                  Location Information Server (LIS) Discovery
                     Using IP Addresses and Reverse DNS

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 is an Internet Standards Track document.

   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). Further information on
   Internet Standards is available in 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/rfc7216.

Table of Contents

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 7.

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.

```
                          _____
                       (/           \)
                      (( Internet ))
                       (_____/)
                            |
                            |
                  .- - - -|- - - - - - - - - - -.
                 (        |                      )
                 (    +--------+     +-------+    )
        Access   (    | Access |. . .|  LIS  |   )
        Network   (   |  Node  |     |       |    )
        (ISP)     (   +--------+     +-------+    )
                 (          \             \       )
                  `- - - -\- - - - - - - -\- - -'
                           \               \
                            \               |
                  .- - - - -\- - - - - - + -.
                 (           \           |    )
                 (    +------------+     :    )
                 (    | Residential |    |    )
    Residential  (    |  Gateway    |    :    )
      Network    (    +------------+     |    )
                 (        /      \      /     )
                 (       /        \    /      )
                 (   +--------+  +--------+   )
                 (   | Device |  | Device |   )
                 (   +--------+  +--------+   )
                 (                            )
                  `- - - - - - - - - - - - - -'
```
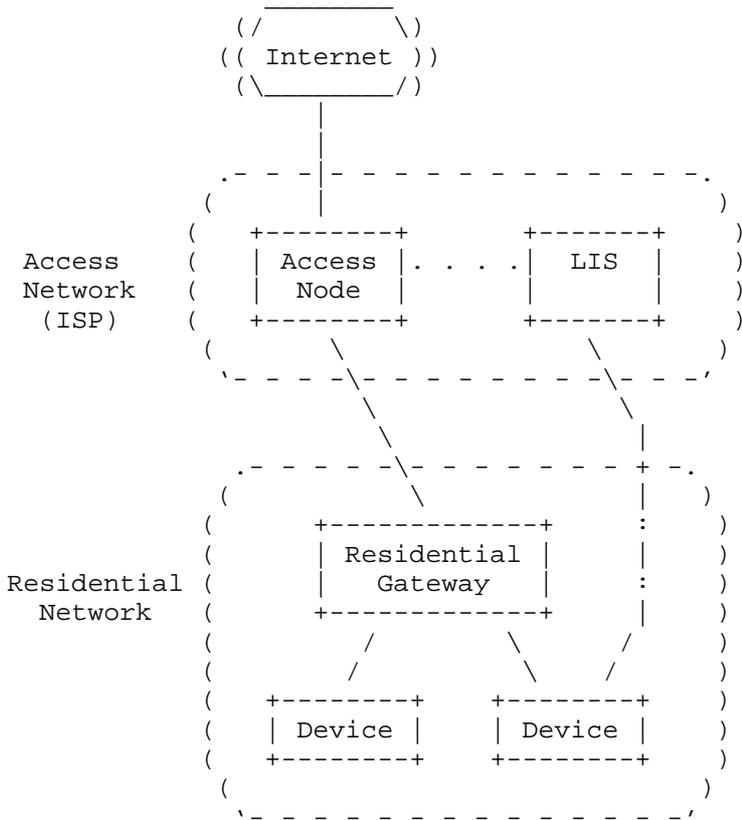
                Figure 1: Simplified Network Topology

   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. Security Features of Residential Gateways

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 from
which it is potentially reachable. 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 favored 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. If enabled, the Port Control Protocol (PCP) [RFC6887],
Universal Plug and Play (UPnP) [UPnP-IGD-WANIPConnection1], and NAT
Port Mapping Protocol (NAT-PMP) [RFC6886] are each able to provide
the external address of a residential gateway device. These, as well
as proprietary methods for determining other addresses, might 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 employed DNS server 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 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.8.2.0.0.0.0.0.0.0.0.8.b.d.0.1.0.0.2
      .ip6.arpa.

   o 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 because 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.

```
   _____        _____
  ( DNS ).....(/           \)      Public
  (__A__)     (( Internet ))      Address
               (_____/)       Space
                     |
                  [NAT]
   _____       _____|_____
  ( DNS )....(/           \)      Private
  (__B__)    (( ISP Network ))    Address Space
              (_____/)      (e.g., 10.0.0.0/8)
                    |
                [Gateway]
               ____|____
              (/          \)       Private
             (( Residence ))      Address Space
              (_____/)        (e.g., 192.168.0.0/16)
```
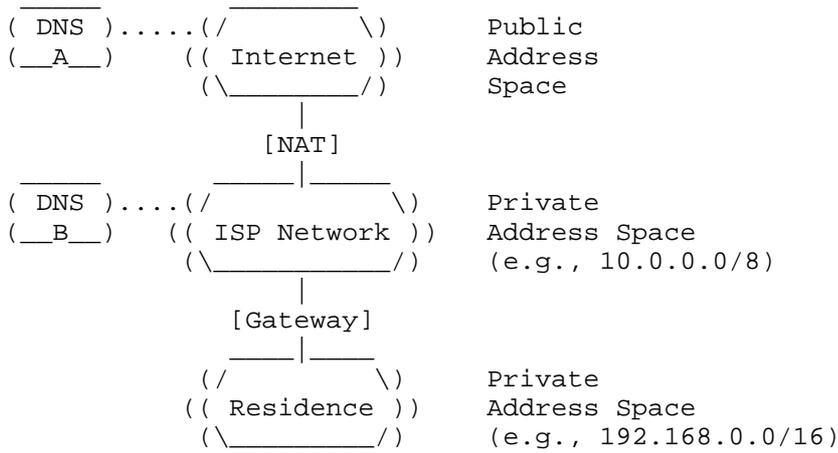
                  Figure 2: Address Space Example

   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 7.

   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 6) 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 address 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 a 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. 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 the 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 (LCP) [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
   [RFC7105], unless it has a secondary method for determining the
   authenticity of the LIS, such as a white list.

6. 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.

   DNS queries that are not blocked as [RFC6303] demands, or directed to
   servers outside the network, can cause the addresses that are in use
   within the network to be exposed outside of the network. For
   resolvers within the network, implementing [RFC6303] avoids this
   issue; otherwise, the problem cannot be avoided without blocking DNS
   queries to external servers.

7. 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 behavior 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 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]).

8. Acknowledgements

   Richard Barnes provided the text in Section 5.

9. References

9.1. 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.

   [RFC7105] Thomson, M. and J. Winterbottom, "Using Device-Provided
              Location-Related Measurements in Location Configuration
              Protocols", RFC 7105, January 2014.

9.2. 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.

   [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.

   [RFC6303] Andrews, M., "Locally Served DNS Zones", BCP 163, RFC
              6303, 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, <http://upnp.org/specs/gw/
              UPnP-gw-WANIPConnection-v1-Service.pdf>.

   [RFC6886] Cheshire, S. and M. Krochmal, "NAT Port Mapping Protocol
              (NAT-PMP)", RFC 6886, April 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/