

GEOPRIV
Internet-Draft
Intended status: Standards Track
Expires: September 5, 2007

H. Schulzrinne
Columbia U.
March 4, 2007

**RELO: Retrieving End System Location Information
draft-schulzrinne-geopriv-relo-03.txt**

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

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

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

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

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

This Internet-Draft will expire on September 5, 2007.

Copyright Notice

Copyright (C) The IETF Trust (2007).

Abstract

In some network configurations, it is desirable for the end system to be able to obtain its geodetic or civic location using an application-layer protocol. This document describes RELO (Retrieving End system LOcation), a simple, HTTP-based stateless protocol profile that fulfills this need.

Table of Contents

- [1. Introduction](#) [3](#)
- [2. Terminology](#) [3](#)
- [3. Protocol Description](#) [3](#)
 - [3.1. Discovery](#) [4](#)
 - [3.2. Query](#) [4](#)
 - [3.3. Response](#) [7](#)
 - [3.4. Signed Location](#) [8](#)
 - [3.5. Error Reporting](#) [8](#)
 - [3.6. Client Authentication](#) [8](#)
- [4. IANA Considerations](#) [8](#)
 - [4.1. S-NAPTR Application Service Tag](#) [8](#)
 - [4.2. HTTP Message Header 'Subscribe'](#) [9](#)
 - [4.3. MIME Type](#) [9](#)
- [5. Security Considerations](#) [10](#)
- [6. Acknowledgments](#) [11](#)
- [7. References](#) [11](#)
 - [7.1. Normative References](#) [11](#)
 - [7.2. Informative References](#) [12](#)
- [Author's Address](#) [12](#)
- [Intellectual Property and Copyright Statements](#) [13](#)

1. Introduction

The RELO HTTP protocol usage allows end systems (devices) to obtain information about their current geodetic (longitude, latitude) or civic (jurisdictional or postal street address) location, based on their Internet Protocol address or possibly other identifiers. The protocol uses HTTP [3] to retrieve the information. The location information can be returned by value or by reference, either for retrieval or for event notification by subscription.

The protocol is motivated by the requirement that end user network-layer equipment, such as DSL modems, routers, NATs and wireless access points, cannot be modified. Hence, a DHCP or PPP based solution cannot be reused. A more detailed problem statement is provided in [11]. To reduce privacy risks, RELO is designed for "first-party" retrieval, i.e., the device obtains its own location or a reference thereto. It is not designed for a third party to retrieve location information about a device. However, RELO may retrieve a reference to location information that can be passed to third parties.

Like other HTTP-based protocols, RELO may fail to deliver the correct location information in some circumstances unless special care is taken. For example, if the ISP only allows HTTP connections that traverse an HTTP proxy, the LIS would return the location of the proxy, not that of the client. In this case, however, the ISP would likely know about the proxy and make appropriate arrangements, e.g., to allow non-proxied connections to the LIS only.

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 [1].

This document reuses terminology introduced by RFC 3693 [5] and [11].

3. Protocol Description

This section describes the Location Information Server (LIS) discovery procedure (see Section 3.1), the query message (see Section 3.2) and the response message (see Section 3.3).

[3.1.](#) Discovery

The URI for the location server is conveyed via DHCP (not described here) or DNS (S-NAPTR) [7]. The domain is determined from the domain name of the end host, typically conveyed as part of the configuration information. In the example below, host dhcp-17.example.com would query the S-NAPTR record for that domain, obtaining the location server name relo.example.com.

```
dhcp-17.example.com.  
;      order pref flags service      regexp  
IN NAPTR 50 50 "a" "Location.relo"  ""  
; replacement  
      relo.example.com
```

If the host does not have a domain name or there is no suitable S-NAPTR record, the host checks whether the PTR record for the IP address exists and uses that domain, e.g., a host with the address 192.168.1.2 would query for the S-NAPTR record of 2.1.168.192.in-addr.arpa.

[3.2.](#) Query

The query is transmitted to the server in an HTTP GET request. The use of TLS [10] is RECOMMENDED. To simplify implementations, the protocol currently transmits all parameters as HTTP query parameters. As always, the order of parameters is immaterial. (Since the query does not change the state of the resource, GET is the appropriate method.)

Unless other identifiers are provided, the end system is identified by its IP address, contained in the IP packets carrying the HTTP request. If the querier is behind a NAT or firewall, the server will see the querier's public IP address and use that address to identify the end system. In those cases, the location of the network termination equipment, such as the DSL modem or 802.11 access point, will be returned, not the actual location of the querier since the LIS generally has no way to estimate that location. Other network identifiers, such as those provided by CDP, LLDP or the MAC address, can be provided; the client SHOULD include all such identifiers it knows about. The server is free to choose the most appropriate identifier to determine the client location information and SHOULD choose the one yielding the highest accuracy and reliability within the time limits provided by the 'within' parameter. If any of the network identifiers or other parameters have the wrong syntax, the server returns a 400 (Bad Request) error, with additional information on the syntax error provided in the entity body and the HTTP Reason-

Phrase.

by The 'by' parameter indicates whether the client would prefer to obtain a value ('value') or a reference ('reference'). The default is 'value' if the LIS supports it, 'reference' otherwise. The client can restrict the type of location information returned via the HTTP Accept header in the request. If the server can only deliver a format not listed, it responds with a 406 (Not Acceptable) status code.

within The 'within' parameter indicates the amount of time that the client is willing to wait for an answer, expressed as a positive decimal integer and measured in seconds, using the canonical representation of the XML 'decimal' primitive data type. If omitted, the LIS SHOULD return the most precise location information available.

type The 'type' parameter indicates whether the client desires a 'civic' or 'geo' address. The default is 'geo' if supported by the server and 'civic' otherwise. If a client requests one type of location information, but the server only has the other, the server MAY return that information instead, as the client can easily determine that this is the case. Alternatively, the LIS MAY return a 404 (Not Found) error, with an appropriate explanation. A client willing to accept both formats can either omit the 'type' parameter if it wants to only receive one type, or query for both types, even if one returns an error.

retransmission-allowed The client uses the 'retransmission-allowed' parameter to request that the PIDF location object contains the corresponding parameter value. Only the string literals 'yes' and 'no' are allowed. The default is 'no'.

retention-expiry The client uses the 'retention-expiry' parameter to request that the PIDF-LO contains the corresponding usage rule. The value is an XML date time, as specified by PIDF-LO. If omitted, the defaults specified for PIDF-LO are used.

external-ruleset The client uses the 'external-ruleset' parameter to request that the PIDF-LO contains the corresponding usage rule. The value is of XML type anyURI, as specified by PIDF-LO.

note-well The client uses the 'note-well' parameter to request that the PIDF-LO contains the corresponding usage rule. The rule is of type XML string.

note-well-lang The client uses the 'note-well-lang' parameter to request that the PIDF-LO 'note-well' element contains the corresponding language indication, using XML conventions.

url The 'url' parameter is used only if a location reference URL is being renewed. It is ignored if the 'by=value' parameter is specified. The expiration time of the URL is updated, assuming that the secret agrees with that stored for the URL. If the parameter is not supplied, a new URL is created.

- expires** The 'expires' parameter contains an XML dateTime string in canonical (UTC) representation. It indicates the time that the requestor would like the location reference or value to expire. For values, the parameter sets the 'retention-expiry' data in PIDF-LO. An expiration date in the past immediately invalidates the URL. By default, the URL expires two hours after being issued.
- secret** The 'secret' parameter allows the client to provide a password that controls access to the URL. When creating a new URL, the server stores that password with the URL for later modification. If not specified upon creation, the URL properties cannot be modified later.
- mac** The 'mac' parameter contains an IEEE IEEE MAC address written in IEEE EUI-64 or EUI-48 notation, with lower-case hexadecimal characters separated by colons. An example is "0:3:fc:0:ca:27". This is a network identifier.
- cdp** The 'cdp' parameter contains a Cisco Discovery Protocol (CDP). The CDP identifier consists of the CDP device id, a colon and the port ID. An example is cepsr-7-1:FastEthernet6/6. This is network identifier.
- msap** The 'msap' parameter identifies a MAC service access point, typically a switch chassis and port. If derived from LLDP (IEEE 802.1ab), it is encoded in base64. This is a network identifier.
- assert** The 'assert' parameter contains a PIDF-LO, e.g., derived via GPS, that the client would like the LIS to sign and store. Depending on the RELO parameters supplied, the server will return either a location reference or a, typically signed, location object. A server MAY return a 403 (Forbidden) response if the LIS does not want to allow this particular client to assert location information. If the assertion is granted, future requests for location for the combination of network identifiers (mac, msap, cdp, etc.) MAY return this location information, but a LIS MAY decide to only allow retrieval from the same IP address used for the assertion.

Thus, a URL without a query string returns the current location value, with a retention period of two hours, based on the client's IP address. If several addresses are provided, it is left to the server to select the one that it has location information for. Due to the use of return routability, the use of the IP address is preferred.

A query example is shown below:

```
http://example.com?type=civic&by=value&secret=bond007
&expires=2007%2D01%2D20T23%3A10%3A01%0D%0A
```

Query URL for location object containing civic

location information

Schulzrinne

Expires September 5, 2007

[Page 6]

This protocol does not provide the ability for the end host to transmit a location estimate as, for example, obtained from a local GPS receiver, to the LIS.

3.3. Response

If the client indicated a preference for location-by-reference, the answer simply contains a URI-list, i.e., media type text/uri-list [2].

For location-by-value, RELO currently returns a PIDF-LO [8] document. (Future extensions of RELO may support other location object formats.)

For PIDF-LO, the entity attribute is pres:anonymous@anonymous.invalid. The <retransmission-allowed> element in the <usage-rules; element is set to 'no'; the <retention-expiry> element is set to the 'expires' attribute in the query or its default value (see above).

Normal HTTP status responses are used to indicate failure conditions, e.g., when the information is unavailable.

The server indicates the validity period of the information using the HTTP Expires header field. If a reference is returned, the reference URL itself is not guaranteed to be valid beyond the expiration time.

The server MAY provide one or more URLs in a new HTTP header field, `Subscribe`, that the client can subscribe to if it wants to receive updates for the object retrieved via HTTP. At least one of the URLs MUST be a SIP URL. For SIP, the event name to be used in the subscription can be encoded in the URL. (An HTTP header field was chosen since the subscription mechanism does not depend on the media type and is equally applicable to other media type. Putting the subscription URL in an HTTP header allows to subscribe to media types where it is difficult to embed SIP URLs, such as a JPEG image.) The server makes no guarantees that the client has the appropriate credentials to subscribe to the object. Clients MAY support this mechanism; all clients that do support subscriptions MUST support the SIP SUBSCRIBE and NOTIFY methods.

The field value consists of one or more absolute URIs:

```
Subscribe = "Subscribe" ":" 1#absoluteURI
```

An example is:

```
Subscribe: sip:data@example.com?Event=location
```


[TBD: Since this mechanism is not limited to location delivery, this might be better separated into a stand-alone draft.]

The response containing the location information is not signed. A response containing a randomized HTTP URL is shown below.

`http://example.com/15555551002adfkafjyonqoijoyukjglky`

Response containing location-by-reference

3.4. Signed Location

RELO uses XML DSIG for digitally signing location objects, as described in [\[12\]](#).

3.5. Error Reporting

RELO uses HTTP status codes in case of errors. In addition to the status code, the response SHOULD contain an entity body explaining the error, in a format corresponding to the Accept header field. For example, a device with a text-only display may allow only textual, rather than HTML, error explanation by listing text/plain in addition to the URI list and location object formats it supports. In addition, the HTTP Reason-Phrase SHOULD identify the error cause, rather than use the generic HTTP response message.

(RELO does not define a range of protocol-specific error conditions since it appears highly unlikely that a client would be able to act on this structured information. The reason phrase and textual information are more likely to be useful to users and for client debugging, as they can represent many more error conditions.)

3.6. Client Authentication

For first-party requests using the IP address as a query parameter, authentication is OPTIONAL, but it is REQUIRED for third-party requests. Where necessary, RELO relies on HTTP authentication mechanisms, such as Digest authentication or TLS client certificates.

4. IANA Considerations

4.1. S-NAPTR Application Service Tag

This document registers the label "RELO" as the S-NAPTR application service tag according to [\[7\]](#) for location lookup services and defines the intended usage, interoperability considerations and security

considerations ([Section 5](#)).

4.2. HTTP Message Header 'Subscribe'

This document requests the registration of a new message header field, 'Subscribe', according to [RFC 3864](#) [6].

Header field name: Subscribe

4.3. MIME Type

This specification also requests the registration of a new MIME type according to the procedures of [RFC 4288](#) [9] and guidelines in [RFC 3023](#) [4].

MIME media type name: application

MIME subtype name: relo+xml

Mandatory parameters: none

Optional parameters: charset

Indicates the character encoding of enclosed XML.

Encoding considerations:

Uses XML, which can employ 8-bit characters, depending on the character encoding used. See [RFC 3023](#) [4], Section 3.2.

Security considerations:

This content type is designed to carry authorization policies. Appropriate precautions should be adopted to limit disclosure of this information. Please refer to [Section 5](#) of RFCXXXX [NOTE TO IANA/RFC-EDITOR: Please replace XXXX with the RFC number of this specification.] and to the security considerations described in [Section 10 of RFC 3023](#) [4] for more information.

Interoperability considerations: None

Published specification: RFCXXXX [NOTE TO IANA/RFC-EDITOR: Please replace XXXX with the RFC number of this specification.] this document

Applications which use this media type:

Presence- and location-based systems

Additional information:

Magic Number: None

File Extension: .reloxml

Macintosh file type code: 'TEXT'

Personal and email address for further information: Henning Schulzrinne, hgs@cs.columbia.edu

Intended usage: LIMITED USE

Author/Change controller:

This specification is a work item of the IETF GEOPRIV working group, with mailing list address <geopriv@ietf.org>.

5. Security Considerations

If IP addresses are used as identifiers, RELO relies on return routability to ensure that only the current owner of an IP address can obtain location information for that host, and assumes that an attacker cannot generate and intercept packets for a spoofed IP address. Note that TLS itself does not prevent client address spoofing if the attacker can intercept and generate IP packets with the victim's IP address.

The victim can be protected against this privacy breach if the client and LIS share a secret, such as a username/password combination, and the LIS can associate an IP address with a particular user, e.g., based on PPP authentication. In that case, HTTP digest authentication can be used to prevent a third party from using a spoofed IP address to fraudulently obtain location information. Unfortunately, such authentication information is not generally available to wireless nodes in residential networks, for example.

To prevent others from accessing location information for a particular host, the reference to a Location Object MUST NOT be guessable. For example, it may contain a random component. It is RECOMMENDED to use TLS with confidentiality protection to prevent eavesdroppers to observe the protocol exchange between the end host and the LIS.

Other identifiers may have different privacy concerns. For example, switch port identifiers, such as those returned by CDP or LLDP, may not pose as grave a risk of disclosing private information by themselves unless they can be linked to an IP address. Thus, in this case, privacy-protecting the RELO query is particularly important. However, no special authorization is needed unless the ability to enumerate the locations of LAN jacks is considered sensitive.

Signing of location information is beyond the scope of this document. Thus, colluding attackers may be able to obtain and replay location information that does not correspond to their true location.

6. Acknowledgments

This document is based on discussions with Hannes Tschofenig and inspired by protocols such as HELD. Jong Yul Kim, Rohan Mahy, Andrew Newton, and Wonsang Song provided helpful input.

7. References

7.1. Normative References

- [1] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [2] Mealling, M. and R. Daniel, "URI Resolution Services Necessary for URN Resolution", [RFC 2483](#), January 1999.
- [3] Fielding, R., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., and T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1", [RFC 2616](#), June 1999.
- [4] Murata, M., St. Laurent, S., and D. Kohn, "XML Media Types", [RFC 3023](#), January 2001.
- [5] Cuellar, J., Morris, J., Mulligan, D., Peterson, J., and J. Polk, "Geopriv Requirements", [RFC 3693](#), February 2004.
- [6] Klyne, G., Nottingham, M., and J. Mogul, "Registration Procedures for Message Header Fields", [BCP 90](#), [RFC 3864](#), September 2004.
- [7] Daigle, L. and A. Newton, "Domain-Based Application Service Location Using SRV RRs and the Dynamic Delegation Discovery Service (DDDS)", [RFC 3958](#), January 2005.

- [8] Peterson, J., "A Presence-based GEOPRIV Location Object Format", [RFC 4119](#), December 2005.
- [9] Freed, N. and J. Klensin, "Media Type Specifications and Registration Procedures", [BCP 13](#), [RFC 4288](#), December 2005.
- [10] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.1", [RFC 4346](#), April 2006.

[7.2.](#) Informative References

- [11] Tschofenig, H. and H. Schulzrinne, "GEOPRIV Layer 7 Location Configuration Protocol; Problem Statement and Requirements", [draft-tschofenig-geopriv-l7-lcp-ps-03](#) (work in progress), October 2006.
- [12] Thomson, M. and J. Winterbottom, "Digital Signature Methods for Location Dependability", [draft-thomson-geopriv-location-dependability-00](#) (work in progress), February 2007.

Author's Address

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

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

Full Copyright Statement

Copyright (C) The IETF Trust (2007).

This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgment

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).

