents of this field is undefined for other values of the Ver field.

Longitude: a 34 bit fixed point value consisting of 9 bits of integer and 25 bits of fraction, interpreted as described in Section 2.3.

AType: Altitude Type (4 bits), defined in Section 2.4.

AltUnc: 6 bits. When the Ver field = 1, this field represents altitude uncertainty. The contents of this field is undefined for other values of the Ver field.

Altitude: A 30 bit value defined by the AType field, described in Section 2.4.

Ver: The Ver field is two bits, providing for four potential versions. This specification defines the behavior of version 1. The Ver field is always located at the same offset from the beginning of the option, regardless of the version in use.

Res: The Res field which is 3 bits, is reserved. These bits have been used by [IEEE-802.11y], but are not defined within this specification.

Datum: 3 bits. The Map Datum used for the coordinates given in this Option.

2.2.3. Option Support

2.2.3.1. Client Support

DHCPv4 clients implementing this specification MUST support receiving the DHCPv4 GeoLoc Option TBD1 (version 1), and MAY support receiving the DHCPv4 GeoConf Option 123 (originally defined in RFC 3825 [RFC3825]).

DHCPv4 clients request the DHCPv4 server to send GeoConf Option 123, GeoLoc Option TBD1 or both via inclusion of the Parameter Request List option. As noted in Section 9.8 of RFC 2132 [RFC2132]:

This option is used by a DHCP client to request values for specified configuration parameters. The list of requested parameters is specified as n octets, where each octet is a valid DHCP option code as defined in this document.

The client MAY list the options in order of preference. The DHCP server is not required to return the options in the requested order, but MUST try to insert the requested options in the order requested by the client.

When DHCPv4 and DHCPv6 clients implementing this specification do not understand a datum value, they MUST assume a World Geodesic System 1984 (WGS84) [WGS84] datum (EPSG [EPSG] 4326 or 4979, depending on whether there is an Altitude value present) and proceed accordingly. Assuming that a less accurate location value is better than none, this ensures that some (perhaps less accurate) location is available to the client.

2.2.3.2. Server Option Selection

A DHCPv4 server implementing this specification **MUST** support sending GeoLoc Option TBD1 version 1 and **SHOULD** support sending GeoConf Option 123 in responses.

A DHCPv4 server that provides location information **SHOULD** honor the Parameter Request List included by the DHCPv4 client in order to decide whether to send GeoConf Option 123, GeoLoc Option TBD1 or both in the Response.

2.3. Latitude and Longitude Fields

The Latitude and Longitude values in this specification are encoded as 34 bit, twos complement, fixed point values with 9 integer bits and 25 fractional bits. The exact meaning of these values is determined by the datum; the description in this section applies to the datums defined in this document. This document uses the same definition for all datums it specifies.

When encoding, Latitude and Longitude values are rounded to the nearest 34-bit binary representation. This imprecision is considered acceptable for the purposes to which this form is intended to be applied and is ignored when decoding.

Positive latitudes are north of the equator and negative latitudes are south of the equator. Positive longitudes are east of the Prime Meridian (Greenwich) and negative (2s complement) longitudes are west of the Prime Meridian.

Within the coordinate reference systems defined in this document (Datum values 1-3), longitude values outside the range of -180 to 180 decimal degrees or latitude values outside the range of -90 to 90 degrees **MUST** be considered invalid. Server implementations **SHOULD** prevent the entry of invalid values within the selected coordinate reference system. Location consumers **MUST** ignore invalid location coordinates and **SHOULD** log invalid location errors.

2.3.1. Latitude and Longitude Resolution

The Latitude (LaRes), Longitude (LoRes) and Altitude (AltRes) Resolution fields are encoded as 6 bit, unsigned integer values. In the DHCPv4 GeoConf Option 123, the LaRes, LoRes and AltRes fields are used to encode the number of bits of resolution. The resolution sub-fields accommodate the desire to easily adjust the precision of a reported location. Contents beyond the claimed resolution **MAY** be randomized to obscure greater precision that might be available.


In the DHCPv4 GeoConf Option 123, the LaRes value encodes the number of high-order latitude bits that should be considered valid. Any bits entered to the right of this limit should not be considered valid and might be purposely false, or zeroed by the sender. The examples in Appendix B illustrate that a smaller value in the resolution field increases the area within which the device is located. A value of 2 in the LaRes field indicates a precision of no greater than 1/6th that of the globe (see the first example of Appendix B). A value of 34 in the LaRes field indicates a precision of about 3.11 mm in latitude at the equator.

In the DHCPv4 GeoConf Option 123, the LoRes value encodes the number of high-order longitude bits that should be considered valid. Any bits entered to the right of this limit should not be considered valid and might be purposely false, or zeroed by the sender. A value of 2 in the LoRes field indicates precision of no greater than 1/6th that of the globe (see the first example of Appendix B). A value of 34 in the LoRes field indicates a precision of about 2.42 mm in Longitude (at the equator). Because lines of longitude converge at the poles, the distance is smaller (better precision) for locations away from the equator.

2.3.2. Latitude and Longitude Uncertainty

In the DHCPv6 GeoLoc Option TBD2 and the DHCPv4 GeoLoc Option TBD1, the Latitude and Longitude Uncertainty fields (LatUnc and LongUnc) quantify the amount of uncertainty in each of the Latitude and Longitude values respectively. A value of 0 is reserved to indicate that the uncertainty is unknown; values greater than 34 are reserved.

A point within the region of uncertainty is selected to be the encoded point; the centroid of the region is often an appropriate choice. The value for uncertainty is taken as the distance from the selected point to the furthest extreme of the region of uncertainty on that axis. This is demonstrated in the figure below, which shows a two-dimensional polygon that is projected on each axis. In the figure, "X" marks the point that is selected; the ranges marked with "U" is the uncertainty.



Key

V, ^ = vertical arrows, delimiting the vertical uncertainty range.
 <> = horizontal arrows, delimiting the horizontal uncertainty range.

Uncertainty applies to each axis independently.

The amount of uncertainty can be determined from the encoding by taking 2 to the power of 8, less the encoded value. As is shown in the following formula, where "x" is the encoded integer value:

$$\text{uncertainty} = 2 ^ (8 - x)$$

The result of this formula is expressed in degrees of latitude or longitude. The uncertainty is added to the base latitude or longitude value to determine the maximum value in the uncertainty range; similarly, the uncertainty is subtracted from the base value to determine the minimum value. Note that because lines of longitude converge at the poles, the actual distance represented by this uncertainty changes with the distance from the equator.

If the maximum or minimum latitude values derived from applying uncertainty are outside the range of -90 to +90, these values are trimmed to within this range. If the maximum or minimum longitude values derived from applying uncertainty are outside the range of -180 to +180, then these values are normalized to this range by adding or subtracting 360 as necessary.

The encoded value is determined by subtracting the next highest whole

integer value for the base 2 logarithm of uncertainty from 8. As is shown by the following formula, where uncertainty is the midpoint of the known range less the lower bound of that range:

$$x = 8 - \text{ceil}(\log_2(\text{uncertainty}))$$

Note that the result of encoding this value increases the range of uncertainty to the next available power of two; subsequent repeated encodings and decodings do not change the value. Only increasing uncertainty means that the associated confidence does not have to decrease.

2.4. Altitude

How the Altitude value is interpreted depends on the Altitude Type (AType) value and the selected datum. Three Altitude Type values are defined in this document: unknown (0), meters (1) and floors (2).

2.4.1. No Known Altitude (AType = 0)

In some cases, the altitude of the location might not be provided. An Altitude Type value of zero indicates that the altitude is not given to the client. In this case, the Altitude and Altitude Uncertainty fields can contain any value and MUST be ignored.

2.4.2. Altitude in Meters (AType = 1)

If the Altitude Type has a value of one, Altitude is measured in meters, in relation to the zero set by the vertical datum. For AType = 1, the Altitude value is expressed as a 30 bit, fixed point, twos complement integer with 22 integer bits and 8 fractional bits.

2.4.3. Altitude in Floors (AType = 2)

A value of two for Altitude Type indicates that the Altitude value is measured in floors. Since altitude in meters may not be known within a building, a floor indication may be more useful. For AType = 2, the Altitude value is expressed as a 30 bit, fixed point, twos complement integer with 22 integer bits and 8 fractional bits.

This value is relevant only in relation to a building; the value is relative to the ground level of the building. Floors located below ground level are represented by negative values. In some buildings it might not be clear which floor is at ground level or an intermediate floor might be hard to identify as such. Determining what floor is at ground level and what constitutes a sub-floor as opposed to an naturally numbered floor is left to local interpretation.

Larger values represent floors that are farther away from floor 0 such that:

- if positive, the floor value is farther above the ground floor.
- if negative, the floor value is farther below the ground floor.

Non-integer values can be used to represent intermediate or sub-floors, such as mezzanine levels. Example: a mezzanine between floor 1 and floor 2 could be represented as a value of 1.25. Example: mezzanines between floor 4 and floor 5 could be represented as values of 4.5 and 4.75.

2.4.4. Altitude Resolution

In the DHCPv4 GeoConf Option 123, the Altitude Resolution (AltRes) value encodes the number of high-order altitude bits that should be considered valid. Values above 30 (decimal) are undefined and reserved.

If the Altitude Type value is one (AType = 1), an AltRes value 0.0 would indicate unknown Altitude. The most precise altitude would have an AltRes value of 30. Many values of AltRes would obscure any variation due to vertical datum differences.

The AltRes field SHOULD be set to maximum precision when AType = 2 (floors) when a floor value is included in the DHCP Reply, or when AType = 0, to denote that the floor isn't known. An altitude coded as AType = 2, AltRes = 30, and Altitude = 0.0 is meaningful even outside a building, and represents ground level at the given latitude and longitude.

2.4.5. Altitude Uncertainty

In the DHCPv6 GeoLoc Option TBD2 or the DHCPv4 GeoLoc Option TBD1, the AltUnc value quantifies the amount of uncertainty in the Altitude value. As with LatUnc and LongUnc, a value of 0 for AltUnc is reserved to indicate that Altitude Uncertainty is not known; values above 30 are also reserved. Altitude Uncertainty only applies to Altitude Type 1.

The amount of Altitude Uncertainty can be determined by the following formula, where x is the encoded integer value:

$$\text{Uncertainty} = 2 ^ { (21 - x)}$$

This value uses the same units as the associated altitude.

Similarly, a value for the encoded integer value can be derived by

the following formula:

$$x = 21 - \text{ceil}(\log_2(\text{uncertainty}))$$

2.5. Datum

The Datum field determines how coordinates are organized and related to the real world. Three datums are defined in this document, based on the definitions in [OGP.Geodesy]:

1: WGS84 (Latitude, Longitude, Altitude):

The World Geodesic System 1984 [WGS84] coordinate reference system.

This datum is identified by the European Petroleum Survey Group (EPSG)/International Association of Oil & Gas Producers (OGP) with the code 4979, or by the URN "urn:ogc:def:crs:EPSG::4979". Without Altitude, this datum is identified by the EPSG/OGP code 4326 and the URN "urn:ogc:def:crs:EPSG::4326".

2: NAD83 (Latitude, Longitude) + NAVD88:

This datum uses a combination of the North American Datum 1983 (NAD83) for horizontal (Latitude and Longitude) values, plus the North American Vertical Datum of 1988 (NAVD88) vertical datum.

This datum is used for referencing location on land (not near tidal water) within North America.

NAD83 is identified by the EPSG/OGP code of 4269, or the URN "urn:ogc:def:crs:EPSG::4269". NAVD88 is identified by the EPSG/OGP code of 5703, or the URN "urn:ogc:def:crs:EPSG::5703".

3: NAD83 (Latitude, Longitude) + MLLW:

This datum uses a combination of the North American Datum 1983 (NAD83) for horizontal (Latitude and Longitude) values, plus the Mean Lower Low Water (MLLW) vertical datum.

This datum is used for referencing location on or near tidal water within North America.

NAD83 is identified by the EPSG/OGP code of 4269, or the URN "urn:ogc:def:crs:EPSG::4269". MLLW does not have a specific code or URN.

All hosts MUST support the WGS84 datum (Datum 1).

3. Security Considerations

Geopriv requirements (including security requirements) are discussed in "Geopriv Requirements" [RFC3693]. A threat analysis is provided in "Threat Analysis of the Geopriv Protocol" [RFC3694].

Since there is no privacy protection for DHCP messages, an eavesdropper who can monitor the link between the DHCP server and requesting client can discover this LCI.

To minimize the unintended exposure of location information, the LCI option SHOULD be returned by DHCP servers only when the DHCP client has included this option in its 'parameter request list' (Section 3.5 [RFC2131], Section 9.8 [RFC2132]).

Where critical decisions might be based on the value of this option, DHCP authentication as defined in "Authentication for DHCP Messages" [RFC3118] and "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)" [RFC3315] SHOULD be used to protect the integrity of the DHCP options.

Link layer confidentiality and integrity protection may also be employed to reduce the risk of location disclosure and tampering.

4. IANA Considerations

4.1. DHCP Options

This document defines the DHCPv6 GeoLoc option (see Section 2.1) which requires assignment of DHCPv6 option code TBD2 [RFC3315]:

Value	Description	Reference
----	-----	-----
TBD2	OPTION_GEOLOCATION	RFC xxxx
[RFC Editor: Please replace xxxx with the RFC number assigned to this document.]		

This document defines the DHCPv4 GeoConf option (see Section 2.2.1) which has been assigned a DHCPv4 option code of 123 from the DHCP Option space.

This document also defines the DHCPv4 GeoLoc option (see Section 2.2.2) which requires assignment of DHCPv4 option code TBD1 [RFC2132][RFC2939]:

Tag	Name	Data Length	Meaning	Reference
TBD1	GeoLoc	16	Geospatial Location with Uncertainty	RFC xxxx

[RFC Editor: Please replace xxxx with the RFC number assigned to this document.]

4.2. Altitude Type Registry

IANA is asked to create and maintain the Altitude Type registry following the guidelines below.

The registry consists of three values: Altitude Type, Description and Reference. These are described below.

Altitude Type: an integer, refers to the value used in the DHCPv4 GeoConf and the DHCPv4 and DHCPv6 GeoLoc Options described in this document. Values from 0 to 15 are assigned.

Description: the description of the altitude described by this code.

Reference: the reference to the document that describes the altitude code. This reference MUST define the way that the 30 bit altitude values and the associated 6 bit uncertainty are interpreted.

Initial values are given below; new assignments are to be made following the "Standards Action" policies [RFC5226].

#	Description	Reference
0	No known altitude	RFC xxxx
1	Altitude in meters	RFC xxxx
2	Altitude in floors	RFC xxxx
3-15	Unassigned	RFC xxxx

[RFC Editor: Please replace xxxx with the RFC number assigned to this document.]

4.3. Datum Registry

IANA is asked to create and maintain the Datum registry following the guidelines below.

The registry consists of three values: Datum, Description and Reference. These are described below.

Datum: an integer, refers to the value used in the DHCPv4 GeoConf and the DHCPv4 and DHCPv6 GeoLoc Options described in this document. Values from 1 to 7 are assigned.

Description: the description of the altitude described by this code.

Reference: the reference to the document that describes the Datum code. This reference MUST include specification of both the horizontal and vertical datum, and MUST define the way that the 34 bit values and the respective 6 bit uncertainties are interpreted.

Initial values are given below; new assignments are to be made following the "Standards Action" policies [RFC5226].

#	Description	Reference
0	Reserved	RFC xxxx
1	Vertical datum WGS 84 defined by EPSG CRS Code 4327	RFC xxxx
2	Vertical datum NAD83 defined by EPSG CRS Code 4269 with North American Vertical Datum of 1988 (NAVD88)	RFC xxxx
3	Vertical datum NAD83, defined by EPSG CRS Code 4269 with Mean Lower Low Water (MLLW) as associated vertical datum	RFC xxxx
4-7	Unassigned	RFC xxxx

[RFC Editor: Please replace xxxx with the RFC number assigned to this document.]

4.4. GeoLoc Option Version Registry

IANA is asked to create and maintain the GeoLoc Option Version registry following the guidelines below.

The registry consists of three values: GeoLoc Option Version, Description and Reference. These are described below.

GeoLoc Option Version: an integer, refers to the version used in the DHCPv4 and DHCPv6 GeoLoc Options described in this document. Values from 1 to 3 are assigned.

Description: the description of the version described by this code.

Reference: the reference to the document that describes the Version code.

Initial values are given below; new assignments are to be made following the "Standards Action" policies [RFC5226].

#	Description	Reference
0	Reserved	RFC xxxx
1	Implementations utilizing uncertainty parameters for both DHCPv4 and DHCPv6 GeoLoc options	RFC xxxx
2-3	Unassigned	RFC xxxx

[RFC Editor: Please replace xxxx with the RFC number assigned to this document.]

5. Acknowledgments

The authors would like to thank Randall Gellens, Patrik Falstrom, Ralph Droms, Ted Hardie, Jon Peterson, Robert Sparks, Ralph Droms, Nadine Abbott and Mykyta Yevstifeyev for their inputs and constructive comments regarding this document. Additionally, the authors would like to thank Greg Troxel for the education on vertical datums, as well as Carl Reed. Thanks to Richard Barnes for his contribution on GML mapping for resolution.

6. References

6.1. Normative References

- [EPSG] European Petroleum Survey Group, <http://www.epsg.org/> and <http://www.epsg-registry.org/>
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2131] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, March 1997.
- [RFC2132] Alexander, S. and R. Droms, "DHCP Options and BOOTP Vendor Extensions", RFC2132, March 1997.
- [RFC2939] Droms, R., "Procedures and IANA Guidelines for Definition of New DHCP Options and Message types", BCP 43, RFC 2939,

September 2000.

- [RFC3046] Patrick, M., "DHCP Relay Agent Information Option", RFC 3046, January 2001.
- [RFC3118] Droms, R. and W. Arbaugh, "Authentication for DHCP Messages", RFC 3118, June 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.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", RFC 5226, May 2008.
- [WGS84] US National Imagery and Mapping Agency, "Department of Defense (DoD) World Geodetic System 1984 (WGS 84), Third Edition", NIMA TR8350.2, January 2000, <https://www1.nga.mil/PRODUCTSSERVICES/GEODESYGEOPHYSICS/WORLDGEODETICSYSTEM/Pages/default.aspx> and <http://www.ngs.noaa.gov/faq.shtml#WGS84>

6.2. Informational References

- [Convey] Polk, J., Rosen, B. and J. Peterson, "Location Conveyance for the Session Initiation Protocol", Internet draft (work in progress), draft-ietf-sipcore-location-conveyance-06.txt, February 23, 2011.
- [GeoShape] Thomson, M. and C. Reed, "GML 3.1.1 PIDF-LO Shape Application Schema for use by the Internet Engineering Task Force (IETF)", Candidate OpenGIS Implementation Specification 06-142, Version: 0.0.9, December 2006.
- [IEEE-802.11y] Information technology - Telecommunications and information exchange between systems - Local and metropolitan area networks - Specific requirements - Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications Amendment 3: 3650-3700 MHz Operation in USA, November 2008.
- [NENA] National Emergency Number Association (NENA) www.nena.org
NENA Technical Information Document on Model Legislation Enhanced 911 for Multi-Line Telephone Systems.
- [RFC3046] Patrick, M., "DHCP Relay Agent Information Option", RFC 3046, January 2001.

- [RFC3693] Cuellar, J., Morris, J., Mulligan, D., Peterson, J. and J. Polk, "Geopriv Requirements", RFC 3693, February 2004.
- [RFC3694] Danley, M., Mulligan, D., Morris, J. and J. Peterson, "Threat Analysis of the Geopriv Protocol", RFC 3694, February 2004.
- [RFC3825] Polk, J., Schnizlein, J. and M. Linsner, "Dynamic Host Configuration Protocol Option for Coordinate-based Location Configuration Information", RFC 3825, July 2004.
- [RFC4119] Peterson, J., "A Presence-based GEOPRIV Location Object Format", RFC 4119, December 2005.
- [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.
- [RFC5491] Winterbottom, J., Thomson, M. and H. Tschafenig, "GEOPRIV PIDF-LO Usage Clarification, Considerations, and Recommendations ", RFC 5491, March 2009

Appendix A. GML Mapping

The GML representation of a decoded DHCP option depends on what fields are specified. The DHCP format for location logically describes a geodetic prism, rectangle, or point, depending on whether Altitude and uncertainty values are provided. In the absence of uncertainty information, the value decoded from the DHCP form can be expressed as a single point; this is true regardless of whether the version 0 or version 1 interpretations of the uncertainty fields are used. If the point includes Altitude, it uses a three dimensional CRS, otherwise it uses a two dimensional CRS. If all fields are included along with uncertainty, the shape described is a rectangular prism. Note that this is necessary given that uncertainty for each axis is provided independently.

If Altitude or Altitude Uncertainty (AltUnc) is not specified, the shape is described as a rectangle using the "gml:Polygon" shape. If Altitude is available, a three dimensional CRS is used, otherwise a two dimensional CRS is used.

For Datum values of 2 or 3 (NAD83), there is no available CRS URN that covers three dimensional coordinates. By necessity, locations described in these datums can be represented by two dimensional shapes only; that is, either a two dimensional point or a polygon.

If the Altitude Type is 2 (floors), then this value can be represented using a civic address object [RFC5139] that is presented alongside the geodetic object.

This Appendix describes how the location value encoded in DHCP format for geodetic location can be expressed in GML. The mapping is valid for the DHCPv6 GeoLoc Option as well as both of the DHCPv4 GeoConf and GeoLoc options, and for the currently-defined datum values (1, 2, and 3). Further version or datum definitions should provide similar mappings.

These shapes can be mapped to GML by first computing the bounds that are described using the coordinate and uncertainty fields, then encoding the result in a GML Polygon or Prism shape.

A.1. GML Templates

If Altitude is provided in meters (AType 1) and the datum value is WGS84 (value 1), then the proper GML shape is a Prism, with the following form (where \$value\$ indicates a value computed from the DHCP option as described below):

```
<gs:Prism srsName="urn:ogc:def:crs:EPSG::4979"
```

```

        xmlns:gs="http://www.opengis.net/pidflo/1.0"
        xmlns:gml="http://www.opengis.net/gml">
<gs:base>
  <gml:Polygon>
    <gml:exterior>
      <gml:LinearRing>
        <gml:posList>
          $lowLatitude$ $lowLongitude$ $lowAltitude$
          $lowLatitude$ $highLongitude$ $lowAltitude$
          $highLatitude$ $highLongitude$ $lowAltitude$
          $highLatitude$ $lowLongitude$ $lowAltitude$
          $lowLatitude$ $lowLongitude$ $lowAltitude$
        </gml:posList>
      </gml:LinearRing>
    </gml:exterior>
  </gml:Polygon>
</gs:base>
<gs:height uom="urn:ogc:def:uom:EPSG::9001">
  $highAltitude - lowAltitude$
</gs:height>
</gs:Prism>

```

The Polygon shape is used if Altitude is omitted or specified in floors, or if either NAD83 datum is used (value 2 or 3). The corresponding GML Polygon has the following form:

```

<gml:Polygon srsName="$2D-CRS-URN$"
  xmlns:gml="http://www.opengis.net/gml">>
  <gml:exterior>
    <gml:LinearRing>
      <gml:posList>
        $lowLatitude$ $lowLongitude$
        $lowLatitude$ $highLongitude$
        $highLatitude$ $highLongitude$
        $highLatitude$ $lowLongitude$
        $lowLatitude$ $lowLongitude$
      </gml:posList>
    </gml:LinearRing>
  </gml:exterior>
</gml:Polygon>

```

The value "2D-CRS-URN" is defined by the datum value: If the datum is WGS84 (value 1), then the 2D-CRS-URN is "urn:ogc:def:crs:EPSG::4326". If the datum is NAD83 (value 2 or 3), then the 2D-CRS-URN is "urn:ogc:def:crs:EPSG::4269".

A Polygon shape with the WGS84 three-dimensional CRS is used if the datum is WGS84 (value 1) and the Altitude is specified in meters

(Altitude type 1), but no Altitude uncertainty is specified (that is, AltUnc is 0). In this case, the value of the Altitude field is added after each of the points above, and the srsName attribute is set to the three-dimensional WGS84 CRS, namely "urn:ogc:def:crs:EPSG::4979".

A simple point shape is used if either Latitude uncertainty (LatUnc) or Longitude uncertainty (LongUnc) is not specified. With Altitude, this uses a three-dimensional CRS; otherwise, it uses a two-dimensional CRS.

```
<gml:Point srsName="$CRS-URN$"  
  xmlns:gml="http://www.opengis.net/gml">  
  <gml:pos>$Latitude$ $Longitude$ $[Altitude]$</gml:pos>  
</gml:Point>
```

A.1.1.1. Finding Low and High Values using Uncertainty Fields

For the DHCPv4 GeoConf Option 123, resolution fields are used (LaRes, LoRes, AltRes), indicating how many bits of a value contain information. Any bits beyond those indicated can be either zero or one.

For the DHCPv6 GeoLoc Option TBD2 and DHCPv4 GeoLoc Option TBD1, the LatUnc, LongUnc and AltUnc fields indicate uncertainty distances, denoting the bounds of the location region described by the DHCP location object.

The two sections below describe how to compute the Latitude, Longitude, and Altitude bounds (e.g., \$lowLatitude\$, \$highAltitude\$) in the templates above. The first section describes how these bounds are computed in the "resolution encoding" (DHCPv4 GeoConf Option 123), while the second section addresses the "uncertainty encoding" (DHCPv6 GeoLoc Option TBD2 and DHCPv4 GeoLoc Option TBD1).

A.1.1.1.1. Resolution Encoding

Given a number of resolution bits (i.e., the value of a resolution field), if all bits beyond those bits are set to zero, this gives the lowest possible value. The highest possible value can be found setting all bits to one.

If the encoded value of Latitude/Longitude and resolution (LaRes, LoRes) are treated as 34-bit unsigned integers, the following can be used (where ">>" is a bitwise right shift, "&" is a bitwise AND, "~" is a bitwise negation, and "|" is a bitwise OR).

```
mask = 0x3fffffff >> resolution  
lowvalue = value & ~mask
```

$$\text{highvalue} = \text{value} \mid \text{mask} + 1$$

Once these values are determined, the corresponding floating point numbers can be computed by dividing the values by 2^{25} (since there are 25 bits of fraction in the fixed-point representation).

Alternatively, the lowest possible value can be found by using resolution to determine the size of the range. This method has the advantage that it operates on the decoded floating point values. It is equivalent to the first mechanism, to a possible error of 2^{-25} (2^{-8} for altitude).

$$\begin{aligned} \text{scale} &= 2^{(9 - \text{resolution})} \\ \text{lowvalue} &= \text{floor}(\text{value} / \text{scale}) * \text{scale} \\ \text{highvalue} &= \text{lowvalue} + \text{scale} \end{aligned}$$

Altitude resolution (AltRes) uses the same process with different constants. There are 22 whole bits in the Altitude encoding (instead of 9) and 30 bits in total (instead of 34).

A.1.1.2. Uncertainty Encoding

In the uncertainty encoding, the uncertainty fields (LongUnc/LatUnc) directly represent the logarithms of uncertainty distances. So the low and high bounds are computed by first computing the uncertainty distances, then adding and subtracting these from the value provided. If "uncertainty" is the unsigned integer value of the uncertainty field and "value" is the value of the coordinate field:

$$\begin{aligned} \text{distance} &= 2^{(8 - \text{uncertainty})} \\ \text{lowvalue} &= \text{value} - \text{distance} \\ \text{highvalue} &= \text{value} + \text{distance} \end{aligned}$$

Altitude uncertainty (AltUnc in version 1) uses the same process with different constants:

$$\begin{aligned} \text{distance} &= 2^{(21 - \text{uncertainty})} \\ \text{lowvalue} &= \text{value} - \text{distance} \end{aligned}$$

Appendix B. Calculations of Resolution

The following examples for two different locations demonstrate how the Resolution values for Latitude, Longitude, and Altitude (used in DHCPv4 GeoConf Option 123) can be calculated. In both examples, the geo-location values were derived from maps using the WGS84 map datum, therefore in these examples, the Datum field would have a value = 1 (00000001, or 0x01).

B.1. Location Configuration Information of "White House" (Example 1)

The grounds of the White House in Washington D.C. (1600 Pennsylvania Ave. NW, Washington, DC 20006) can be found between 38.895375 and 38.898653 degrees North and 77.037911 and 77.035116 degrees West. In this example, we assume that we are standing on the sidewalk on the north side of the White House, between driveways. Since we are not inside a structure, we assume an Altitude value of 15 meters, interpolated from the US Geological survey map, Washington Washington West quadrangle.

The address was NOT picked for any political reason and can easily be found on the Internet or mapping software, but was picked as an easily identifiable location on our planet.

In this example, the requirement of emergency responders in North America via their NENA Model Legislation [NENA] could be met by a LaRes value of 21 and a LoRes value of 20. This would yield a geo-location that is Latitude 38.8984375 north to Latitude 38.8988616 north and Longitude -77.0371094 to Longitude -77.0375977. This is an area of approximately 89 feet by 75 feet or 6669 square feet, which is very close to the 7000 square feet requested by NENA. In this example, a service provider could enforce that a device send a Location Configuration Information with this minimum amount of resolution for this particular location when calling emergency services.

An approximate representation of this location might be provided using the DHCPv4 GeoConf Option 123 encoding as follows:

```

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
+-----+-----+-----+-----+-----+-----+-----+-----+
| Code (123) | OptLen (16) | LaRes | Latitude |
| 0 1 1 1 1 0 1 1 | 0 0 0 1 0 0 0 0 | 0 1 0 0 1 0 | 0 0 0 1 0 0 1 1 0 1 |
+-----+-----+-----+-----+-----+-----+-----+-----+
.
. Latitude (cont'd) | LoRes |
. 1 1 0 0 1 0 1 1 1 0 0 1 1 0 0 0 0 1 1 0 0 0 1 1 | 0 1 0 0 0 1 | 1 1 |
+-----+-----+-----+-----+-----+-----+-----+-----+
.
. Longitude (cont'd) |
. 0 1 1 0 0 1 0 1 1 1 1 0 1 1 0 1 0 1 0 0 0 0 1 0 1 1 0 0 0 1 0 0 |
+-----+-----+-----+-----+-----+-----+-----+-----+
| AType | AltRes | Altitude |
| 0 0 0 1 | 0 1 0 0 0 1 | 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 1 1 1 1 |
+-----+-----+-----+-----+-----+-----+-----+-----+
. Alt (cont'd) | Res | Datum |
. 0 0 0 0 0 0 0 0 0 0 | 0 0 0 0 0 0 | 0 0 1 |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

In hexadecimal, this is 7B10484D CB986347 65ED42C4 1440000F 0001.

Decoding Location Configuration Information with Resolution

Decoding this option gives a latitude of 38.897647 (to 7 decimal places) with 18 bits of resolution; a longitude of -77.0366000 with 17 bits of resolution; an altitude type of meters with a value of 15 and 17 bits of resolution; version 0 (resolution) and the WGS84 datum.

For the latitude value, 18 bits of resolution allow for values in the range from 38.8964844 to 38.8984375. For the longitude value, 17 bits of resolution allow for values in the range from -77.0390625 to -77.0351563. Having 17 bits of resolution in the altitude allows for values in the range from 0 to 32 meters.

GML Representation of Decoded Location Configuration Information

The following GML shows the value decoded in the previous example as a point in a three dimensional CRS:

```
<gml:Point srsName="urn:ogc:def:crs:EPSG::4979"
  xmlns:gml="http://www.opengis.net/gml">
  <gml:pos>38.897647 -77.0366 15</gml:pos>
</gml:Point>
```

This representation ignores the values included in the resolution parameters. If resolution values are provided, a rectangular prism can be used to represent the location.

The following example uses all of the decoded information from the previous example:

```
<gs:Prism srsName="urn:ogc:def:crs:EPSG::4979"
  xmlns:gs="http://www.opengis.net/pidflo/1.0"
  xmlns:gml="http://www.opengis.net/gml">
  <gs:base>
    <gml:Polygon>
      <gml:exterior>
        <gml:LinearRing>
          <gml:posList>
            38.8964844 -77.0390625 0
            38.8964844 -77.0351563 0
            38.8984375 -77.0351563 0
            38.8984375 -77.0390625 0
            38.8964844 -77.0390625 0
          </gml:posList>
        </gml:LinearRing>
      </gml:exterior>
    </gml:Polygon>
  </gs:base>
</gs:Prism>
```

```

        </gml:exterior>
      </gml:Polygon>
    </gs:base>
    <gs:height uom="urn:ogc:def:uom:EPSG::9001">
      32
    </gs:height>
  </gs:Prism>

```

B.2. Location Configuration Information of "Sears Tower" (Example 2)

Postal Address:

```

Sears Tower
103rd Floor
233 S. Wacker Dr.
Chicago, IL 60606

```

Viewing the Chicago area from the Observation Deck of the Sears Tower.

```

Latitude 41.87884 degrees North (or +41.87884 degrees)
Using 2s complement, 34 bit fixed point, 25 bit fraction
Latitude = 0x053clf751,
Latitude = 0001010011110000011111011101010001
Longitude 87.63602 degrees West (or -87.63602 degrees)
Using 2s complement, 34 bit fixed point, 25 bit fraction
Longitude = 0xf50ba5b97,
Longitude = 1101010000101110100101101110010111

```

Altitude 103

In this example, we are inside a structure, therefore we will assume an Altitude value of 103 to indicate the floor we are on. The Altitude Type value is 2, indicating floors. The AltRes field would indicate that all bits in the Altitude field are true, as we want to accurately represent the floor of the structure where we are located.

```

AltRes = 30, 0x1e, 011110
AType = 2, 0x02, 000010
Altitude = 103, 0x00006700, 00000000000000001100111000000000

```

For the accuracy of the Latitude and Longitude, the best information available to us was supplied by a generic mapping service that shows a single geo-loc for all of the Sears Tower. Therefore we are going to show LaRes as value 18 (0x12 or 010010) and LoRes as value 18 (0x12 or 010010). This would be describing a geo-location area that is Latitude 41.8769531 to Latitude 41.8789062 and extends from -87.6367188 degrees to -87.6347657 degrees Longitude. This is an area of approximately 373412 square feet (713.3 ft. x 523.5 ft.).

Appendix C. Calculations of Uncertainty

The following example demonstrates how uncertainty values for Latitude, Longitude, and Altitude (LatUnc, LongUnc and AltUnc used in the DHCPv6 GeoLoc Option TBD2 as well as DHCPv4 GeoLoc Option TBD1) can be calculated.

C.1. Location Configuration Information of "Sydney Opera House" (Example 3)

This section describes an example of encoding and decoding the geodetic DHCP Option. The textual results are expressed in GML [OGC.GML-3.1.1] form, suitable for inclusion in PIDF-LO [RFC4119].

These examples all assume a datum of WGS84 (datum = 1) and an Altitude type of meters (AType = 1).

C.1.1. Encoding a Location into DHCP Geodetic Form

This example draws a rough polygon around the Sydney Opera House. This polygon consists of the following six points:

```
33.856625 S, 151.215906 E
33.856299 S, 151.215343 E
33.856326 S, 151.214731 E
33.857533 S, 151.214495 E
33.857720 S, 151.214613 E
33.857369 S, 151.215375 E
```

The top of the building 67.4 meters above sea level, and a starting Altitude of 0 meters above the WGS84 geoid is assumed.

The first step is to determine the range of Latitude and Longitude values. Latitude ranges from -33.857720 to -33.856299; Longitude ranges from 151.214495 to 151.215906.

For this example, the point that is encoded is chosen by finding the middle of each range, that is (-33.8570095, 151.2152005). This is encoded as (1110111100010010010011011000001101, 010010111001101110001011011000011) in binary, or (3BC49360D, 12E6E2EC3) in hexadecimal notation (with an extra 2 bits of leading padding on each). Altitude is set at 33.7 meters, which is 0000000000000000010000110110011 (binary) or 000021B3 (hexadecimal).

The Latitude Uncertainty (LatUnc) is given by inserting the difference between the center value and the outer value into the formula from Section 2.3.1. This gives:

$$x = 8 - \text{ceil}(\log_2(-33.8570095 - -33.857720))$$

The result of this equation is 18, therefore the uncertainty is encoded as 010010 in binary.

Similarly, Longitude Uncertainty (LongUnc) is given by the formula:

$$x = 8 - \text{ceil}(\log_2(151.2152005 - 151.214495))$$

The result of this equation is also 18, or 010010 in binary.

Altitude Uncertainty (AltUnc) uses the formula from Section 2.4.4:

$$x = 21 - \text{ceil}(\log_2(33.7 - 0))$$

The result of this equation is 15, which is encoded as 001111 in binary.

Adding an Altitude Type of 1 (meters) and a Datum of 1 (WGS84), this gives the following DHCPv4 GeoLoc Option TBD1 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																																
Code (TBD1)																OptLen (16)																LatUnc																Latitude															
0	1	1	1	1	0	1	1	0	0	0	1	0	0	0	0	0	1	0	0	1	0	1	1	1	0	1	1	1	1	0	0																																
Latitude (cont'd)																LongUnc																																															
0	1	0	0	1	0	0	1	0	0	1	1	0	1	1	0	0	0	0	0	0	1	1	0	1	0	1	0	1	0	1																																	
Longitude (cont'd)																																																															
0	0	1	0	1	1	1	0	0	1	1	0	1	1	1	0	0	0	1	0	1	1	1	0	1	1	0	0	0	0	1	1																																
AType																AltUnc																Altitude																															
0	0	0	1	0	0	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	0	0	0	1																																	
Alt (cont'd)																Ver																Res																Datum															
1	0	1	1	0	0	1	1	0	1	0	1	0	0	0	0	0	0	0	0	1	0	0	0	0	0	0	0	0	0	1																																	

In hexadecimal, this is 7B104BBC 49360D49 2E6E2EC3 13C00021 B341.
The DHCPv6 form only differs in the code and option length portion.

C.1.2. Decoding a Location from DHCP Geodetic Form

If receiving the binary form created in the previous section, this section describes how that would be interpreted. The result is then represented as a GML object, as defined in [GeoShape].

A Latitude value of 1110111100010010010011011000001101 decodes to a value of -33.8570095003 (to 10 decimal places). The Longitude value of 0100101110011011100010111011000011 decodes to 151.2152005136.

Decoding Tip: If the raw values of Latitude and Longitude are placed in integer variables, the actual value can be derived by the following process:

1. If the highest order bit is set (i.e. the number is a twos complement negative), then subtract 2 to the power of 34 (the total number of bits).
2. Divide the result by 2 to the power of 25 (the number of fractional bits) to determine the final value.

The same principle can be applied when decoding Altitude values, except with different powers of 2 (30 and 8 respectively).

The Latitude and Longitude Uncertainty are both 18, which gives an uncertainty value using the formula from Section 2.3.1 of 0.0009765625. Therefore, the decoded Latitudes is -33.8570095003 +/- 0.0009765625 (or the range from -33.8579860628 to -33.8560329378) and the decoded Longitude is 151.2152005136 +/- 0.0009765625 (or the range from 151.2142239511 to 151.2161770761).

The encoded Altitude of 000000000000000010000110110011 decodes to 33.69921875. The encoded uncertainty of 15 gives a value of 64, therefore the final uncertainty is 33.69921875 +/- 64 (or the range from -30.30078125 to 97.69921875).

C.1.2.1. GML Representation of Decoded Locations

The following GML shows the value decoded in the previous example as a point in a three dimensional CRS:

```
<gml:Point srsName="urn:ogc:def:crs:EPSG::4979"
  xmlns:gml="http://www.opengis.net/gml">
  <gml:pos>-33.8570095003 151.2152005136 33.69921875</gml:pos>
</gml:Point>
```

The following example uses all of the decoded information from the previous example:

```
<gs:Prism srsName="urn:ogc:def:crs:EPSG::4979"
  xmlns:gs="http://www.opengis.net/pidflo/1.0"
  xmlns:gml="http://www.opengis.net/gml">
  <gs:base>
    <gml:Polygon>
```

```
<gml:exterior>
  <gml:LinearRing>
    <gml:posList>
      -33.8579860628 151.2142239511 -30.30078125
      -33.8579860628 151.2161770761 -30.30078125
      -33.8560329378 151.2161770761 -30.30078125
      -33.8560329378 151.2142239511 -30.30078125
      -33.8579860628 151.2142239511 -30.30078125
    </gml:posList>
  </gml:LinearRing>
</gml:exterior>
</gml:Polygon>
</gs:base>
<gs:height uom="urn:ogc:def:uom:EPSG::9001">
  128
</gs:height>
</gs:Prism>
```

Note that this representation is only appropriate if the uncertainty is sufficiently small. [GeoShape] recommends that distances between polygon vertices be kept short. A GML representation like this one is only appropriate where uncertainty is less than 1 degree (an encoded value of 9 or greater).

Appendix D. Changes from RFC 3825

This section lists the major changes between RFC 3825 and this document. Minor changes, including style, grammar, spelling and editorial changes are not mentioned here.

- o Section 1 now includes clarifications on wired and wireless uses.
- o The former Sections 1.2 and 1.3 have been removed. Section 1.2 now defines the concepts of uncertainty and resolution, as well as conversion between the DHCP option formats and PIDF-LO.
- o A DHCPv6 GeoLoc Option is now defined (Section 2.1) as well as a new DHCPv4 GeoLoc Option (Section 2.2.2).
- o The former Datum field has been split into three fields: Ver, Res and Datum. These fields are used in both the DHCPv4 GeoLoc Option and the DHCPv6 GeoLoc Option.
- o Section 2.2.3 has been added, describing option support requirements on DHCP clients and servers.
- o Section 2.3 has been added, describing the Latitude and Longitude fields.
- o Section 2.3.1 has been added, covering Latitude and Longitude resolution.
- o Section 2.3.2 has been added, covering Latitude and Longitude uncertainty.
- o Section 2.4 has been added, covering values of the Altitude field (Sections 2.4.1, 2.4.2 and 2.4.3), Altitude resolution (Section 2.4.4), and Altitude uncertainty (Section 2.4.5).
- o Section 2.5 has been added, covering the Datum field.
- o Section 3 (Security Considerations) has added a recommendation on link layer confidentiality.
- o Section 4 (IANA Considerations) has consolidated material relating to parameter allocation for both the DHCPv4 and DHCPv6 option parameters, and has been rewritten to conform to the practices recommended in RFC 5226.
- o The material formerly in Appendix A has been updated and shortened and has been moved to Appendix B.
- o An Appendix A on GML mapping has been added.
- o Appendix C has been added, providing an example of uncertainty encoding.
- o Appendix D has been added, detailing the changes from RFC 3825.

Authors' Addresses

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

EMail: jmpolk@cisco.com

Marc Linsner
Cisco Systems
Marco Island, FL 34145 USA
USA

EMail: marc.linsner@cisco.com

Martin Thomson
Andrew Corporation
PO Box U40
Wollongong University Campus, NSW 2500
AU

EMail: martin.thomson@andrew.com

Bernard Aboba
Microsoft Corporation
One Microsoft Way
Redmond, WA 98052 USA
USA

EMail: bernard_aboba@hotmail.com