

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

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

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

Abstract

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

Status of This Memo

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

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

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

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

Copyright Notice

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

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

Table of Contents

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

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

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

1. Introduction

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

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

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

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

2. Conventions used in this document

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

This document also uses the following definitions:

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

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

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

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

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

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

3. Location-Related Measurements in LCPs

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

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

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

Figure 1: HELD Location Request with Measurement Data

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

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

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

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

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

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

4. Location-Related Measurement Data Types

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

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

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

Figure 2: Measurement Example

4.1. Measurement Container

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

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

4.1.1. Time of Measurement

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

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

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

4.1.2. Expiry Time on Location-Related Measurement Data

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

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

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

4.2. RMS Error and Number of Samples

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

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

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

Two attributes are provided for use on some measurement values:

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

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

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

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

4.2.1. Time RMS Error

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

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

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

4.3. Measurement Request

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

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

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

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

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

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

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

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

Figure 3: HELD Error Requesting Measurement Data

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

4.4. Identifying Location Provenance

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

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

The following types of measurement source are identified:

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

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

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

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

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

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

PIDF-LO document with source labels

5. Location-Related Measurement Data Types

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

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

5.1. LLDP Measurements

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

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

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

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

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

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

Figure 4: LLDP Measurement Example

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

5.2. DHCP Relay Agent Information Measurements

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

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

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

Figure 5: DHCP Relay Agent Information Measurement Example

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

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

5.3. 802.11 WLAN Measurements

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

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

Figure 6: 802.11 WLAN Measurement Example

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

5.3.1. Wifi Measurement Requests

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

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

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

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

5.4. Cellular Measurements

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

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

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

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

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

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

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

Figure 7: Example LTE Cellular Measurement

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

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

Figure 8: Example UMTS Cellular Measurement

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

```

    </cellular>
</measurements>

```

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

Figure 9: Example GSM Cellular Measurement

```

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

```

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

Figure 10: Example CDMA Cellular Measurement

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

```

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

```

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

Figure 11: Example Observed Cellular Measurement

5.4.1. Cellular Measurement Requests

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

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

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

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

5.5. GNSS Measurements

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

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

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

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

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

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

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

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

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

Figure 12: Example GNSS Measurement

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

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

5.5.1. GNSS System and Signal

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

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

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

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

5.5.2. Time

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

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

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

5.5.3. Per-Satellite Measurement Data

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

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

The GNSS receiver measures the following parameters for each satellite:

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

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

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

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

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

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

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

5.5.4. GNSS Measurement Requests

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

5.6. DSL Measurements

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

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

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

5.6.1. L2TP Measurements

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

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

Figure 13: Example DSL L2TP Measurement

5.6.2. RADIUS Measurements

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

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

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

Figure 14: Example DSL RADIUS Measurement

5.6.3. Ethernet VLAN Tag Measurements

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

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

Figure 15: Example DSL VLAN Tag Measurement

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

5.6.4. ATM Virtual Circuit Measurements

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

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

Figure 16: Example DSL ATM Measurement

6. Privacy Considerations

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

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

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

6.1. Measurement Data Privacy Model

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

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

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

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

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

6.2. LIS Privacy Requirements

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

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

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

6.3. Measurement Data and Location URIs

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

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

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

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

6.4. Third-Party-Provided Measurement Data

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

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

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

7. Security Considerations

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

7.1. Threat Model

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

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

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

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

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

7.1.1. Acquiring Location Information Without Authorization

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

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

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

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

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

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

7.1.2. Extracting Network Topology Data

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

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

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

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

7.1.3. Exposing Network Topology Data

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

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

7.1.4. Lying By Proxy

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

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

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

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

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

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

7.1.5. Measurement Replay

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

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

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

7.1.6. Environment Spoofing

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

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

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

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

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

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

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

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

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

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

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

7.2. Mitigation

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

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

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

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

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

7.2.1. Measurement Validation

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

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

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

7.2.1.1. Effectiveness

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

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

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

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

7.2.1.2. Limitations (Unique Observer)

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

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

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

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

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

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

7.2.2. Location Validation

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

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

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

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

7.2.2.1. Effectiveness

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

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

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

7.2.2.2. Limitations

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

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

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

7.2.3. Supporting Observations

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

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

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

7.2.3.1. Effectiveness

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

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

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

7.2.3.2. Limitations

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

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

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

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

7.2.4. Attribution

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

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

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

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

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

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

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

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

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

7.2.5. Stateful Correlation of Location Requests

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

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

7.3. An Unauthorized or Compromised LIS

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

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

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

8. Measurement Schemas

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

8.1. Measurement Container Schema

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

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

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

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

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

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

```
</xs:complexType>

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

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

Measurement Container Schema

8.2. Measurement Source Schema

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

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

```

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

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

```

Measurement Source PIDF-LO Extension Schema

8.3. Base Type Schema

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

```

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

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

```

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

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

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

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

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

```

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

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

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

```

```

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

</xs:schema>

```

Base Type Schema

8.4. LLDP Measurement Schema

```

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

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

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

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

```



```

</xs:complexType>

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

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

</xs:schema>

```

LLDP measurement schema

8.5. DHCP Measurement Schema

```

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

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

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

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

```

```

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

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

</xs:schema>

```

DHCP measurement schema

8.6. WiFi Measurement Schema

```

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

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

```

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

</xs:schema>
```

WiFi measurement schema

8.7. Cellular Measurement Schema

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

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

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

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

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

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

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

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

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

```

```
</xs:simpleType>

<!-- Measurement Request elements -->

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

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

</xs:schema>
```

Cellular measurement schema

8.8. GNSS Measurement Schema

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

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



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

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

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

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

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

```

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

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

```

GNSS measurement Schema

8.9. DSL Measurement Schema

```

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

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

```

```
</xs:annotation>

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

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

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

</xs:schema>
```

DSL measurement schema

9. IANA Considerations

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

9.1. IANA Registry for GNSS Types

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

Each entry in the registry requires the following information:

GNSS name: the name of the GNSS

Brief description: a brief description of the GNSS

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

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

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

The registry initially includes two registrations:

GNSS name: Global Positioning System (GPS)

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

GNSS token: gps

Signals: L1, L2, L1C, L2C, L5

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

GNSS name: Galileo

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

GNSS Token: galileo

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

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

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

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

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

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

XML:

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

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

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

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

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

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

XML:

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

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

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

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

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

XML:

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

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

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

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

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

XML:

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

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

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

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

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

XML:

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



```
</html>
END
```

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

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

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

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

XML:

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

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

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

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

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

XML:

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

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

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

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

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

XML:

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

```

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

```

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

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

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

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

XML:

```

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

```

9.11. XML Schema Registration for Measurement Source Schema

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

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

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

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

9.12. XML Schema Registration for Measurement Container Schema

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

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

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

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

9.13. XML Schema Registration for Base Types Schema

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

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

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

Schema: The XML for this schema can be found in Section 8.3 of this document.

9.14. XML Schema Registration for LLDP Schema

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

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

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

Schema: The XML for this schema can be found in Section 8.4 of this document.

9.15. XML Schema Registration for DHCP Schema

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

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

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

Schema: The XML for this schema can be found in Section 8.5 of this document.

9.16. XML Schema Registration for WiFi Schema

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

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

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

Schema: The XML for this schema can be found in Section 8.6 of this document.

9.17. XML Schema Registration for Cellular Schema

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

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

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

Schema: The XML for this schema can be found in Section 8.7 of this document.

9.18. XML Schema Registration for GNSS Schema

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

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

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

Schema: The XML for this schema can be found in Section 8.8 of this document.

9.19. XML Schema Registration for DSL Schema

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

URI: urn:iETF:params:xml:schema:lm:dsl

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

Schema: The XML for this schema can be found in Section 8.9 of this document.

10. Acknowledgements

Thanks go to Simon Cox for his comments relating to terminology that have helped ensure that this document is aligned with ongoing work in the Open Geospatial Consortium (OGC). Thanks to Neil Harper for his review and comments on the GNSS sections of this document. Thanks to Noor-E-Gagan Singh, Gabor Bajko, Russell Priebe, and Khalid Al-Mufti for their significant input to and suggestions for improving the 802.11 measurements. Thanks to Cullen Jennings for feedback and suggestions. Bernard Aboba provided review and feedback on a range of measurement data definitions. Mary Barnes and Geoff Thompson provided a review and corrections. David Waitzman and John Bressler both noted shortcomings with 802.11 measurements. Keith Drage, Darren Pawson provided expert LTE knowledge.

11. References

11.1. Normative References

[ASCII] , "US-ASCII. Coded Character Set - 7-Bit American Standard Code for Information Interchange. Standard ANSI X3.4-1986, ANSI, 1986.", .

[GPS.ICD] , "Navstar GPS Space Segment/Navigation User Interface", ICD GPS-200, Apr 2000.

[Galileo.ICD]
GJU, "Galileo Open Service Signal In Space Interface Control Document (SIS ICD)", May 2006.

[IANA.enterprise]
IANA, "Private Enterprise Numbers", 2011,
<<http://www.iana.org/assignments/enterprise-numbers>>.

[IEEE.80211]

IEEE, "Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications", IEEE Std 802.11-2012, March 2012.

[IEEE.8021AB]

IEEE, "IEEE Standard for Local and Metropolitan area networks, Station and Media Access Control Connectivity Discovery", IEEE Std 802.1AB-2009, September 2009.

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3046] Patrick, M., "DHCP Relay Agent Information Option", RFC 3046, January 2001.
- [RFC3315] Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, July 2003.
- [RFC3629] Yergeau, F., "UTF-8, a transformation format of ISO 10646", STD 63, RFC 3629, November 2003.
- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, January 2005.
- [RFC3993] Johnson, R., Palaniappan, T., and M. Stapp, "Subscriber-ID Suboption for the Dynamic Host Configuration Protocol (DHCP) Relay Agent Option", RFC 3993, March 2005.
- [RFC4119] Peterson, J., "A Presence-based GEOPRIV Location Object Format", RFC 4119, December 2005.
- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, February 2006.
- [RFC4580] Volz, B., "Dynamic Host Configuration Protocol for IPv6 (DHCPv6) Relay Agent Subscriber-ID Option", RFC 4580, June 2006.
- [RFC4649] Volz, B., "Dynamic Host Configuration Protocol for IPv6 (DHCPv6) Relay Agent Remote-ID Option", RFC 4649, August 2006.
- [RFC5491] Winterbottom, J., Thomson, M., and H. Tschofenig, "GEOPRIV Presence Information Data Format Location Object (PIDF-LO) Usage Clarification, Considerations, and Recommendations", RFC 5491, March 2009.

- [RFC5952] Kawamura, S. and M. Kawashima, "A Recommendation for IPv6 Address Text Representation", RFC 5952, August 2010.
- [RFC5985] Barnes, M., "HTTP-Enabled Location Delivery (HELD)", RFC 5985, September 2010.
- [RFC5986] Thomson, M. and J. Winterbottom, "Discovering the Local Location Information Server (LIS)", RFC 5986, September 2010.
- [TIA-2000.5] TIA/EIA, "Upper Layer (Layer 3) Signaling Standard for cdma2000(R) Spread Spectrum Systems", TIA-2000.5-D, March 2004.
- [TS.3GPP.23.003] 3GPP, "Numbering, addressing and identification", 3GPP TS 23.003 9.4.0, September 2010.

11.2. Informative References

- [ANSI-TIA-1057] ANSI/TIA, "Link Layer Discovery Protocol for Media Endpoint Devices", TIA 1057, April 2006.
- [DSL.TR025] Wang, R., "Core Network Architecture Recommendations for Access to Legacy Data Networks over ADSL", September 1999.
- [DSL.TR101] Cohen, A. and E. Shrum, "Migration to Ethernet-Based DSL Aggregation", April 2006.
- [GPS.SPOOF] Scott, L., "Anti-Spoofing and Authenticated Signal Architectures for Civil Navigation Signals", ION-GNSS Portland, Oregon, 2003.
- [HARPER] Harper, N., Dawson, M., and D. Evans, "Server-side spoofing and detection for Assisted-GPS", Proceedings of International Global Navigation Satellite Systems Society (IGNSS) Symposium 2009 16, December 2009, <<http://ignss.org/files/Paper16.pdf>>.
- [RFC2661] Townsley, W., Valencia, A., Rubens, A., Pall, G., Zorn, G., and B. Palter, "Layer Two Tunneling Protocol "L2TP"", RFC 2661, August 1999.

- [RFC2865] Rigney, C., Willens, S., Rubens, A., and W. Simpson, "Remote Authentication Dial In User Service (RADIUS)", RFC 2865, June 2000.
- [RFC3688] Mealling, M., "The IETF XML Registry", BCP 81, RFC 3688, January 2004.
- [RFC3693] Cuellar, J., Morris, J., Mulligan, D., Peterson, J., and J. Polk, "Geopriv Requirements", RFC 3693, February 2004.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.
- [RFC6155] Winterbottom, J., Thomson, M., Tschofenig, H., and R. Barnes, "Use of Device Identity in HTTP-Enabled Location Delivery (HELD)", RFC 6155, March 2011.
- [RFC6280] Barnes, R., Lepinski, M., Cooper, A., Morris, J., Tschofenig, H., and H. Schulzrinne, "An Architecture for Location and Location Privacy in Internet Applications", BCP 160, RFC 6280, July 2011.

Authors' Addresses

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

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

James Winterbottom
Unaffiliated
AU

Email: a.james.winterbottom@gmail.com

GEOPRIV
Internet-Draft
Intended status: Standards Track
Expires: February 22, 2013

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

M. Thomson
Microsoft
August 21, 2012

Geolocation Policy: A Document Format for Expressing Privacy Preferences
for Location Information
draft-ietf-geopriv-policy-27

Abstract

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

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

Status of this Memo

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

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

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any

time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on February 22, 2013.

Copyright Notice

Copyright (c) 2012 IETF Trust and the persons identified as the document authors. All rights reserved.

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

Table of Contents

1.	Introduction	5
2.	Terminology	7
3.	Generic Processing	8
3.1.	Structure of Geolocation Authorization Documents	8
3.2.	Rule Transport	8
4.	Location-specific Conditions	9
4.1.	Geodetic Location Condition Profile	9
4.2.	Civic Location Condition Profile	10
5.	Actions	11
6.	Transformations	12
6.1.	Set Retransmission-Allowed	12
6.2.	Set Retention-Expiry	12
6.3.	Set Note-Well	12
6.4.	Keep Ruleset Reference	13
6.5.	Provide Location	13
6.5.1.	Civic Location Profile	14
6.5.2.	Geodetic Location Profile	15
7.	Examples	18
7.1.	Rule Example with Civic Location Condition	18
7.2.	Rule Example with Geodetic Location Condition	19
7.3.	Rule Example with Civic and Geodetic Location Condition	19
7.4.	Rule Example with Location-based Transformations	20
7.5.	Location Obfuscation Example	22
8.	XML Schema for Basic Location Profiles	26
9.	XML Schema for Geolocation Policy	27
10.	XCAP Usage	29
10.1.	Application Unique ID	29
10.2.	XML Schema	29
10.3.	Default Namespace	29
10.4.	MIME Media Type	29
10.5.	Validation Constraints	29
10.6.	Data Semantics	29
10.7.	Naming Conventions	29
10.8.	Resource Interdependencies	30
10.9.	Authorization Policies	30
11.	IANA Considerations	31
11.1.	Geolocation Policy XML Schema Registration	31
11.2.	Geolocation Policy Namespace Registration	31
11.3.	Geolocation Policy Location Profile Registry	32
11.4.	Basic Location Profile XML Schema Registration	32
11.5.	Basic Location Profile Namespace Registration	33
11.6.	XCAP Application Usage ID	34
12.	Internationalization Considerations	35
13.	Security Considerations	36
13.1.	Introduction	36
13.2.	Obfuscation	36

13.3. Algorithm Limitations 38
13.4. Usability 38
13.5. Location Obscuring Limitations 39
14. References 41
14.1. Normative References 41
14.2. Informative References 41
Appendix A. Acknowledgments 44
Appendix B. Pseudo-Code 45
Authors' Addresses 49

1. Introduction

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

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

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

In the typical sequence of operations, a Location Server receives a query for location information for a particular Target. The requestor's identity will likely be revealed as part of this request for location information. The authenticated identity of the Location Recipient, together with other information provided with the request or generally available to the server, is then used for searching through the rule set. If more than one rule matches the condition element, then the combined permission is evaluated according to the description in Section 10 of [RFC4745]. The result of the rule evaluation is applied to the location information, yielding a possibly modified Location Object that is delivered to the Location Recipient.

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

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

This document defines two algorithms for reducing the granularity of returned location information. The first algorithm is defined for usage with civic location information (see Section 6.5.1) while the other one applies to geodetic location information (see Section 6.5.2). Both algorithms come with limitations, i.e. they provide location obfuscation under certain conditions and may therefore not be appropriate for all application domains. These limitations are documented within the security consideration section (see Section 13). It is worth pointing out that the geodetic transformation algorithm Section 6.5.2 deals with privacy risks related to targets that are stationary, as well as to moving targets. However, with respect to movement there are restriction as to what information can be hidden from an adversary. To cover applications that have more sophisticated privacy requirements additional algorithms may need to be defined. This document foresees extensions in the form of new algorithms and therefore defines a registry (see Section 11.3).

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

2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

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

Presentity or Target:

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

Watcher or Location Recipient:

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

Authorization policy:

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

Permission:

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

In this document we use the term Location Servers as the entities that evaluate the geolocation authorization policies. The geolocation privacy architecture is, as described in RFC 4079 [RFC4079], aligned with the presence architecture and a Presence Server is therefore an entity that distributes location information along with other presence-specific XML data elements.

3. Generic Processing

3.1. Structure of Geolocation Authorization Documents

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

3.2. Rule Transport

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

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

4. Location-specific Conditions

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

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

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

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

The `<location-condition>` and the `<location>` elements provide extension points. An extension that is not understood by the entity evaluating the rules then this rule evaluates to FALSE. This causes a `>conditions>` element to evaluate to FALSE if a `>location-condition>` element is unsupported, but allows a `>location-condition>` to be TRUE if an child `>location>` is not understood as long as an understood `>location>` is TRUE.

4.1. Geodetic Location Condition Profile

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

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

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

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

A CRS MUST be specified using the above URN notation only, implementations do not need to support user-defined CRSs.

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

4.2. Civic Location Condition Profile

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

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

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

5. Actions

This document does not define location-specific actions.

6. Transformations

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

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

6.1. Set Retransmission-Allowed

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

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

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

6.2. Set Retention-Expiry

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

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

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

6.3. Set Note-Well

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

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

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

6.4. Keep Ruleset Reference

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

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

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

6.5. Provide Location

The <provide-location> element contains child elements of a specific location profile that controls the granularity of returned location information. This form of location granularity reduction is also called 'obfuscation' and is defined in [duckham05] as

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

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

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

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

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

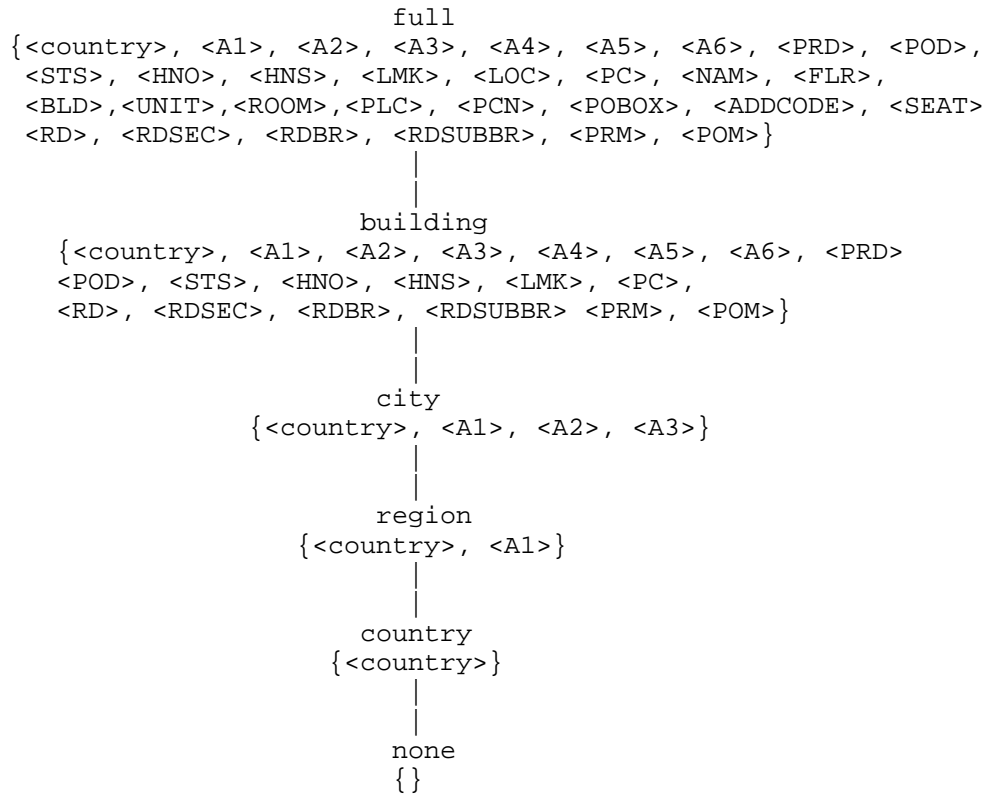
The <provide-location> element MUST contain the 'profile' attribute if it contains child elements and the 'profile' attribute MUST match with the contained child elements.

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

6.5.1. Civic Location Profile

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

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



The default value is "none".

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

6.5.2. Geodetic Location Profile

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

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

The algorithm proceeds in 6 steps. The first two steps are independent of the measured position to be obscured. Those two steps should be run only once or rather seldom (for every region and desired uncertainty). The steps are:

1. Choose a geodesic projection with Cartesian coordinates and a surface you want to cover. The maximal distortion of the map may not be too much (see notes below).
2. Given uncertainty "d", choose a grid of so called "landmarks" at a distance (maximal) d of each other.
3. Given a measured location $M=(m,n)$ in the surface, calculate its 4 closest landmarks on the grid, with coordinates: SW = (l,b), SE=(r,b), NW=(l,t), NE=(r,t). Thus $l \leq m < r$ and $b \leq n < t$. See notes below.
4. Let $x=(m-l)/(r-l)$ and $y=(n-b)/(t-b)$

x and y are thus the local coordinates of the point M in the small grid square that contains it. $0 \leq x, y < 1$.

5. Let $p = 0.2887$ ($=\sqrt{3}/6$) and $q = 0.7113$ ($=1-p$), determine which of the following 8 cases holds:

- C1. $x < p$ and $y < p$
- C2. $p \leq x < q$ and $y < x$ and $y < 1-x$
- C3. $q \leq x$ and $y < p$
- C4. $p \leq y < q$ and $x \leq y$ and $y < 1-x$
- C5. $p \leq y < q$ and $y < x$ and $1-x \leq y$
- C6. $x < p$ and $q \leq y$
- C7. $p \leq x < q$ and $x \leq y$ and $1-x \leq y$
- C8. $q \leq x$ and $q \leq y$

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

C1: SW
 C2: SW or SE
 C3: SE

C4: SW or NW
 C5: SE or NE

C6: NW
 C7: NW or NE
 C8: NE

Return the circle with center C and radius d.

Notes:

Regarding Step 1:

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

Regarding Step3:

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

7. Examples

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

7.1. Rule Example with Civic Location Condition

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

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

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

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

7.2. Rule Example with Geodetic Location Condition

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

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

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

7.3. Rule Example with Civic and Geodetic Location Condition

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

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

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

7.4. Rule Example with Location-based Transformations

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

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

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

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

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

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

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

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

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

7.5. Location Obfuscation Example

Suppose you want to obscure positions in the continental USA.

Step 1:

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

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

As origin of our grid, we choose the point at latitude 25 degrees and longitude 0 (Greenwich). The latitude 25 degrees is chosen to be just south of Florida and thus south of the continental US. (On the south hemisphere the origin should be north of the region to be covered; if the region crosses the Equator, the origin should be on the Equator. In this way it is guaranteed that the latitudinal degree has largest distance at the latitude of the origin).

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

The same procedure, basically, produces grids for

- * 45 degrees south to 45 degrees north Tropics and subtropics
- * 25 to 50 degrees (both north or south) Continental US
- * 35 to 55 degrees (both north or south) South and Central Europe
- * 45 to 60 degrees (both north or south) Central and North Europe
- * 55 to 65 degrees (both north or south) Scandinavia

* 60 to 70 degrees (both north or south)

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

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

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

Step 2:

To construct the grid points, we start with our chosen origin and place the along the main axes (NS and EW) grid points at a distance d of each other.

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

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

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

Step 3:

Now suppose you measure a position at M , with longitude -105 (the minus sign is used to denote 105 degrees *west*; without minus, the point is in China, 105 degrees east) and latitude 40 degrees

(just north of Denver, CO). The point M is 105 degrees west and 15 degrees north of our origin (which has longitude 0 and latitude 25).

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

$$i = \text{floor}(m/d1) = \text{floor}(-105/0.993) = -106$$

$$j = \text{floor}(n-o/d2) = \text{floor}(15/0.904) = 16$$

Those are the indexes of SW on the grid. The coordinates of SW are then: $(d1*i, 25+d2*j) = (-105.242, 39.467)$.

Thus:

$$l = d1 * \text{floor}(m/d1) = -105.243$$

$$r = l + d1 = -105.243 + 0.993 = -104.250$$

$$b = o + d2 * \text{floor}(n-o/d2) = 39.467$$

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

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

Step 4:

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

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

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

Step 5:

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

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

We are in case C4 ($p \leq y < q$ and $x \leq y$ and $y < 1-x$).

Step 6:

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

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

8. XML Schema for Basic Location Profiles

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

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

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

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

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

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

</xs:schema>
```

9. XML Schema for Geolocation Policy

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

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

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

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

  <!-- Geopriv Conditions -->

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

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

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

```
        <xs:attribute name="label" type="xs:string" />
        <xs:attribute ref="xml:lang" />
    </xs:restriction>
</xs:complexContent>
</xs:complexType>

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

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

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

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

</xs:schema>
```

10. XCAP Usage

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

10.1. Application Unique ID

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

10.2. XML Schema

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

10.3. Default Namespace

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

10.4. MIME Media Type

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

10.5. Validation Constraints

This specification does not define additional constraints.

10.6. Data Semantics

This document discusses the semantics of a geolocation privacy authorization.

10.7. Naming Conventions

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

found beneath that point to guide authorization policy.

10.8. Resource Interdependencies

This application usage does not define additional resource interdependencies.

10.9. Authorization Policies

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

11. IANA Considerations

There are several IANA considerations associated with this specification.

11.1. Geolocation Policy XML Schema Registration

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

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

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

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

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

and its last line is

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

11.2. Geolocation Policy Namespace Registration

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

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

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

XML:


```
BEGIN
<?xml version="1.0"?>
<!DOCTYPE html PUBLIC "-//W3C//DTD XHTML Basic 1.0//EN"
  "http://www.w3.org/TR/xhtml-basic/xhtml-basic10.dtd">
<html xmlns="http://www.w3.org/1999/xhtml">
<head>
  <meta http-equiv="content-type"
    content="text/html; charset=iso-8859-1"/>
  <title>Geolocation Policy Namespace</title>
</head>
<body>
  <h1>Namespace for Geolocation Authorization Policies</h1>
  <h2>urn:ietf:params:xml:schema:geolocation-policy</h2>
<p>See <a href="[URL of published RFC]">RFCXXXX
  [NOTE TO IANA/RFC-EDITOR:
    Please replace XXXX with the RFC number of this
    specification.]</a>.</p>
</body>
</html>
END
```

11.3. Geolocation Policy Location Profile Registry

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

This document defines the following profile names:

geodetic-condition: Defined in Section 4.1.

civic-condition: Defined in Section 4.2.

geodetic-transformation: Defined in Section 6.5.2.

civic-transformation: Defined in Section 6.5.1.

11.4. Basic Location Profile XML Schema Registration

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

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

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

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

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

and its last line is

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

11.5. Basic Location Profile Namespace Registration

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

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

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

XML:

BEGIN

```
<?xml version="1.0"?>
<!DOCTYPE html PUBLIC "-//W3C//DTD XHTML Basic 1.0//EN"
  "http://www.w3.org/TR/xhtml-basic/xhtml-basic10.dtd">
<html xmlns="http://www.w3.org/1999/xhtml">
<head>
  <meta http-equiv="content-type"
    content="text/html; charset=iso-8859-1"/>
  <title>Basic Location Profile Namespace</title>
</head>
<body>
  <h1>Namespace for Basic Location Profile</h1>
  <h2>urn:ietf:params:xml:ns:basic-location-profiles</h2>
  <p>See <a href="[URL of published RFC]">RFCXXXX
    [NOTE TO IANA/RFC-EDITOR:
      Please replace XXXX with the RFC number of this
      specification.]</a>.</p>
</body>
</html>
END
```

11.6. XCAP Application Usage ID

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

Name of the AUID: geolocation-policy

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

12. Internationalization Considerations

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

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

13. Security Considerations

13.1. Introduction

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

In addition to the authorization policies, mechanisms for obfuscating location information are described. A theoretical treatment of location obfuscation is provided in [duckham05] and in [ifip07]. [duckham05] provides the foundation and [ifip07] illustrates three different types of location obfuscation by enlarging the radius, by shifting the center, and by reducing the radius. The algorithm in Section 6.5.2 for geodetic location information obfuscation uses of these techniques.

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

13.2. Obfuscation

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

Consider the drawing in Figure 1, which shows three ellipses, a dotted area in the middle, and the Target's true location marked as 'x'. The ellipses illustrate the location shapes as received by a potential location recipient over time for requests of a target's location information. Collecting information about the returned

location information over time allows the location recipient to narrow the potential location of the target down to the dotted area in the center of the graph.

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




Figure 1: Obfuscation: A Static Target

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

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

13.3. Algorithm Limitations

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

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

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

In addition to these concerns, the combination of an obscured location with public geographic information (highways, lakes, mountains, cities, etc) may render a much more precise location information than is desired. But even without it, just observing movements, once or multiple times, any obscuring algorithm can leak information about velocities or positions. Suppose a user wants to disclose location information with a radius of r . The privacy region, a circle with that radius, has an area of $A = \pi * r^2$. An adversary, observing the movements, will deduce that the information that the target is, was, or regularly visits, a region of size A_1 , smaller than A . The quotient of the sizes A_1/A should be, even in the worst case, larger than a fixed known number, in order that the user knows what is the maximal information leakage he has. The choices of Section 6.5.2 are such that this maximum leakage can be established: by any statistical procedures, without using external information (highways, etc. as discussed above), the quotient A_1/A is larger than 0.13 (= $1/(5*1.5)$). Thus, for instance, when choosing a provided location of size 1000 km², he will be leaking, in worst case, the location within a region of size 130 km².

13.4. Usability

There is the risk that end users are specifying their location-based policies in such a way that very small changes in location yields a significantly different level of information disclosure. For example, a user might want to set authorization policies differently when they are in a specific geographical area (e.g., at home, in the office). Location might be the only factor in the policy that triggers a very different action and transformation to be executed. The accuracy of location information is not always sufficient to

unequivocally determine whether a location is within a specific boundary [I-D.thomson-geopriv-uncertainty]. In some situations uncertainty in location information could produce unexpected results for end users. Providing adequate user feedback about potential errors arising from these limitation can help prevent unintentional information leakage.

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

13.5. Location Obscuring Limitations

Location obscuring attempts to remove information about the location of a Target. The effectiveness of location obscuring is determined by how much uncertainty a Location Recipient has about the location of the Target. A location obscuring algorithm is effective if the Location Recipient cannot recover a location with better uncertainty than the obscuring algorithm was instructed to add.

Effective location obscuring is difficult. The amount of information that can be recovered by a determined and resourceful Location Recipient can be considerably more than is immediately apparent. A concise summary of the challenges is included in [duckham10].

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

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

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

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

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

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

14. References

14.1. Normative References

- [GML] OpenGIS, "OpenGIS Geography Markup Language (GML) Implementation Specification, Version 3.1.1, OGC 03-105r1",
http://portal.opengeospatial.org/files/?artifact_id=4700,
July 2004.
- [NIMA.TR8350.2-3e] OpenGIS, "US National Imagery and Mapping Agency,
"Department of Defense (DoD) World Geodetic System 1984
(WGS 84), Third Edition, NIMA TR8350.2", , January 2000.
- [OGC-06-103r4] OpenGIS, "OpenGIS Implementation Standard for Geographic
information - Simple feature access - Part 1: Common
architecture",
<http://www.opengeospatial.org/docs/06-103r4.pdf>,
May 2011.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate
Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3688] Mealling, M., "The IETF XML Registry", BCP 81, RFC 3688,
January 2004.
- [RFC4745] Schulzrinne, H., Tschofenig, H., Morris, J., Cuellar, J.,
Polk, J., and J. Rosenberg, "Common Policy: A Document
Format for Expressing Privacy Preferences", RFC 4745,
February 2007.
- [RFC5139] Thomson, M. and J. Winterbottom, "Revised Civic Location
Format for Presence Information Data Format Location
Object (PIDF-LO)", RFC 5139, February 2008.
- [RFC5491] Winterbottom, J., Thomson, M., and H. Tschofenig, "GEOPRIV
Presence Information Data Format Location Object (PIDF-LO)
Usage Clarification, Considerations, and Recommendations",
RFC 5491, March 2009.

14.2. Informative References

- [I-D.thomson-geopriv-geo-shape] Thomson, M., "Geodetic Shapes for the Representation of
Uncertainty in PIDF-LO",
draft-thomson-geopriv-geo-shape-03 (work in progress),

December 2006.

- [I-D.thomson-geopriv-uncertainty]
Thomson, M. and J. Winterbottom, "Representation of Uncertainty and Confidence in PIDF-LO", draft-thomson-geopriv-uncertainty-07 (work in progress), March 2012.
- [RFC2392] Levinson, E., "Content-ID and Message-ID Uniform Resource Locators", RFC 2392, August 1998.
- [RFC2778] Day, M., Rosenberg, J., and H. Sugano, "A Model for Presence and Instant Messaging", RFC 2778, February 2000.
- [RFC3694] Danley, M., Mulligan, D., Morris, J., and J. Peterson, "Threat Analysis of the Geopriv Protocol", RFC 3694, February 2004.
- [RFC4079] Peterson, J., "A Presence Architecture for the Distribution of GEOPRIV Location Objects", RFC 4079, July 2005.
- [RFC4119] Peterson, J., "A Presence-based GEOPRIV Location Object Format", RFC 4119, December 2005.
- [RFC4825] Rosenberg, J., "The Extensible Markup Language (XML) Configuration Access Protocol (XCAP)", RFC 4825, May 2007.
- [RFC5025] Rosenberg, J., "Presence Authorization Rules", RFC 5025, December 2007.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.
- [RFC6280] Barnes, R., Lepinski, M., Cooper, A., Morris, J., Tschofenig, H., and H. Schulzrinne, "An Architecture for Location and Location Privacy in Internet Applications", BCP 160, RFC 6280, July 2011.
- [duckham05]
Duckham, M. and L. Kulik, "A formal model of obfuscation and negotiation for location privacy. In Proc. of the 3rd International Conference PERVASIVE 2005, Munich, Germany", May 2005.
- [duckham10]
Duckham, M., "Moving forward: Location privacy and

location awareness. In Proc. 3rd ACM SIGSPATIAL GIS Workshop on Security and Privacy in GIS and LBS (SPRINGL), ACM.", Nov 2010.

- [ifip07] Ardagna, C., Cremonini, M., Damiani, E., De Capitani di Vimercati, S., and S. Samarati, "Location-privacy protection through obfuscation-based techniques, in: Proceedings of the 21st Annual IFIP WG 11.3 Working Conference on Data and Applications Security, Redondo Beach, CA, USA", July 2007.

Appendix A. Acknowledgments

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

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

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

This document uses text from [I-D.thomson-geopriv-geo-shape]. Therefore, we would like to thank Martin Thomson for his work in [I-D.thomson-geopriv-geo-shape]. We would also like to thank Martin Thomson, Matt Lepinski and Richard Barnes for their comments regarding the geodetic location transformation procedure. Richard provided us with a detailed text proposal.

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

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

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

Appendix B. Pseudo-Code

This section provides an informal description for the algorithm described in Section 6.5.2 in form of pseudo-code.

Constants

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

Parameters

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

Inputs

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

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

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

o : real

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

//      +-----+-----+-----+-----+-----+
//      | min   | max   |               |               |
//      |               |               |               |
```

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

Outputs

```

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

Local Variables

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

temp : Integer[1..8]

Function
choose(Ma, Mb: real * real): real * real;
    // This function chooses either Ma or Mb
    // depending on the parameter 'prob'
    // and on prev-M1, the previous value of M1:
    // If prev-M1 == Ma choose Ma with probability 'prob'
    // If prev-M1 == Mb choose Mb with probability 'prob'
    // Else choose Ma or Mb with probability 1/2
Begin
rand:= Random[0,1];
    // a real random number between 0 and 1
If     prev-M1 == Ma Then
        If rand < prob Then choose := Ma;
                                Else choose := Mb;   EndIf
Elseif prev-M1 == Mb Then
        If rand < prob Then choose := Mb;
                                Else choose := Ma;   EndIf
Else
        If rand < 0.5 Then choose := Ma;
                                Else choose := Mb;   EndIf
End // Function choose

Main // main procedure
Begin
d := radius/1000; // uncertainty, measured in km

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

d2:= d / 110.6;

l := d1*floor(m/d1)
    // "floor" returns the largest integer
    // smaller or equal to a floating point number
r := l+d1;
b := o+d2*floor(n-o/d2);
t := b+d2;

```



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

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

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

End // Main
```

Authors' Addresses

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

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

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

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

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

Email: Jorge.Cuellar@siemens.com

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

Email: jmpolk@cisco.com

John B. Morris, Jr.

Email: ietf@jmorris.org

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

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

GEOPRIV
Internet-Draft
Intended status: Standards Track
Expires: June 22, 2014

M. Thomson
Mozilla
R. Bellis
Nominet UK
December 19, 2013

Location Information Server (LIS) Discovery using IP address and Reverse
DNS
draft-ietf-geopriv-res-gw-lis-discovery-08

Abstract

The residential gateway is a device that has become an integral part of home networking equipment. Discovering a Location Information Server (LIS) is a necessary part of acquiring location information for location-based services. However, discovering a LIS when a residential gateway is present poses a configuration challenge, requiring a method that is able to work around the obstacle presented by the gateway.

This document describes a solution to this problem. The solution provides alternative domain names as input to the LIS discovery process based on the network addresses assigned to a Device.

Status of This Memo

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

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

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

This Internet-Draft will expire on June 22, 2014.

Copyright Notice

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

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

Table of Contents

1. Introduction	2
2. Conventions used in this document	3
3. Problem Statement	3
3.1. Residential Gateway	5
3.2. Residential Gateway Security Features	6
4. IP-based DNS Solution	6
4.1. Identification of IP Addresses	7
4.2. Domain Name Selection	8
4.3. Shortened DNS Names	8
4.4. When To Use The Reverse DNS Method	9
4.5. Private Address Spaces	9
4.6. Necessary Assumptions and Restrictions	10
4.7. Failure Modes	10
4.8. Deployment Considerations	11
5. IANA Considerations	12
6. Privacy Considerations	12
7. Security Considerations	12
8. IAB Considerations	13
9. Acknowledgements	15
10. References	15
10.1. Normative References	15
10.2. Informative References	16
Authors' Addresses	17

1. Introduction

A Location Information Server (LIS) is a service provided by an access network. The LIS uses knowledge of the access network topology and other information to generate location information for Devices. Devices within an access network are able to acquire location information from a LIS.

The relationship between a Device and an access network might be transient. Configuration of the correct LIS at the Device ensures that accurate location information is available. Without location information, some network services are not available.

The configuration of a LIS IP address on a Device requires some automated process. This is particularly relevant when one considers that Devices might move between different access networks served by different LISs. LIS Discovery [RFC5986] describes a method that employs the Dynamic Host Configuration Protocol (DHCPv4 [RFC2131], DHCPv6 [RFC3315]) as input to U-NAPTR [RFC4848] discovery.

A residential gateway, or home router, provides a range of networking functions for Devices within the network it serves. Unfortunately in most cases these functions effectively prevent the successful use of DHCP for LIS discovery.

One drawback with DHCP is that universal deployment of a new option takes a considerable amount of time. Often, networking equipment needs to be updated in order to support the new option. Of particular concern are the millions of residential gateway devices used to provide Internet access to homes and businesses. While [RFC5986] describes functions that can be provided by residential gateways to support LIS discovery, gateways built before the publication of this specification are not expected (and are likely not able) to provide these functions.

This document explores the problem of configuring Devices with a LIS address when a residential gateway is interposed between the Device and access network. Section 3 defines the problem and Section 4 describes a method for determining a domain name that can be used for discovery of the LIS.

In some cases, the solution described in this document is based on a UNilateral Self-Address Fixing (UNSAF) [RFC3424] method. For those cases, this solution is considered transitional until such time as the recommendations for residential gateways in [RFC5986] are more widely deployed. Considerations relating to UNSAF applications are described in Section 8.

2. Conventions used in this document

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

This document uses terminology established in [RFC6280] and [RFC5012]. The terms Device and LIS are capitalized throughout when they are used to identify the roles defined in [RFC6280].

3. Problem Statement

Figure 1 shows a simplified network topology for fixed wire-line Internet access. This arrangement is typical when wired Internet access is provided. The diagram shows two network segments: the access network provided by an internet service provider (ISP), and the residential network served by the residential gateway.

There are a number of variations on this arrangement, as documented in Section 3.1 of [RFC5687]. In each of these variations the goal of LIS discovery is to identify the LIS in the access network.




Figure 1: Simplified Network Topology

A particularly important characteristic of this arrangement is the relatively small geographical area served by the residential gateway. Given a small enough area, it is reasonable to delegate the responsibility for providing Devices within the residential network with location information to the ISP. The ISP is able to provide location information that identifies the residence, which should be adequate for a wide range of purposes.

A residential network that covers a larger geographical area might require a dedicated LIS, a case that is outside the scope of this document.

The goal of LIS discovery is to identify a LIS that is able to provide the Device with accurate location information. In the network topology described, this means identifying the LIS in the access network. The residential gateway is a major obstacle in achieving this goal.

3.1. Residential Gateway

A residential gateway can encompass several different functions including: modem, Ethernet switch, wireless access point, router, network address translation (NAT), DHCP server, DNS relay and firewall. Of the common functions provided, the NAT function of a residential gateway has the greatest impact on LIS discovery.

An ISP is typically parsimonious about their IP address allocations; each customer is allocated a limited number of IP addresses. Therefore, NAT is an extremely common function of gateways. NAT enables the use of multiple Devices within the residential network. However NAT also means that Devices within the residence are not configured by the ISP directly.

When it comes to discovering a LIS, the fact that Devices are not configured by the ISP causes a significant problem. Configuration is the ideal method of conveying the information necessary for discovery. Devices attached to residential gateways are usually given a generic configuration that includes no information about the ISP network. For instance, DNS configuration typically points to a DNS relay on the gateway device. This approach ensures that the local network served by the gateway is able to operate without a connection to the ISP, but it also means that Devices are effectively ignorant of the ISP network.

[RFC5986] describes several methods that can be applied by a residential gateway to assist Devices in acquiring location information. For instance, the residential gateway could forward LIS address information to hosts within the network it serves.

Unfortunately, such an active involvement in the discovery process only works for new residential gateway devices that implement those recommendations.

Where residential gateways already exist, direct involvement of the gateway in LIS discovery requires that the residential gateway be updated or replaced. The cost of replacement is difficult to justify to the owner of the gateway, especially when it is considered that the gateway still fills its primary function: Internet access. Furthermore, updating the software in such devices is not feasible in many cases. Even if software updates were made available, many residential gateways cannot be updated remotely, inevitably leading to some proportion that is not updated.

This document therefore describes a method that can be used by Devices to discover their LIS without any assistance from the network.

3.2. Residential Gateway Security Features

A network firewall function is often provided by residential gateways as a security measure. Security features like intrusion detection systems help protect users from attacks. Amongst these protections is a port filter that prevents both inbound and outbound traffic on certain TCP and UDP ports. Therefore, any solution needs to consider the likelihood of traffic being blocked.

4. IP-based DNS Solution

LIS discovery [RFC5986] uses a DNS-based Dynamic Delegation Discovery Service (DDDS) system as the basis of discovery. Input to this process is a domain name. Use of DHCP for acquiring the domain name is specified, but alternative methods of acquisition are permitted.

This document specifies a means for a Device to discover several alternative domain names that can be used as input to the DDDS process. These domain names are based on the IP address of the Device. Specifically, the domain names are a portion of the reverse DNS trees - either the ".in-addr.arpa." or ".ip6.arpa." tree.

The goal of this process is to make a small number of DDDS queries in order to find a LIS. After LIS discovery using the DHCP-based process in [RFC5986] has failed, a Device can:

1. Collect a set of IP addresses that refer to the Device (Section 4.1).

2. Convert each IP address into DNS names in the "in-addr.arpa." or "ip6.arpa." tree (Section 4.2).
3. Perform the DDDS process for LIS discovery on those DNS names ([RFC5986]).
4. Shorten the DNS names by some number of labels and repeat the DDDS process (Section 4.3).

A Device might be reachable at one of a number of IP addresses. In the process described, a Device first identifies each IP address that it is potentially reachable from. From each of these addresses, the Device then selects up to three domain names for use in discovery. These domain names are then used as input to the DDDS process.

4.1. Identification of IP Addresses

A Device identifies a set of potential IP addresses that currently result in packets being routed to it. These are ordered by proximity, with those addresses that are used in adjacent network segments being favoured over those used in public or remote networks. The first addresses in the set are those that are assigned to local network interfaces.

A Device can use the Session Traversal Utilities for NAT (STUN) [RFC5389] mechanism to determine its public reflexive transport address. The host uses the "Binding Request" message and the resulting "XOR-MAPPED-ADDRESS" parameter that is returned in the response.

Alternative methods for determining other IP addresses MAY be used by the Device. Port Control Protocol (PCP) [RFC6887], Universal Plug and Play (UPnP) [UPnP-IGD-WANIPConnection1] and NAT Port Mapping Protocol (NAT-PMP) [I-D.cheshire-nat-pmp] are both able to provide the external address of a residential gateway device when enabled. These as well as proprietary methods for determining other addresses might also be available. Because there is no assurance that these methods will be supported by any access network, these methods are not mandated. Note also that in some cases, methods that rely on the view of the network from the residential gateway device could reveal an address in a private address range (see Section 4.6).

In many instances, the IP address produced might be from a private address range. For instance, the address on a local network interface could be from a private range allocated by the residential gateway. In other cases, methods that rely on the view of the network (UPnP, NAT-PMP) from the residential gateway device could reveal an address in a private address range if the access network

also uses NAT. For a private IP address, the derived domain name is only usable where the DNS server used contains data for the corresponding private IP address range.

4.2. Domain Name Selection

The domain name selected for each resulting IP address is the name that would be used for a reverse DNS lookup. The domain name derived from an IP version 4 address is in the ".in-addr.arpa." tree and follows the construction rules in Section 3.5 of [RFC1035]. The domain name derived from an IP version 6 address is in the ".ip6.arpa." tree and follows the construction rules in Section 2.5 of [RFC3596].

4.3. Shortened DNS Names

Additional domain names are added to allow for a single DNS record to cover a larger set of addresses. If the search on the domain derived from the full IP address does not produce a NAPTR record with the desired service tag (e.g., "LIS:HELD"), a similar search is repeated based on a shorter domain name, using a part of the IP address:

- o For IP version 4, the resulting domain name SHOULD be shortened successively by one and two labels and the query repeated. This corresponds to a search on a /24 or /16 network prefix. This allows for fewer DNS records in the case where a single access network covering an entire /24 or /16 network is served by the same LIS.
- o For IP version 6, the resulting domain name SHOULD be shortened successively by 16, 18, 20 and 24 labels and the query repeated. This corresponds to a search on a /64, /56, /48 or /32 network prefix.

This set of labels is intended to provide network operators with a degree of flexibility in where LIS discovery records can be placed without significantly increasing the number of DNS names that are searched. This does not attach any other significance to these specific zone cuts, or create a classful addressing hierarchy based on the reverse DNS tree.

For example, the IPv4 address "192.0.2.75" could result in queries to:

- o 75.2.0.192.in-addr.arpa.
- o 2.0.192.in-addr.arpa.

- o 0.192.in-addr.arpa.

Similarly, the IPv6 address "2001:DB8::28e4:3a93:4429:dfb5" could result in queries to:

- o 5.b.f.d.9.2.4.4.3.9.a.3.4.e.8.2.0.0.0.0.0.0.0.0.0.8.b.d.0.1.0.0.2.ip6.arpa.
- o 0.0.0.0.0.0.0.0.0.8.b.d.0.1.0.0.2.ip6.arpa.
- o 0.0.0.0.0.0.8.b.d.0.1.0.0.2.ip6.arpa.
- o 0.0.0.0.8.b.d.0.1.0.0.2.ip6.arpa.
- o 8.b.d.0.1.0.0.2.ip6.arpa.

The limited number of labels by which each name is shortened is intended to limit the number of DNS queries performed by Devices. If no LIS is discovered by this method, the result will be that no more than five U-NAPTR resolutions are invoked for each IP address.

4.4. When To Use The Reverse DNS Method

The DHCP method described in [RFC5986] MUST be attempted on all local network interfaces before attempting this method. This method is employed either because DHCP is unavailable, when the DHCP server does not provide a value for the access network domain name option, or if a request to the resulting LIS results in a HELD "notLocatable" error or equivalent.

4.5. Private Address Spaces

Addresses from a private use address space can be used as input to this method. In many cases, this applies to addresses defined in [RFC1918], though other address ranges could have limited reachability where this advice also applies. This is only possible if a DNS server with a view of the same address space is used. Public DNS servers cannot provide useful records for private addresses.

Using an address from a private space in discovery can provide a more specific answer if the DNS server has records for that space. Figure 2 shows a network configuration where addresses from an ISP network could better indicate the correct LIS. Records in DNS B can be provided for the 10.0.0.0/8 range, potentially dividing that range so that a more local LIS can be selected.




Figure 2: Address Space Example

The goal of automatic DNS configuration is usually to select a local DNS, which suits configurations like the one shown. However, use of public DNS or STUN servers means that a public IP address is likely to be found. For STUN in particular, selecting a public server minimizes the need for reconfiguration when a Device moves. Adding records for the public address space used by an access network ensures that the discovery process succeeds when a public address is used.

4.6. Necessary Assumptions and Restrictions

When used by a Device for LIS discovery this is an UNSAF application and is subject to the limitations described in Section 8.

It is not necessary that the IP address used is unique to the Device, only that the address can be somehow related to the Device or the access network that serves the Device. This allows a degree of flexibility in determining this value, although security considerations (Section 7) might require that the address be verified to limit the chance of falsification.

This solution assumes that the public reflexive transport address used by a Device is in some way controlled by the access network provider, or some other related party. This implies that the corresponding ".in-addr.arpa." or ".ip6.arpa." record can be updated by that entity to include a useful value for the LIS address.

4.7. Failure Modes

Successful use of private addresses relies on a DNS server that has records for the address space that is used. Using a public IP address increases the likelihood of this. This document relies on STUN to provide the Device with a public reflexive transport address. Configuration of a STUN server is necessary to ensure that this is successful.

In cases where a virtual private network (VPN) or other tunnel is used, the entity providing a public IP address might not be able to provide the Device with location information. It is assumed that this entity is able to identify this problem and indicate this to the Device (using the "notLocatable" HELD error, or similar). This problem is described in more detail in [RFC5985].

4.8. Deployment Considerations

An access network provider SHOULD provide NAPTR records for each public IP address that is used for Devices within the access network.

Any DNS server internal to a NAT SHOULD also include records for the private address range. These records might only be provided to clients making requests from the private address range. Doing so allows clients within the private address range to discover a LIS based on their IP address prior to any address translation. In geographically distributed networks that use a private address range, this enables the use of a different LIS for different locations, based on the IP address range used at each location. Use of a public, translated IP address for the network can still work, but it might result in a suboptimal LIS selection.

A network that operates network address translation SHOULD provide NAPTR records that reference a LIS endpoint with a public address. This requires the reservation of an IP and port for the LIS. To ensure requests toward the LIS from within the private address space do not traverse the NAT and have source addresses mapped by the NAT, networks can provide direct route to the LIS. Clients that perform discovery based on public DNS or STUN servers are thereby easier to trace based on source address information.

NAPTR records can be provided for individual IP addresses. To limit the proliferation of identical records, a single record can be placed at higher nodes of the tree (corresponding to /24 and /16 for IPv4; /64, /56, /48 and /32 for IPv6). A record at a higher point in the tree (those with a shorter prefix) applies to all addresses lower in the tree (those with a longer prefix); records at the lower point override those at higher points, thus allowing for exceptions to be specified.

5. IANA Considerations

This document has no IANA actions.

6. Privacy Considerations

As with all uses of geolocation information, it is very important that measures be taken to ensure that location information is not provided to unauthorized parties. The mechanism defined in this document is focused on the case where a device is learning its own location, so that it can provide that location information to applications. We assume that the device learning its own location is not a privacy risk. There are then two remaining privacy risks: The use of geolocation by applications, and abuse of the location configuration protocol.

The privacy considerations around the use of geolocation by applications vary considerably by application context. A framework for location privacy in applications is provided in [RFC6280].

The mechanism specified in this document allows a device to discover its local LIS, from which it then acquires its location using a Location Configuration Protocol [RFC5687]. If an unauthorized third party can spoof the LCP to obtain a target's location information, then the mechanism in this document could allow them to discover the proper server to attack for a given IP address. Thus, it is important that a LIS meet the security requirements of the LCP it implements. For HELD, these requirements are laid out in Section 9 of [RFC5985].

A Device that discovers a LIS using the methods in this document MUST NOT provide that LIS with additional information that might reveal its position, such as the location measurements described in [I-D.ietf-geopriv-held-measurements], unless it has a secondary method for determining the authenticity of the LIS, such as a white list.

7. Security Considerations

The security considerations described in [RFC5986] apply to the discovery process as a whole. The primary security concern is with the potential for an attacker to impersonate a LIS.

The added ability for a third party to discover the identity of a LIS does not add any concerns, since the identity of a LIS is considered public information.

In addition to existing considerations, this document introduces further security considerations relating to the identification of the IP address. It is possible that an attacker could attempt to provide a falsified IP address in an attempt to subvert the rest of the process.

[RFC5389] describes attacks where an attacker is able to ensure that a Device receives a falsified reflexive address. An on-path attacker might be able to ensure that a falsified address is provided to the Device. Even though STUN messages are protected by a STUN MESSAGE-INTEGRITY attribute, which is an HMAC that uses a shared secret, an on-path attacker can capture and modify packets, altering source and destination addresses to provide falsified addresses.

This attack could result in an effective means of denial of service, or a means to provide a deliberately misleading service. Notably, any LIS that is identified based on a falsified IP address could still be a valid LIS for the given IP address, just not one that is useful for providing the Device with location information. In this case, the LIS provides a HELD "notLocatable" error, or an equivalent. If the falsified IP address is under the control of the attacker, it is possible that misleading (but verifiable) DNS records could indicate a malicious LIS that provides false location information.

In all cases of falsification, the best remedy is to perform some form of independent verification of the result. No specific mechanism is currently available to prevent attacks based on falsification of reflexive addresses; it is suggested that Devices attempt to independently verify that the reflexive transport address provided is accurate. An independent, trusted source of location information could aid in verification, even if the trusted source is unable to provide information with the same degree of accuracy as the discovered LIS.

Use of private address space effectively prevents use of the usual set of trust anchors for DNSSEC. Only a DNS server that is able to see the same private address space can provide useful records. A Device that relies on DNS records in the private address space portion of the ".in-addr.arpa." or ".ip6.arpa." trees MUST either use an alternative trust anchor for these records or rely on other means of ensuring the veracity of the DNS records.

8. IAB Considerations

The IAB has studied the problem of Unilateral Self-Address Fixing (UNSAF) [RFC3424], which is the general process by which a client attempts to determine its address in another realm on the other side of a NAT through a collaborative protocol reflection mechanism, such as STUN.

This section only applies to the use of this method of LIS discovery by Devices and does not apply to its use for third-party LIS discovery.

The IAB requires that protocol specifications that define UNSAF mechanisms document a set of considerations.

1. Precise definition of a specific, limited-scope problem that is to be solved with the UNSAF proposal.

Section 3 describes the limited scope of the problem addressed in this document.

2. Description of an exit strategy/transition plan.

[RFC5986] describes behaviour that residential gateways require in order for this short term solution to be rendered unnecessary. When implementations of the recommendations in LIS discovery are widely available, this UNSAF mechanism can be made obsolete.

3. Discussion of specific issues that may render systems more "brittle".

A description of the necessary assumptions and limitations of this solution are included in Section 4.6.

Use of STUN for discovery of a reflexive transport address is inherently brittle in the presence of multiple NATs or address realms. In particular, brittleness is added by the requirement of using a DNS server that is able to view the address realm that contains the IP address in question. If address realms use overlapping addressing space, then there is a risk that the DNS server provides information that is not useful to the Device.

4. Identify requirements for longer term, sound technical solutions; contribute to the process of finding the right longer term solution.

A longer term solution is already provided in [RFC5986]. However, that solution relies on widespread deployment. The UNSAF solution provided here is provided as an interim solution that enables LIS access for Devices that are not able to benefit from deployment of the recommendations in [RFC5986].

5. Discussion of the impact of the noted practical issues with existing deployed NATs and experience reports.

The UNSAF mechanism depends on the experience in deployment of STUN [RFC5389]. On the whole, existing residential gateway devices are able to provide access to STUN and DNS service reliably, although regard should be given to the size of the DNS response (see [RFC5625]).

9. Acknowledgements

Richard Barnes provided the text in Section 6.

10. References

10.1. Normative References

- [RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, November 1987.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3596] Thomson, S., Huitema, C., Ksinant, V., and M. Souissi, "DNS Extensions to Support IP Version 6", RFC 3596, October 2003.
- [RFC5986] Thomson, M. and J. Winterbottom, "Discovering the Local Location Information Server (LIS)", RFC 5986, September 2010.
- [I-D.ietf-geopriv-held-measurements] Thomson, M. and J. Winterbottom, "Using Device-provided Location-Related Measurements in Location Configuration Protocols", draft-ietf-geopriv-held-measurements-09 (work in progress), September 2013.

10.2. Informative References

- [RFC1918] Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996.
- [RFC2131] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, March 1997.
- [RFC3315] Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, July 2003.
- [RFC3424] Daigle, L. and IAB, "IAB Considerations for UNilateral Self-Address Fixing (UNSAF) Across Network Address Translation", RFC 3424, November 2002.
- [RFC3693] Cuellar, J., Morris, J., Mulligan, D., Peterson, J., and J. Polk, "Geopriv Requirements", RFC 3693, February 2004.
- [RFC4848] Daigle, L., "Domain-Based Application Service Location Using URIs and the Dynamic Delegation Discovery Service (DDDS)", RFC 4848, April 2007.
- [RFC5012] Schulzrinne, H. and R. Marshall, "Requirements for Emergency Context Resolution with Internet Technologies", RFC 5012, January 2008.
- [RFC5389] Rosenberg, J., Mahy, R., Matthews, P., and D. Wing, "Session Traversal Utilities for NAT (STUN)", RFC 5389, October 2008.
- [RFC5687] Tschofenig, H. and H. Schulzrinne, "GEOPRIV Layer 7 Location Configuration Protocol: Problem Statement and Requirements", RFC 5687, March 2010.
- [RFC6280] Barnes, R., Lepinski, M., Cooper, A., Morris, J., Tschofenig, H., and H. Schulzrinne, "An Architecture for Location and Location Privacy in Internet Applications", BCP 160, RFC 6280, July 2011.
- [RFC6887] Wing, D., Cheshire, S., Boucadair, M., Penno, R., and P. Selkirk, "Port Control Protocol (PCP)", RFC 6887, April 2013.
- [UPnP-IGD-WANIPConnection1]
UPnP Forum, "Internet Gateway Device (IGD) Standardized Device Control Protocol V 1.0: WANIPConnection:1 Service

Template Version 1.01 For UPnP Version 1.0", DCP 05-001,
Nov 2001.

[I-D.cheshire-nat-pmp]
Cheshire, S. and M. Krochmal, "NAT Port Mapping Protocol
(NAT-PMP)", draft-cheshire-nat-pmp-07 (work in progress),
January 2013.

[RFC5625] Bellis, R., "DNS Proxy Implementation Guidelines", BCP
152, RFC 5625, August 2009.

[RFC5985] Barnes, M., "HTTP-Enabled Location Delivery (HELD)", RFC
5985, September 2010.

Authors' Addresses

Martin Thomson
Mozilla
Suite 300
650 Castro Street
Mountain View, CA 94041
US

Email: martin.thomson@gmail.com

Ray Bellis
Nominet UK
Edmund Halley Road
Oxford OX4 4DQ
United Kingdom

Phone: +44 1865 332211
Email: ray.bellis@nominet.org.uk
URI: <http://www.nominet.org.uk/>