

Internet Engineering Task Force (IETF)
Request for Comments: 5687
Category: Informational
ISSN: 2070-1721

H. Tschofenig
Nokia Siemens Networks
H. Schulzrinne
Columbia University
March 2010

GEOPRIV Layer 7 Location Configuration Protocol:
Problem Statement and Requirements

Abstract

This document provides a problem statement, lists requirements, and captures design aspects for a GEOPRIV Layer 7 (L7) Location Configuration Protocol (LCP). This protocol aims to allow an end host to obtain location information, by value or by reference, from a Location Information Server (LIS) that is located in the access network. The obtained location information can then be used for a variety of different protocols and purposes. For example, it can be used as input to the Location-to-Service Translation (LoST) Protocol or to convey location within the Session Initiation Protocol (SIP) to other entities.

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/rfc5687>.

Copyright Notice

Copyright (c) 2010 IETF Trust and the persons identified as the document authors. All rights reserved.

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

Table of Contents

1. Introduction	3
2. Terminology	3
3. Scenarios	4
3.1. Fixed-Wired Environment	4
3.2. Mobile Network	7
3.3. Wireless Access	8
4. Discovery of the Location Information Server	9
5. Identifier for Location Determination	11
6. Requirements	14
7. Security Considerations	16
8. Contributors	17
9. Acknowledgements	18
10. References	18
10.1. Normative References	18
10.2. Informative References	18

1. Introduction

This document provides a problem statement, lists requirements, and captures design aspects for a GEOPRIV Layer 7 (L7) Location Configuration Protocol (LCP). The protocol has two purposes:

- o It is used by a device to obtain its own location (referred as "Location by Value" or LbyV) from a dedicated node, called the Location Information Server (LIS).
- o It enables the device to obtain a reference to location information (referred as "Location by Reference" or LbyR). This reference can take the form of a subscription URI, such as a SIP presence-based Uniform Resource Identifier (URI), an HTTP/HTTPS URI, or another URI. The requirements related to the task of obtaining an LbyR are described in a separate document, see [LBYR-REQS].

The need for these two functions can be derived from the scenarios presented in Section 3.

For this document, we assume that the GEOPRIV Layer 7 LCP runs between the device and the LIS.

This document is structured as follows. Section 4 discusses the challenge of discovering the LIS in the access network. Section 5 compares different types of identifiers that can be used to retrieve location information. A list of requirements for the L7 LCP can be found in Section 6.

This document does not describe how the access network provider determines the location of the device since this is largely a matter of the capabilities of specific link-layer technologies or certain deployment environments.

2. Terminology

In this document, the key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" are to be interpreted as described in RFC 2119 [RFC2119], with the qualification that unless otherwise stated these words apply to the design of the GEOPRIV Layer 7 Location Configuration Protocol.

The term Location Information Server (LIS) refers to an entity capable of determining the location of a device and of providing that location information, a reference to it, or both via the Location Configuration Protocol (LCP) to the Target.

This document also uses terminology from [RFC5012] (such as Internet Access Provider (IAP), Internet Service Provider (ISP), and Application Service Provider (ASP)).

With the term "Access Network Provider" we refer to the IAP and the ISP) without further distinguishing these two entities, as it is not relevant for the purpose of this document. An additional requirements document on LIS-to-LIS protocol [LIS2LIS] shows a scenario where the separation between IAP and ISP is important.

3. Scenarios

This section describes a few network scenarios where the L7 LCP may be used. Note that this section does not aim to exhaustively list all possible deployment environments. Instead, we focus on the following environments:

- o DSL/Cable networks, WiMAX-like (Worldwide Interoperability for Microwave Access) fixed access
- o Airport, city, campus wireless networks, such as 802.11a/b/g, 802.16e/WiMAX
- o 3G networks
- o Enterprise networks

Note that we use the term 'host' instead of device for better readability.

3.1. Fixed-Wired Environment

Figure 1 shows a Digital Subscriber Line (DSL) network scenario with the Access Network Provider and the customer premises. The Access Network Provider operates link- and network-layer devices (represented as a node) and the LIS.

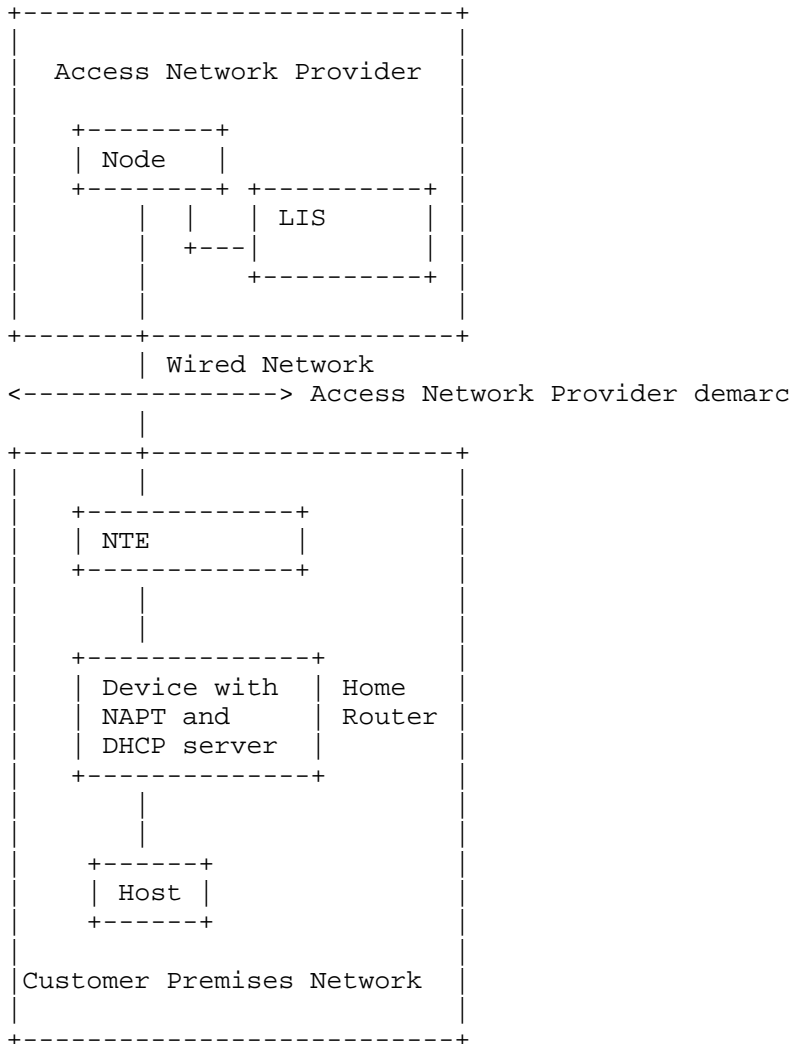


Figure 1: Fixed-Wired Scenario

The customer premises network consists of a router with a Network Address Translator with Port Address Translation (NAPT) and a DHCP server as used in most Customer Premises Networks (CPNs) and the Network Termination Equipment (NTE) where Layer 1 and sometimes Layer 2 protocols are terminated. The router in the home network (e.g., broadband router, cable or DSL router) typically runs a NAPT and a DHCP server. The NTE is a legacy device and in many cases cannot be modified for the purpose of delivering location information to the host. The same is true of the device with the NAPT and DHCP server.

It is possible for the NTE and the home router to physically be in the same box, or for there to be no home router, or for the NTE and host to be in the same physical box (with no home router). An example of this last case is where Ethernet service is delivered to customers' homes, and the Ethernet network interface card (NIC) in their PC serves as the NTE.

Current CPN deployments generally fall into one of the following classifications:

1. Single PC

1. with Ethernet network interface card (NIC), with Point-to-Point Protocol Over Ethernet (PPPoE), or Dynamic Host Configuration Protocol (DHCP) on PC; there may be a bridged DSL or cable modem as the NTE, or the Ethernet NIC might be the NTE.
2. with USB-based DSL access or a cable modem access using Point-to-Point Protocol over ATM (PPPoA), PPPoE, or DHCP on PC.

Note that the device with NAPT and DHCP of Figure 1 is not present in such a scenario.

2. One or more hosts with at least one router (DHCP client or PPPoE, DHCP server in router; Voice over IP (VoIP) can be a soft client on a PC, a stand-alone VoIP device, or an Analog Terminal Adaptor (ATA) function embedded in a router):
 1. combined router and NTE.
 2. separate router with NTE in bridged mode.
 3. separate router with NTE (NTE/router does PPPoE or DHCP to WAN, router provides DHCP server for hosts in LAN; double NAT).

The majority of fixed-access broadband customers use a router. The placement of the VoIP client is mentioned to describe what sorts of hosts may need to be able to request location information. Soft clients on PCs are frequently not launched until long after bootstrapping is complete, and are not able to control any options that may be specified during bootstrapping. They also cannot control whether a VPN client is running on the end host.

3.2. Mobile Network

One example of a moving network is a WiMAX-fixed wireless scenario. This also applies to "pre-WiMAX" and "WiMAX-like" fixed wireless networks. In implementations intended to provide broadband service to a home or other stationary location, the customer-side antenna/NTE tends to be rather small and portable. The LAN-side output of this device is an Ethernet jack, which can be used to feed a PC or a router. The PC or router then uses DHCP or PPPoE to connect to the access network, the same as for wired access networks. Access providers who deploy this technology may use the same core network (including network elements that terminate PPPoE and provide IP addresses) for DSL, fiber to the premises (FTTP), and fixed wireless customers.

Given that the customer antenna is portable and can be battery-powered, it is possible for a user to connect a laptop to it and move within the coverage area of a single base antenna. This coverage area can be many square kilometers in size. In this case, the laptop (and any SIP client running on it) would be completely unaware of their mobility. Only the user and the network are aware of the laptop's mobility.

Further examples of moving networks (where end devices may not be aware that they are moving) can be found in busses, trains, and airplanes.

Figure 2 shows an example topology for a moving network.

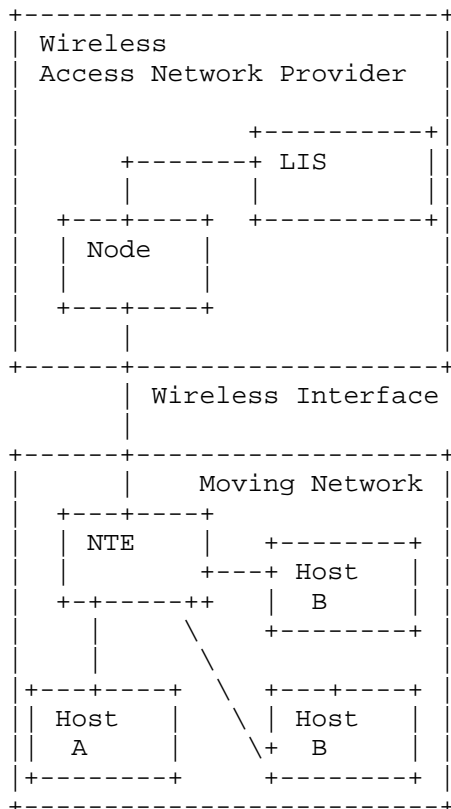


Figure 2: Moving Network

3.3. Wireless Access

Figure 3 shows a wireless access network where a moving host obtains location information or references to location information from the LIS. The access equipment uses, in many cases, link-layer devices. Figure 3 represents a hotspot network found, for example, in hotels, airports, and coffee shops. For editorial reasons we only describe a single access point and do not depict how the LIS obtains location information since this is very deployment specific.




Figure 3: Wireless Access Scenario

4. Discovery of the Location Information Server

Note that this section lists mechanisms that were discussed in the GEOPRIV Layer 7 Location Configuration Protocol design team. They are included to show challenges in the problem space and are listed for completeness reasons. They do not in any way mean that there is consensus about any of the mechanisms or that the IETF recommends any of the procedures described in this section.

When a device wants to retrieve location information from the LIS, it first needs to discover it. Based on the problem statement of determining the location of the device, which is known best by entities close to the device itself, we assume that the LIS is located in the local subnet or in the access network. Several procedures have been investigated that aim to discover the LIS in such an access network.

DHCP-based Discovery:

In some environments, the Dynamic Host Configuration Protocol (DHCP) might be a good choice for discovering the fully-qualified domain name (FQDN) or the IP address of the LIS. In environments where DHCP can be used, it is also possible to use the already defined location extensions. In environments with legacy devices, such as the one shown in Section 3.1, a DHCP-based discovery solution may not be possible.

DNS-based Discovery:

Before a Domain Name System (DNS) lookup can be started, it is necessary to learn the domain name of the access network that runs an LIS. Several ways to learn the domain name exist. For example, the end host obtains its own public IP address via Simple Traversal of the UDP Protocol through NAT (STUN) [RFC5389], and performs a reverse DNS lookup (assuming the data is provisioned into the DNS). Then, the DNS Service (SRV) record or the DNS Naming Authority Pointer (NAPTR) record for that domain is retrieved. A more detailed description of this approach can be found in [LIS-DISC].

Redirect Rule:

A redirect rule at an entity in the access network could be used to redirect the L7 LCP signaling messages (destined to a specific port) to the LIS. The device could then discover the LIS by sending a packet with a specific (registered) port number to almost any address as long as the destination IP address does not target an entity in the local network. The packet would be redirected to the respective LIS being configured. The same procedure is used by captive portals whereby any HTTP traffic is intercepted and redirected.

To some extent, this approach is similar to packets that are marked with a Router Alert option [RFC2113] and intercepted by entities that understand the specific marking. In the above-mentioned case, however, the marking is provided via a registered port number instead of relying on a Router Alert option.

This solution approach would require a deep packet inspection capability at an entity in the access provider's networks that scans for the occurrence of particular destination port numbers.

Multicast Query:

A device could also discover an LIS by sending a DNS query to a well-known address. An example of such a mechanism is multicast DNS (see [RFC4795] and [mDNS]). Unfortunately, these mechanisms only work on the local link.

Anycast:

With this solution, an anycast address is defined (for IPv4 and IPv6) in the style of [RFC3068] that allows the device to route discovery packets to the nearest LIS. Note that this procedure would be used purely for discovery and is therefore similar to the local Teredo server discovery approach outlined in Section 4.2 of [TEREDO-SEL].

The LIS discovery procedure raises deployment and security issues. The access network needs to be designed to prevent man-in-the-middle adversaries from presenting themselves as an LIS to devices. When a device discovers an LIS, it needs to ensure (and be able to ensure) that the discovered entity is indeed an authorized LIS.

5. Identifier for Location Determination

Note that this section lists mechanisms that were discussed in the GEOPRIV Layer 7 Location Configuration Protocol design team. They are included to show challenges in the problem space and are listed for completeness reasons. They do not in any way mean that there is consensus about any of the mechanisms or that the IETF recommends any of the procedures described in this section.

The LIS returns location information to the device when it receives a request. Some form of identifier is therefore needed to allow the LIS to retrieve the device's current location, or a good approximation, from a database.

The chosen identifier needs to have the following properties:

Ability for Device to learn or know the identifier:

The device MUST know or MUST be able to learn of the identifier (explicitly or implicitly) in order to send it to the LIS. Implicitly refers to the situation where a device along the path between the device and the LIS modifies the identifier, as it is done by a NAT when an IP address based identifier is used.

Ability to use the identifier for location determination:

The LIS MUST be able to use the identifier (directly or indirectly) for location determination. Indirectly refers to the case where the LIS uses other identifiers internally for location determination, in addition to the one provided by the device.

Security properties of the identifier:

Misuse needs to be minimized whereby an off-path adversary **MUST NOT** be able to obtain location information of other devices. An on-path adversary in the same subnet **SHOULD NOT** be able to spoof the identifier of another device in the same subnet.

The following list discusses frequently mentioned identifiers and their properties:

Media Access Control (MAC) Address:

The MAC address is known to the device itself, but not carried beyond a single IP hop and therefore not accessible to the LIS in most deployment environments (unless carried in the L7 LCP itself).

Asynchronous Transfer Mode (ATM) Virtual Path Identifier / Virtual Circuit Identifier (VPI/VCI):

The VCI/VPI is generally only seen by the DSL modem. Almost all routers in the United States use 1 of 2 VPI/VCI value pairs: 0/35 and 8/35. This VC is terminated at the digital subscriber line access multiplexer (DSLAM), which uses a different VPI/VCI (per end customer) to connect to the ATM switch. Only the network provider is able to map VPI/VCI values through its network. With the arrival of Very high rate Digital Subscriber Line (VDSL), ATM will slowly be phased out in favor of Ethernet.

Ethernet Switch (Bridge)/Port Number:

This identifier is available only in certain networks, such as enterprise networks, typically available via the IEEE 802.1AB protocol [802.1AB] or proprietary protocols like the Cisco Discovery Protocol (CDP) [CDP].

Cell ID:

This identifier is available in cellular data networks and the cell ID may not be visible to the device.

Host Identifier:

The Host Identifier introduced by the Host Identity Protocol (HIP) [RFC5201] allows identification of a particular host. Unfortunately, the network can only use this identifier for location determination if the operator already stores a mapping of host identities to location information. Furthermore, there is a deployment problem since the host identities are not used in today's networks.

Cryptographically Generated Address (CGA):

The concept of a Cryptographically Generated Address (CGA) was introduced by [RFC3972]. The basic idea is to put the truncated hash of a public key into the interface identifier part of an IPv6 address. In addition to the properties of an IP address, it allows a proof of ownership. Hence, a return routability check can be omitted. It is only available for IPv6 addresses.

Network Access Identifiers:

A Network Access Identifier [RFC4282] is used during the network access authentication procedure, for example, in RADIUS [RFC2865] and Diameter [RFC3588]. In DSL networks, the user credentials are, in many cases, only known by the home router and not configured at the device itself. To the network, the authenticated user identity is only available if a network access authentication procedure is executed. In case of roaming, the user's identity might not be available to the access network since security protocols might offer user identity confidentiality and thereby hide the real identity of the user allowing the access network to only see a pseudonym or a randomized string.

Unique Client Identifier

The Broadband Forum has defined that all devices that expect to be managed by the TR-069 interface, see [TR069], have to be able to generate an identifier that uniquely identifies the device. It also has a requirement that routers that use DHCP to the WAN use RFC 4361 [RFC4361] to provide the DHCP server with a unique client identifier. This identifier is, however, not visible to the device when legacy NTE devices are used.

IP Address:

The device's IP address may be used for location determination. This IP address is not visible to the LIS if the device is behind one or multiple NATs. This may not be a problem since the location of a device that is located behind a NAT cannot be determined by the access network. The LIS would in this case only see the public IP address of the NAT binding allocated by the NAT, which is the expected behavior. The property of the IP address for a return routability check is attractive to return location information only to the address that submitted the request. If an adversary wants to learn the location of a device (as identified by a particular IP address), then it does not see the response message (unless it is on the subnetwork or at a router along the path towards the LIS).

On a shared medium, an adversary could ask for location information of another device. The adversary would be able to see the response message since it is sniffing on the shared medium unless security mechanisms, such as link-layer encryption, are in place. With a network deployment as shown in Section 3.1 with multiple devices in the Customer Premises being behind a NAT, the LIS is unable to differentiate the individual devices. For WLAN deployments as found in hotels, as shown in Section 3.3, it is possible for an adversary to eavesdrop data traffic and subsequently to spoof the IP address in a query to the LIS to learn more detailed location information (e.g., specific room numbers). Such an attack might, for example, compromise the privacy of hotel guests.

6. Requirements

The following requirements and assumptions have been identified:

Requirement L7-1: Identifier Choice

The L7 LCP MUST be able to carry different identifiers or MUST define an identifier that is mandatory to implement. Regarding the latter aspect, such an identifier is only appropriate if it is from the same realm as the one for which the location information service maintains identifier-to-location mapping.

Requirement L7-2: Mobility Support

The L7 LCP MUST support a broad range of mobility from devices that can only move between reboots, to devices that can change attachment points with the impact that their IP address is changed, to devices that do not change their IP address while roaming, to devices that continuously move by being attached to the same network attachment point.

Requirement L7-3: ASP and Access Network Provider Relationship

The design of the L7 LCP MUST NOT assume that a business or trust relationship between the Application Service Provider (ASP) and the Access Network Provider. Requirements for resolving a reference to location information are not discussed in this document.

Requirement L7-4: Layer 2 and Layer 3 Provider Relationship

The design of the L7 LCP MUST assume that there is a trust and business relationship between the L2 and the L3 provider. The L3 provider operates the LIS that the device queries. It, in turn, needs to obtain location information from the L2 provider since this one is closest to the device. If the L2 and L3 provider for the same device are different entities, they cooperate for the purposes needed to determine locations.

Requirement L7-5: Legacy Device Considerations

The design of the L7 LCP MUST consider legacy devices, such as residential NAT devices and NTEs in a DSL environment, that cannot be upgraded to support additional protocols, for example, to pass additional information towards the device.

Requirement L7-6: Virtual Private Network (VPN) Awareness

The design of the L7 LCP MUST assume that at least one end of a VPN is aware of the VPN functionality. In an enterprise scenario, the enterprise side will provide the LIS used by the device and can thereby detect whether the LIS request was initiated through a VPN tunnel.

Requirement L7-7: Network Access Authentication

The design of the L7 LCP MUST NOT assume that prior network access authentication.

Requirement L7-8: Network Topology Unawareness

The design of the L7 LCP MUST NOT assume that devices are aware of the access network topology. Devices are, however, able to determine their public IP address(es) via mechanisms, such as Simple Traversal of User Datagram Protocol (UDP) Through Network Address Translators (NATs) (STUN) [RFC5389] or Next Steps in Signaling (NSIS) NAT/Firewall NSIS Signaling Layer Protocol (NSLP) [NSLP].

Requirement L7-9: Discovery Mechanism

The L7 LCP MUST define a mandatory-to-implement LIS discovery mechanism.

Requirement L7-10: PIDF-LO Creation

When an LIS creates a Presence Information Data Format (PIDF) Location Object (LO) [RFC4119], then it MUST put the <geopriv> element into the <device> element of the presence document (see [RFC4479]). This ensures that the resulting PIDF-LO document, which is subsequently distributed to other entities, conforms to the rules outlined in [RFC5491].

7. Security Considerations

By using a Geolocation L7 Location Configuration Protocol, the device (and a human user of such a device, if applicable) exposes themselves to a privacy risk whereby an unauthorized entity receives location information. Providing confidentiality protected location to the requestor depends on the success of four steps:

1. The client MUST have a means to discover a LIS.
2. The client MUST authenticate the discovered LIS.
3. The LIS MUST be able to determine location and return it to the authorized entity.
4. The LIS MUST securely exchange messages without intermediaries eavesdropping or tampering with them.

This document contains various security-related requirements throughout the document addressing the above-mentioned steps. For a broader security discussion of the overall geolocation privacy architecture, the reader is referred to [GEOPRIV-ARCH].

8. Contributors

This contribution is a joint effort of the GEOPRIV Layer 7 Location Configuration Requirements Design Team of the IETF GEOPRIV Working Group. The contributors include Henning Schulzrinne, Barbara Stark, Marc Linsner, Andrew Newton, James Winterbottom, Martin Thomson, Rohan Mahy, Brian Rosen, Jon Peterson, and Hannes Tschofenig.

We would like to thank the GEOPRIV Working Group Chairs, Andy Newton, Randy Gellens, and Allison Mankin, for creating the design team. Furthermore, we would like thank Andy Newton for his support during the design team mailing list, for setting up Jabber chat conferences, and for participating in the phone conference discussions.

The design team members can be reached at:

Marc Linsner: mlinsner@cisco.com

Rohan Mahy: rohan@ekabal.com

Andrew Newton: andy@hxr.us

Jon Peterson: jon.peterson@neustar.biz

Brian Rosen: br@brianrosen.net

Henning Schulzrinne: hgs@cs.columbia.edu

Barbara Stark: Barbara.Stark@bellsouth.com

Martin Thomson: Martin.Thomson@andrew.com

Hannes Tschofenig: Hannes.Tschofenig@nsn.com

James Winterbottom: James.Winterbottom@andrew.com

9. Acknowledgements

We would also like to thank Murugaraj Shanmugam, Ted Hardie, Martin Dawson, Richard Barnes, James Winterbottom, Tom Taylor, Otmar Lendl, Marc Linsner, Brian Rosen, Roger Marshall, Guy Caron, Doug Stuard, Eric Arolick, Dan Romascanu, Jerome Grenier, Martin Thomson, Barbara Stark, Michael Haberler, and Mary Barnes for their WGLC review comments.

The authors would like to thank NENA for their work on [NENA] as it helped to provide some of the initial thinking.

The authors would also like to thank Cullen Jennings for his feedback as part of the IESG processing. Additionally, we would like to thank Alexey Melnikov, Dan Romascanu, and Robert Sparks.

10. References

10.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC5012] Schulzrinne, H. and R. Marshall, "Requirements for Emergency Context Resolution with Internet Technologies", RFC 5012, January 2008.

10.2. Informative References

- [802.1AB] "IEEE 802.1AB-2005 IEEE Standard for Local and Metropolitan Area Networks Station and Media Access Control Connectivity Discovery", May 2005, <<http://standards.ieee.org/getieee802/download/802.1AB-2005.pdf>>.
- [CDP] Wikipedia, "Cisco Discovery Protocol (CDP)", <http://en.wikipedia.org/wiki/Cisco_Discovery_Protocol>.
- [GEOPRIV-ARCH] Barnes, R., Lepinski, M., Cooper, A., Morris, J., Tschofenig, H., and H. Schulzrinne, "An Architecture for Location and Location Privacy in Internet Applications", Work in Progress, October 2009.
- [LBYR-REQS] Marshall, R., Ed., "Requirements for a Location-by-Reference Mechanism", Work in Progress, November 2009.

- [LIS-DISC] Thomson, M. and J. Winterbottom, "Discovering the Local Location Information Server (LIS)", Work in Progress, February 2010.
- [LIS2LIS] Winterbottom, J. and S. Norreys, "LIS to LIS Protocol Requirements", Work in Progress, November 2007.
- [NENA] "NENA 08-505, Issue 1, 2006 (December 21, 2006), NENA Recommended Method(s) for Location Determination to Support IP-Based Emergency Services - Technical Information Document (TID)", December 2006, <http://www.nena.org/sites/default/files/08-505_20061221.pdf>.
- [NSLP] Stiernerling, M., Tschofenig, H., Aoun, C., and E. Davies, "NAT/Firewall NSIS Signaling Layer Protocol (NSLP)", Work in Progress, February 2010.
- [RFC2113] Katz, D., "IP Router Alert Option", RFC 2113, February 1997.
- [RFC2865] Rigney, C., Willens, S., Rubens, A., and W. Simpson, "Remote Authentication Dial In User Service (RADIUS)", RFC 2865, June 2000.
- [RFC3068] Huitema, C., "An Anycast Prefix for 6to4 Relay Routers", RFC 3068, June 2001.
- [RFC3588] Calhoun, P., Loughney, J., Guttman, E., Zorn, G., and J. Arkko, "Diameter Base Protocol", RFC 3588, September 2003.
- [RFC3972] Aura, T., "Cryptographically Generated Addresses (CGA)", RFC 3972, March 2005.
- [RFC4119] Peterson, J., "A Presence-based GEOPRIV Location Object Format", RFC 4119, December 2005.
- [RFC4282] Aboba, B., Beadles, M., Arkko, J., and P. Eronen, "The Network Access Identifier", RFC 4282, December 2005.
- [RFC4361] Lemon, T. and B. Sommerfeld, "Node-specific Client Identifiers for Dynamic Host Configuration Protocol Version Four (DHCPv4)", RFC 4361, February 2006.
- [RFC4479] Rosenberg, J., "A Data Model for Presence", RFC 4479, July 2006.

- [RFC4795] Aboba, B., Thaler, D., and L. Esibov, "Link-local Multicast Name Resolution (LLMNR)", RFC 4795, January 2007.
- [RFC5201] Moskowitz, R., Nikander, P., Jokela, P., and T. Henderson, "Host Identity Protocol", RFC 5201, April 2008.
- [RFC5389] Rosenberg, J., Mahy, R., Matthews, P., and D. Wing, "Session Traversal Utilities for NAT (STUN)", RFC 5389, October 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.
- [TEREDO-SEL] Ward, N., "Teredo Server Selection", Work in Progress, July 2007.
- [TR069] "TR-069, CPE WAN Management Protocol v1.1, Version: Issue 1 Amendment 2", December 2007, <http://www.broadband-forum.org/technical/download/TR-069_Amendment-2.pdf>.
- [mDNS] Cheshire, S. and M. Krochmal, "Multicast DNS", Work in Progress, September 2009.

Authors' Addresses

Hannes Tschofenig
Nokia Siemens Networks
Linnoitustie 6
Espoo 02600
Finland

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

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

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