Internet Engineering Task Force (IETF)
Request for Comments: 5808
Category: Informational
ISSN: 2070-1721

R. Marshall, Ed.
TCS

May 2010

Requirements for a Location-by-Reference Mechanism

Abstract

This document defines terminology and provides requirements relating to the Location-by-Reference approach using a location Uniform Resource Identifier (URI) to handle location information within signaling and other Internet messaging.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at http://www.rfc-editor.org/info/rfc5808.

Table of Contents

1. Introduction 3
2. Terminology 5
3. Overview of Location-by-Reference 6
   3.1. Location URI Usage 7
   3.2. Location URI Expiration 8
   3.3. Location URI Authorization 8
   3.4. Location URI Construction 9
4. High-Level Requirements 9
   4.1. Requirements for a Location Configuration Protocol 9
   4.2. Requirements for a Location Dereference Protocol 11
5. Security Considerations 12
6. Acknowledgements 13
7. References 13
   7.1. Normative References 13
   7.2. Informative References 13

1. Introduction

All location-based services rely on ready access to location information. Location information can be used in either a direct, Location-by-Value (LbyV) approach or an indirect, Location-by-Reference (LbyR) approach.

For LbyV, location information is conveyed directly in the form of a Presence Information Data Format Location Object (PIDF-LO) [RFC4119]. Using LbyV might be either infeasible or undesirable in some circumstances. There are cases where LbyR is better able to address location requirements for a specific architecture or application. This document provides a list of requirements for use with the LbyR approach, and leaves the LbyV model explicitly out of scope.

As justification for an LbyR model, consider the circumstance that in some mobile networks it is not efficient for the end host to periodically query the Location Information Server (LIS) for up-to-date location information. This is especially the case when power availability is a constraint or when a location update is not immediately needed. Furthermore, the end host might want to delegate the task of retrieving and publishing location information to a third party, such as to a presence server. Additionally, in some deployments, the network operator may not want to make location information widely available. These kinds of location scenarios form the basis of motivation for the LbyR model.

The concept of an LbyR mechanism is simple. An LbyR is made up of a URI scheme, a domain, and a randomized component. This combination of data elements, in the form of a URI, is referred to specifically as a "location URI".

A location URI is thought of as a reference to the current location of the Target, yet the location value might remain unchanged over specific intervals of time for several reasons. The type of location information returned as part of the dereferencing step may, for example, be influenced by the following factors:

- Limitations in the process used to generate location information mean that cached location might be used.

- Policy constraints may dictate that the location provided remains fixed over time for specified Location Recipients. Without additional information, a Location Recipient cannot assume that the location information provided by any location URI is static, and will never change.

The LbyR mechanism works according to an information life cycle. Within this life cycle, location URIs are considered temporary identifiers, each undergoing the following uses: Creation; Distribution; Conveyance; Dereference; and Termination. The use of a location URI according to these various states is generally applied in one of the following ways:

1. Creation of a location URI, within a location server, based on some request for its creation.

2. Distribution of a location URI, via a Location Configuration Protocol, between a Target and a location server.

3. Conveyance, applied to LbyR, for example in SIP (Session Initiation Protocol), is the transporting of the location URI, in this case, between any successive signaling nodes.

4. Dereference of a location URI, a request/response between a client having a location URI and a location server holding the location information that the location URI references.

5. Termination of a location URI, due to either expiration or cancellation within a location server, and that is based on a Target cancellation request or some other action, such as timer expiration.

Note that this document makes no functional differentiation between a Location Server (LS), per [RFC3693], and a Location Information Server (LIS), as shown in [RFC5687], but may refer to either of them as a location server interchangeably.

Location determination, as distinct from location configuration or dereferencing, often includes topics related to manual provisioning processes, automated location calculations based on a variety of measurement techniques, and/or location transformations (e.g., geo-coding), and is beyond the scope of this document.

Location Conveyance for either LbyR or LbyV, as defined within SIP signaling is considered out of scope for this document. (See [LOC-CONVEY] for an explanation of location conveyance for either LbyR or LbyV scenarios.)

Except for location conveyance, the above stages in the LbyR life cycle fall into one of two general categories of protocols, either a Location Configuration Protocol or a Location Dereference Protocol. The stages of LbyR Creation, Distribution, and Termination, are each

found within the set of Location Configuration Protocols (LCPs). The Dereference stage belongs solely to the set of Location Dereference Protocols.

The issues around location configuration protocols have been documented in a location configuration protocol problem statement and requirements document [RFC5687]. There are currently several examples of documented location configuration protocols, namely DHCP [DHCP-LOC-URI], LLDP-MED [LLDP-MED], and HELD [HELD].

For dereferencing a location URI, depending on the type of reference used, such as a HTTP/HTTPS or SIP Presence URI, different operations can be performed. While an HTTP/HTTPS URI can be resolved to location information, a SIP Presence URI provides further benefits from the SUBSCRIBE/NOTIFY concept that can additionally be combined with location filters [LOC-FILTERS].

The structure of this document includes terminology, Section 2, followed by a discussion of the basic elements that surround how a location URI is used. These elements, or actors, are discussed in an overview section, Section 3, accompanied by a graph, associated processing steps, and a brief discussion around the use, expiration, authorization, and construction of location URIs.

Requirements are outlined accordingly, separated as location configuration requirements, Section 4.1, and location dereference requirements, Section 4.2.

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], with the important qualification that, unless otherwise stated, these terms apply to the design of the Location Configuration Protocol and the Location Dereferencing Protocol, not its implementation or application.

This document reuses the terminology of [RFC3693], such as Location Server (LS), Location Recipient (LR), Rule Maker (RM), Target, and Location Object (LO). Furthermore, the following terms are defined in this document:

Location-by-Value (LbyV): Using location information in the form of a location object (LO), such as a PIDF-LO.

Location-by-Reference (LbyR): Representing location information indirectly using a location URI.

Location Configuration Protocol: A protocol that is used by a Target to acquire either a location object or a location URI from a location configuration server, based on information unique to the Target.

Location Dereference Protocol: A protocol that is used by a client to query a location server, based on the location URI input, and that returns location information.

Location URI: As defined within this document, an identifier that serves as a reference to location information. A location URI is provided by a location server, and is later used as input by a dereference protocol to retrieve location information.

3. Overview of Location-by-Reference

This section describes the entities and interactions involved in the LbyR model.

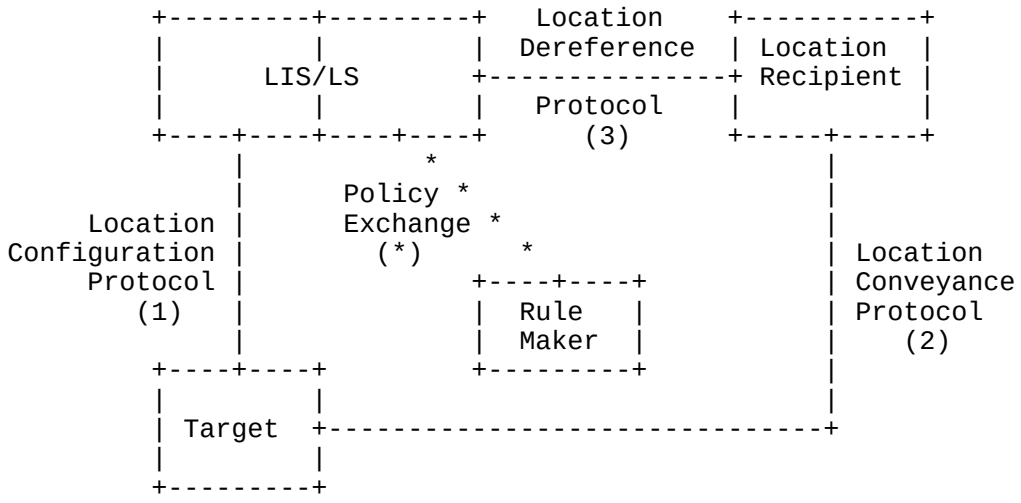

Figure 1: Location Reference Entities and Interactions

Figure 1 shows the assumed communication model for both a Layer 7 location configuration protocol and a location dereference protocol.

(1) The Target (an end device) uses a location configuration protocol to acquire a location reference from a LIS, which acts as (or is able to access) an LS.

In the case where the Target is also a Rule Maker, the location configuration protocol can be used to convey policy information.

In the case where possession of a location URI is the only required form of authorization (see Section 3.3), a policy is implied whereby any requester is granted access to location information. This does not preclude other means of providing authorization policies.

A Target could also acquire a location URI from the LS directly using alternative means, for example, the acquisition of a presence Address of Record (AoR) to be used for location information, in which case, it could be regarded as a location URI.

(2) The Target conveys the location URI to the Location Recipient (interface out of scope).

(3) The Location Recipient dereferences the location URI to acquire location information from the LS.

The LS controls access to location information based on the policy provided by the Rule Maker.

Note A. There is no requirement for using the same protocol in (1) and (3).

Note B. Figure 1 includes the interaction between the owner of the Target and the LIS to obtain Rule Maker policies. This interaction needs to happen before the LIS will authorize anything other than what is allowed based on default policies in order to dereference a location request of the Target. This communication path is out of scope for this document.

Note C. The Target might take on the role of the Location Recipient, in which case, it could attempt to dereference the location URI itself, in order to obtain its own location information.

3.1. Location URI Usage

An example scenario of how the above location configuration and location dereference steps might work using SIP is where a Target obtains a location URI in the form of a subscription URI (e.g., a SIP URI) via a location configuration protocol. In this case, the Target is the same as the Recipient; therefore, the Target can subscribe to the URI in order to be notified of its current location based on subscription parameters. In the example, parameters are set up for a specific Target/Recipient along with an expressed geospatial boundary, so that the Target/Recipient receives an updated location notification once the boundary is crossed (see [LOC-FILTERS]).

3.2. Location URI Expiration

Location URIs may have an expiry associated with them, primarily for security considerations, and generally in order for the LIS to keep track of the location URIs that have been handed out, to know whether a location URI is still valid once the LIS receives it in a request, and for preventing a recipient of such a URI from being able to (in some cases) permanently track a host. Expiration of a location URI limits the time that accidental leaking of a location URI introduces. Other justifications for expiration of location URIs include the ability for a LIS to do garbage collection.

3.3. Location URI Authorization

How a location URI will ultimately be used within the dereference step is an important consideration at the time the location URI is requested via a location configuration protocol. The process of dereferencing location URIs will be influenced by the specific authorization model applied by the Location Information Server and the URI scheme that indicates the protocol to be used to resolve the reference to a location object.

Location URIs manifest themselves in a few different forms. The different ways that a location URI can be represented are based on local policy, and are depicted in the following four scenarios.

1. No location information included in the URI: As is typical, a location URI is used to get location information. However, in this case, the URI representation itself does not need to reveal any specific information at all. Location information is acquired by the dereferencing operation using a location URI.

2. URI does not identify a Target: By default, a location URI MUST NOT reveal any information about the Target other than location information. This is true for the URI itself (or in the document acquired by dereferencing), unless policy explicitly permits otherwise.

3. Access control authorization model: If this model is used, the location URI MUST NOT include any location information in its representation. Location URIs operating under this model could be widely published to recipients that are not authorized to receive this information.

4. Possession authorization model (the URI itself is a secret): If this model is used, the location URI is confidential information shared between the LIS/LS, the Target, and all authorized Location Recipients. In this case, possession implies

authorization. Because knowledge of the location URI is used to authenticate and authorize access to location information, the URI needs to include sufficient randomness to make guessing its value difficult. A possession model URI can include location information in its representation.

3.4. Location URI Construction

Given scenarios 2 and 4, above, and depending on local policy, a location URI may be constructed in such a way as to make it difficult to guess. Accordingly, the form of the URI is then constrained by the degree of randomness and uniqueness applied to it. In this case, it may be important to protect the actual location information from inspection by an intermediate node. Construction of a location URI in such a way as to not reveal any Target-specific information (e.g., user or device information), with the goal of making the location URI appear bland, uninteresting, and generic, may be helpful to some degree in order to keep location information more difficult to detect. Thus, obfuscating the location URI in this way may provide some level of safeguard against the undetected inspection and unintended use of what would otherwise be evident location information, since it forces a dereference operation at the location dereference server, an important step for the purpose of providing statistics, audit trails, and general logging for many different kinds of location-based services.

4. High-Level Requirements

This document outlines the requirements for a Location by Reference mechanism that can be used by a number of underlying protocols. Requirements here address two general types of such protocols, a general location configuration protocol and a general location dereferencing protocol.

The requirements are broken into two sections.

4.1. Requirements for a Location Configuration Protocol

Below, we summarize high-level design requirements needed for a location-by-reference mechanism as used within the location configuration protocol.

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

Motivation: A standardized location reference mechanism increases interoperability.

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

Motivation: A location URI may not intend to represent a location forever, and the identifier eventually may need to be recycled, or may be subject to a specific window of validity, after which the location reference fails to yield a location, or the location is determined to be kept confidential.

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

Motivation: If the Target determines that a location URI should no longer be used to dereference a location, then there should be a way to request that the location URI be nullified.

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

Motivation: It is important to keep any location information masked from a casual observing node.

C5. Target identity protection: The location URI MUST NOT contain information that identifies the Target (e.g., user or device). Examples include phone extensions, badge numbers, and first or last names.

Motivation: It is important to protect caller identity or contact address from being included in the form of the location URI itself when it is generated.

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

Motivation: The Target requesting a location URI may request a location URI that has a 'one-time-use' only characteristic, as opposed to a location URI having multiple reuse capability. This would allow the server to return an error with or without location information during the subsequent dereference operation.

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

Motivation: The Rule Maker has to be in control of how much information is revealed during the dereferencing step as part of the privacy features.

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

Motivation: Location URIs should be constructed in such a way that an adversary cannot guess them and dereference them without having previously obtained them from the Target.

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

Motivation: Users don’t always have such strict privacy requirements, but may opt to specify their own location URI or components to be included within a location URI.

4.2. Requirements for a Location Dereference Protocol

Below, we summarize high-level design requirements needed for a location-by-reference mechanism as used within the location dereference protocol.

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

Motivation: It is required that there be consistency of use between location URI formats used in a configuration protocol and those used by a dereference protocol.

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

Motivation: Although the implementations must support authentication of both parties, any given transaction has the option not to authenticate one or both parties.

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

Motivation: This is in order to ensure that adequate privacy rules can be adhered to, since the PIDF-LO format comprises the necessary structures to maintain location privacy.

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

Motivation: Through dereference server configuration, for example, it may be useful to not only allow more than one dereference request, but, in some cases, to also limit the number of dereferencing attempts by a client.

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

Motivation: The location URI indicates what type of security protocol has to be provided. An example is a location URI using a HTTPS URI scheme.

5. Security Considerations

The method of constructing the location URI to include randomized components helps to prevent adversaries from obtaining location information without ever retrieving a location URI. In the possession model, a location URI, regardless of its construction, if made publicly available, implies no safeguard against anyone being able to dereference and get the location. Care has to be paid when distributing such a location URI to the trusted location recipients. When this aspect is of concern, the authorization model has to be chosen. Even in this model, care has to be taken on how to construct the authorization policies to ensure that only those parties have access to location information that are considered trustworthy enough to enforce the basic rule set that is attached to location information in a PIDF-LO document.

Any location URI, by necessity, indicates the server (name) that hosts the location information. Knowledge of the server in some specific domain could therefore reveal something about the location of the Target. This kind of threat may be mitigated somewhat by introducing another layer of indirection: namely the use of a (remote) presence server.

A covert channel for protocol message exchange is an important consideration, given an example scenario where user A subscribes to location information for user B, then every time A gets a location update, an (external) observer of the subscription notification may

know that B has moved. One mitigation of this is to have periodic notification, so that user B may appear to have moved even when static.

6. Acknowledgements

I would like to thank the present IETF GEOPRIV working group chairs, Alissa Cooper and Richard Barnes, past chairs, Robert Sparks, Andy Newton, Allison Mankin, and Randall Gellens, who established a design team that initiated this requirements work. I'd also like to thank those original design team participants for their inputs, comments, and insightful reviews. The design team included the following folks: Richard Barnes, Martin Dawson, Keith Drage, Randall Gellens, Ted Hardie, Cullen Jennings, Marc Linsner, Rohan Mahy, Allison Mankin, Andrew Newton, Jon Peterson, James M. Polk, Brian Rosen, John Schnizlein, Henning Schulzrinne, Barbara Stark, Hannes Tschofenig, Martin Thomson, and James Winterbottom.

7. References

7.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

7.2. Informative References

[DHCP-LOC-URI] Polk, J., "Dynamic Host Configuration Protocol (DHCP) IPv4 and IPv6 Option for a Location Uniform Resource Identifier (URI)", Work in Progress, March 2010.

[HELD] Barnes, M., Winterbottom, J., Thomson, M., and B. Stark, "HTTP Enabled Location Delivery (HELD)", Work in Progress, August 2009.

[LLDP-MED] Telecommunications Industry Association (TIA), "ANSI/TIA-1057 Link Layer Discovery Protocol - Media Endpoint Discovery", 2006.

[LOC-FILTERS] Mahy, R., Rosen, B., and H. Tschofenig, "Filtering Location Notifications in the Session Initiation Protocol (SIP)", Work in Progress, March 2010.

[LOC-CONVEY] Polk, J. and B. Rosen, "Location Conveyance for the Session Initiation Protocol", Work in Progress, February 2010.

[RFC3693] Cuellar, J., Morris, J., Mulligan, D., Peterson, J., and J. Polk, "Geopriv Requirements", RFC 3693, February 2004.

[RFC4119] Peterson, J., "A Presence-based GEOPRIV Location Object Format", RFC 4119, December 2005.

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

Author's Address

Roger Marshall (editor)
TeleCommunication Systems, Inc.
2401 Elliott Avenue
2nd Floor
Seattle, WA 98121
US

Phone: +1 206 792 2424
EMail: rmarshall@telecomsys.com
URI: http://www.telecomsys.com