ALTO                                                         S. Kiesel
Internet-Draft                                   University of Stuttgart
Intended status: Standards Track                          M. Stiemerling
Expires: September 15, 2011                               NEC Europe Ltd.
                                                               N. Schwan
                                                               M. Scharf
                                                 Alcatel-Lucent Bell Labs
                                                                M. Tomsu
                                                           Alcatel-Lucent
                                                                 H. Song
                                                                  Huawei
                                                          March 14, 2011


                       **ALTO Server Discovery Protocol**
                        **draft-kiesel-alto-3pdisc-05**

Abstract

   The goal of Application-Layer Traffic Optimization (ALTO) is to
   provide guidance to applications, which have to select one or several
   hosts from a set of candidates that are able to provide a desired
   resource.

   Entities seeking guidance need to discover and possibly select an
   ALTO server to ask.  This is called ALTO server discovery.  This memo
   describes an ALTO server discovery mechanism based on several
   alternative mechanisms that are applicable in a diverse set of ALTO
   deployments.

Status of this Memo

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

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

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

   This Internet-Draft will expire on September 15, 2011.

Copyright Notice

Table of Contents

## 1.  Introduction

   The goal of Application-Layer Traffic Optimization (ALTO) is to
   provide guidance to applications, which have to select one or several
   hosts from a set of candidates, that are able to provide a desired
   resource [RFC5693].  The requirements for ALTO are itemized in
   [I-D.ietf-alto-reqs].  ALTO is realized by a client-server protocol.
   ALTO clients send queries to ALTO servers, in order to solicit
   guidance.

   ALTO clients have to discover suitable ALTO servers.  Therefore the
   output of the herein defined ALTO discovery procedure tells the ALTO
   client which ALTO servers to send the queries to.  The ALTO discovery
   procedure, as part of the the ALTO client, can be embedded in the
   resource consumer, which will eventually access the desired resource.
   As an alternative, they can be embedded in a resource directory,
   which assists resource consumers in finding appropriate resource
   providers.  In some specific peer-to-peer application protocols these
   resource directories are called "trackers".  Finally the ALTO server
   discovery procedure can be embedded in the resource consumer, whereas
   the ALTO client is embedded in the resource directory.  ALTO queries,
   which are issued by a resource directory on behalf of a resource
   consumer, are referred to as third-party ALTO queries.  The various
   possibilities to place ALTO servers and the placement of ALTO clients
   is discussed in [I-D.stiemerling-alto-deployments].

   No matter where ALTO server and client are located, clients have to
   first find out if there is an ALTO server deployed that is in charge
   for them, and second they have to get the contact information of that
   server, i.e., the IP address, port number, and probably transport
   protocol (which defaults to TCP for [I-D.ietf-alto-protocol]).

   The goal of this memo is to propose a uniform mechanism for all types
   of ALTO client deployments that is implementable and deployable at a
   fast pace, i.e., without creating other deployment dependecies for
   ALTO.  We propose to use a combination of DHCP and DNS to retrieve
   the URL of the resposnsible ALTO server.

   Comments and discussions about this memo should be directed to the
   ALTO working group: alto@ietf.org.

## 1.1.  History

   This document represents a merge of features from two previous
   drafts:

   o  draft-kiesel-alto-3pdisc-04

   o   draft-song-alto-server-discovery-03

## 1.2.  Discovery Scenarios

   Figure 1 below shows an overview on the different entities of a
   generic ALTO framework.  The ALTO Server discovery mechanism is used
   by the p2p application in order retrieve the point of contact of the
   ALTO Service.

```
                                        +------+
                               +-----+  |  Peers
       +-----+      +------+   +=====|      |--+-oo
       |     |......|      |   |====+   oo+--*--+    o
       +-----+      +------+   |   o    *  ooooooo
       Source of      ALTO     |   o    *
       Topological   Service   |   o +--*--+
       information             +===o=|     | Tracker
                                   o +-----+ /Super-peer
                    ALTO Discovery o    o
                        Service    o    o
                      +------+    o    o
                      |      |oooooooooo
                      +------+
```

              Legend:

              ===    ALTO query protocol
              ooo    ALTO service discovery protocol
              ***    Application protocol (out of scope)
              ...    Provisioning or initialization (out of scope)

         Figure 1:  ALTO Discovery Overview

   Hereby the ALTO service discovery scenarios are classified into two
   types: one is the ALTO server discovery by the resource consumer, and
   the other is the ALTO server discovery by a third party, such as
   application trackers.  Before the specification of the discovery
   mechanism the following section illustrates and discusses both
   scenarios.

## 1.2.1.  ALTO Server Discovery by Resource Consumers

   The ALTO service discovery in some scenarios needs to be performed by
   the resource consumer itself.  In particular in p2p applications
   without a tracker like DHTs and other conventional client/server
   applications.

In addition also p2p application which are tracker based may embed
the ALTO client into the resurce consumer to allow peers a peer
selection after retrieving peer list from the application tracker.
Another option is that the reource consumer peer sends its ALTO
server address information to the application tracker or any other
third party entity, which in turn will contact the specific ALTO
server and in order to retrieve ALTO guidance on behalf of the
resource consumer.

The following figure illustrates this scenario, showing the
relationship between the different entities as discussed before.

```
                 +---------------+
                 |  ALTO Server  |
                 +---------------+
                     ^       ^         +-----------+
                     |       |         | ALTO      |
                     |   +---+---+     | Service   |
                     |   |tracker|    | Discovery |
                     |   +-------+     +---------+-+
                     |       |            o       o
      +------------+--+      |            o       o
      |P2P Application|ooooo|oooooooooooo      o
      |   Client A    |      |                  o
      +---------------+      |                  o
                             |                  o
                     +--+-------------+    o
                     | P2P Application|oooooo
                     |    Client B    |
                     +----------------+
```
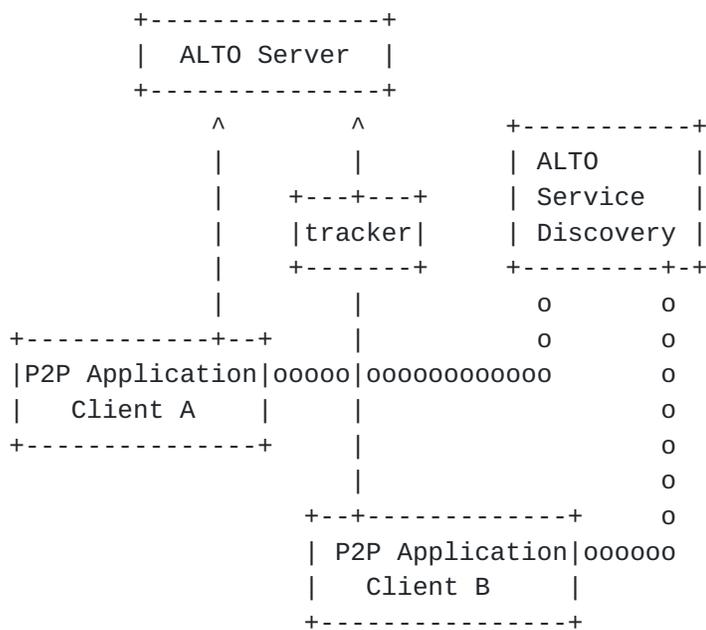
Figure 2: Resource Consumer ALTO Server Discovery (Example)

## 1.2.2.  ALTO Server Discovery by a Third Party

Some p2p applications have trackers, and these applications might not
need to have their clients looking for the ALTO server guidance.  In
these scenarios trackers query the ALTO servers for guidance
themselves, and then return the final ranked result to the
application clients.  However, application clients are distributed
among different network operators and autonomous systems.  Trackers
thus need to find different ALTO servers for the clients located in
different operator networks or autonomous systems.  In such scenarios
the discuvery is thus not performed by the resource consumer, but a
third party entity on behalf of the resource consumer.

Figure 3 shows an example for a third party ALTO server discovery.
For client 1, the tracker has not cached yet the mapping between

client 1's network operator and its ALTO server address, so it
queries the DNS server for the ALTO server address in that operator's
domain.  And then the tracker interacts with the ALTO server on
behalf of client 1(to get the network map and cost map), finally, the
ranked list is sent back to client 1.  For client 2, the tracker has
cached the mapping between client 2's network operator and its ALTO
server address, so it does not need to query the DNS for the address
of ALTO server 2.  If the Application tracker already has the network
map and cost map from ALTO Server2, then it does not to query the
ALTO Server for network map and cost map frequently.

```
              +-------------+    +-------------+
              | ALTO Server1|    | ALTO Server2|
              +-------------+    +-------------+
                 ^     |            ^    |
                3|    4|           b|   |c
                 |     |            |   |
                    v /----------+-\ v
     +---+         //////            \\\\\
     |   |      |||                        |||
     |   |<===>|                               |
     |DNS|  2  |      Application Tracker      |
     |   |      |||                        |||
     |   |         \\\\\\               /////
     +---+     ^     |     \-----------/   |
              |    |5                ^    |d
             1|    |                a|    |
              |    v                 |    v
            +-+---------+        +---+--------+
            |Application|        |Application |
            |client1    |        |client2     |
            +-----------+        +------------+
```
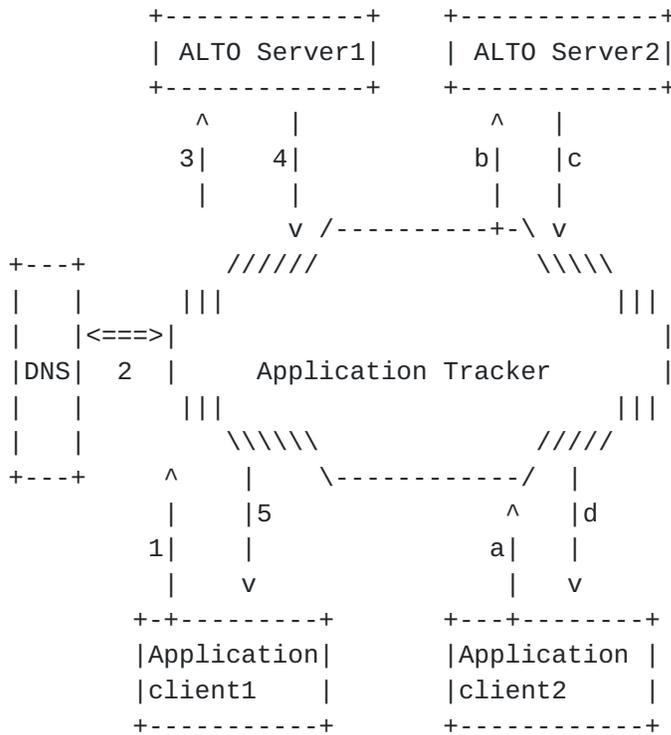
      Figure 3:  Third Party ALTO Server Discovery (Example)

## 1.3.  Pre-Conditions

The whole document assumes certain pre-conditions, such as:

o   The ALTO server discovery procedure is executed on a per IP
    address base.  Multiple IP addresses per interface or multiple IP
    addresses assigned to different IP interfaces require to repeat
    the procedure for every IP address.  It may be fine to group IP
    addresses according their domain suffixes and to perfom the
    procedure for such a group.  However, this is out of scope of this
    document.

o  The ALTO server discovery procedure is executed on a per IP family
   base, i.e., seperate for IPv4 and IPv6.  It is up to the ALTO
   client to decide which of the possible multiple results of
   different IP address families to use.  The choice of whether to
   use IPv4 or IPv6 is out of scope of this document.

o  A change of the IP address at an interface invalidates the result
   of the ALTO server discovery procedure.  For instance, if the IP
   address assigned to a mobile host changes due to host mobility, it
   is required to run the ALTO server discovery procedure for the new
   IP address without relying on earlier gained information.

[2](#). **Protocol Overview**

   We define multiple alternatives to discover the IP address of the
   ALTO server, as there are a number of ways possible how such
   information can be provided to the ALTO client.  The choice of method
   is up to the local network deployment.  For instance, there can be
   deployments where the ALTO server in charge for ALTO client is
   provisioned by the network operator and communicated to the ALTO
   client's host via a DHCP option, while in other deployments no such
   means may exist.

   The following figure illustrates the different protocols that are
   used to find the URI of a suitable ALTO server.

```
                   Descending order of Preference
             ~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~>


             --------------   --------   ---------------
             | User Input |   | DHCP |   | Reverse DNS |
             --------------   --------   ---------------
                   |            |              |
                    \           |             /
                     Retrieving the DNS suffix
                      \         |          /
                       --------+----------
                               |
                          -----------
                          | U-NAPTR |
                          -----------
                               |
                      Retrieving ALTO Server URI
                               |
                               |
                         Final DNS lookup
```
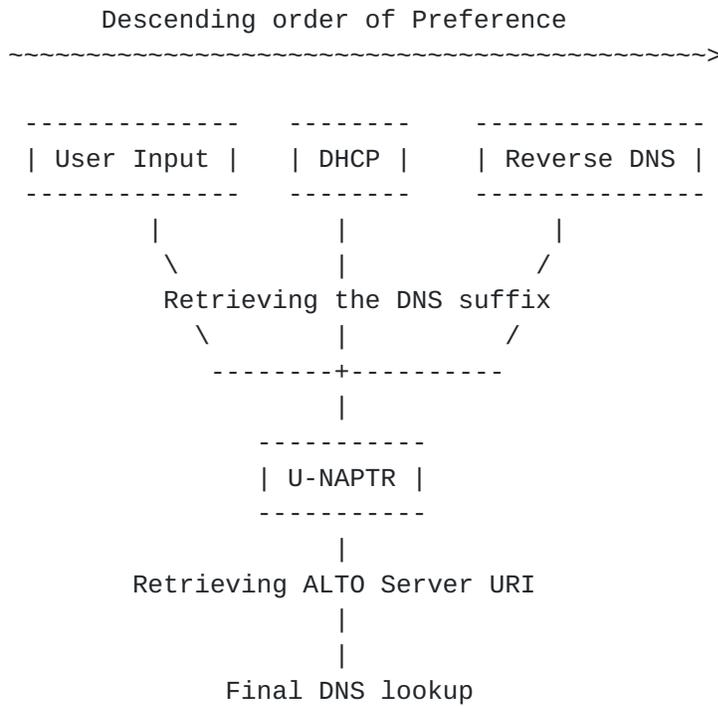
                     Figure 1: Protocol Overview

   The figure above illustrates the U-NAPTR based resolution process to
   retrieve the ALTO Server URL.  As a precondition for resolution the
   U-NAPTR process needs the right domain name as input.  This domain
   name is determined by the IP address of the client and the DNS suffix
   of the access network where the client is registered in.  In order to
   retrieve the DNS suffix we specify three options:

   User input:  a user may manually specify the DNS suffix on its own,
      either to access a 3rd party ALTO service provider or as it does
      know such information.

   DHCP:  a network provider provides the DNS suffix through a DHCP
      option.

   Reverse DNS:  the DNS system can be used to retrieve the DNS suffix
      through reverse lookup of an FQDN associated with an IP address.
      This is the last resort if all other options failed.

## 3.  Retrieving the URI by U-NAPTR

This section specifies the U-NAPTR based resolution process.  To
start the U-NAPTR resolution process a domain name as input is
needed.  Thus the section is devided into two parts: Section 3.1
describes the U-NAPTR resolution process itself.  How the client
identifies this DNS suffix of the access network where the resource
consumer is registered in is described in Section 3.2.

### 3.1.  U-NAPTR Resolution

ALTO servers are identified by U-NAPTR/DDDS (URI-Enabled NAPTR/
Dynamic Delegation Discovery Service) [RFC4848] application unique
strings, in the form of a DNS name.  An example is
'altoserver.example.com'.

Clients need to use the U-NAPTR [RFC4848] specification described
below to obtain a URI (indicating host and protocol) for the
applicable ALTO service.  In this document, only the HTTP and HTTPS
URL schemes are defined.  Note that the HTTP URL can be any valid
HTTP URL, including those containing path elements.

The following two DNS entries show the U-NAPTR resolution for
"example.com" to the HTTPS URL https://altoserver.example.com/secure
or the HTTP URL http://altoserver.example.com, with the former being
preferred.

```
    example.com.

    IN NAPTR 100  10   "u"    "ALTO:https"
         "!.*!https://altoserver.example.com/secure!"  ""

    IN NAPTR 200  10   "u"    "ALTO:http"
         "!.*!http://altoserver.example.com!"  ""
```

### 3.2.  Retrieving the Domain Name

The U-NAPTR resolution process requires a domain name as input.  The
algorithm that is applied to determine this domain name is described
in this section.  We specify three different options.  In option 1
the user manually configures a specific ALTO service instance that he
wants to use.  Option 2 defines a DHCP option to allow the network
service provider a remote configuration of the client.  In option 3
the client tries to get the domain name by performing a reverse DNS
lookup on its IP address.

The resource consumer may have private IP addresses and public IP

addresses and depending on the deployment it might be necessary to
determine for all IP addresses the ALTO server in charge of.  To
determine its public IP address the resource consumer may need to use
STUN[RFC5389] or BEP24[bep24].  For the following examples we assume
that the IP address of the resource consumer is a.b.c.d.

### 3.2.1.  Option 1: User input

A user may want to use a third party ALTO service instance.
Therefore we allow the user to specify a DNS suffix on its own, for
example in a config file option.  The DNS suffix given by the user is
combined with the IP address of the resource consumer to allow the
third party ALTO service to direct the client to a suitable ALTO
server based on the location of the client.  A possible DNS suffix
entered by the user may be:

    myaltoprovider.org

This DNS suffix is prepended with the IP address of the resource
consumer in reverse order to compose the domain name used for the
final U-NAPTR lookup Section 3.1.  In case there are multiple ALTO
servers deployed, the third party ALTO service instance can direct
the ALTO client to the ALTO server closest to the client based on the
IP address.

Multiple lookups with different domain names might be necessary to
complete the U-NAPTR resolution process.  If there is no response for
a lookup the domain name is shortened by one part for the succeeding
lookup, until a lookup is successful, as for example

    d.c.b.a.myaltoprovider.org.

    c.b.a.myaltoprovider.org.

    b.a.myaltoprovider.org.

    a.myaltoprovider.org.

    myaltoprovider.org.

### 3.2.2.  Option 2: DHCP

As a second option network operators can configure the domain name to
be used for service discovery within an access network.  RFC
5986[RFC5986] defines DHCP IPv4 and IPv6 access network domain name
options that identify a domain name that is suitable for service
discovery within the access network.  The ALTO server discovery
procedure uses these DHCP options to retrieve the domain name as an

input for the U-NAPTR resolution.  One example could be:

    example.com

### 3.2.3.  Option 3: Reverse DNS Lookup

The last option to get the domain name is to use a DNS PTR query for
the IP address of the resource consumer.  The local DNS server
resolves the IP address to the FQDN that also contains the DNS suffix
for the respective IP address.  A possible answer for a PTR lookup
for d.c.b.a.in-addr.apra might be, for example:

    d-c-b-a.dsl.westcoast.myisp.net

This domain name can be used for the final U-NAPTR lookup
Section 3.1.  If there is no response to the lookup the domain name
is shortened by one part for one succeeding lookup.  If there is
still no response we consider the reverse lookup being failed.  The
domain names used for the example as described above are:

    d-c-b-a.dsl.westcoast.myisp.net.

    dsl.westcoast.myisp.net.

## 4.  Applicability

This section discusses the applicability of the proposed solution
with respect to the resource consumer server discovery and the third
party deployment scenarios.  Each section discusses the proposed
steps that are needed to determine the ALTO Server URI.

### 4.1.  Applicability for Resource Consumer Server Discovery

In this scenario the ALTO server discovery procedure is performed by
the resource consumer, for example a peer in a P2P system.  After the
discovery the peer does the ALTO query on its own, or it might share
the ALTO server contact information with a third party, for example a
tracker, which then does the ALTO query on behalf of the peer.

To complete the ALTO server discovery process the resource consumer
first SHOULD check whether the user has provider the domain name
through manual configuration.  If this is not the case the next step
SHOULD be to check for the access network domain name DHCP option
(Section 3.2.2).  Finally the client SHOULD try to retrieve the
domain name by the last option, the DNS reverse lookup on its IP
address as described in Section 3.2.3.

In case the ALTO discovery client has determined the domain name
through one of the described options it proceedes with the U-NAPTR
lookup as described in Section 3.1.

### 4.2.  Applicability for Third Party Server Discovery

In case of the third party server discovery deployment scenario the
entity performing the ALTO server discovery process is different from
the resource consumer.  Typically the resource consumer is a peer
whereas the ALTO client is a resource directory which seeks for ALTO
guidance on behalf of the peer.  Another use case for the third party
discovery is an application that looks for ALTO guidance
transparently for the resource consumer, for example a CDN.

Here the ALTO server discovery process can also retrieve guidance
through the DHCP option or manual user configuration, but only if the
provided discovery information is forwarded by the resource consumer
to the third party entity.  In this case, additional mechanisms for
the forwarding of this discovery information need to be specified.
However these mechanisms are out of scope of this doument.

If the third party entity cannot obtain this discovery information,
the ALTO server discovery process relies on retrieving the domain
name used as input to the U-NAPTR lookup through reverse DNS lookup
of the IP address of the resource consumer as described in

Section 3.2.3.  Usually the third party entity already knows the IP
address of the resource consumer which was used to establish the
initial connection.  In general this IP address is a public address,
either of the resource consumer or of the last NAT on the path to the
ALTO client.  This makes the IP address a good candidate for the DNS
PTR query.  Thus, we expect that the DNS query will be successfully
resolved to the FQDN of the domain where the resource consumer is
registered in.

In case the resource consumer needs guidance for a different IP
address, for example one from a private network, we recommend that
the resource consumer discovers the server itself and forwards the
ALTO server contact information directly to the third party entity,
which in turn can then do the third party ALTO query.  Again,
forwarding the contact information from the resource consumer to the
third party entity is out of scope of this document.

5.  Deployment Considerations

   The mechanism specified in this document needs some configuration
   effort in order to work properly.  Especially the domain name
   retrieved through the reverse DNS lookup (PTR records) and the
   U-NAPTR entry need to be coordinated.  In this section we discuss
   this configuration for different scenarios.

5.1.  Private customers or very small businesses

   For private customers and very small businesses that are DSL or cable
   customers often a dynamically assigned IP address is provisioned.
   Here, the reverse DNS lookup (PTR records) are controlled by the ISP
   and they point to the ISP's domain, e.g.:

      p5B203EA1.dip.t-dialin.net.

      dslb-084-056-144-100.pools.arcor-ip.net.

      187-4-222-157.bnut3700.dsl.brasiltelecom.net.br.

      65-154-39-69.ispnetbilling.com.

      197-151-94-178.pool.ukrtel.net.

   In this case, it would be the responsibility of the respective ISP to
   provide U-NAPTR entries for the DNS suffix without the endhost part,
   e.g.:

      dip.t-dialin.net.

      pools.arcor-ip.net.

      bnut3700.dsl.brasiltelecom.net.br.

      ispnetbilling.com.

      pool.ukrtel.net.

5.2.  Medium-size customer networks

   The second class of customers have their own DNS domain but only one
   single upstream ISP, e.g.:

   (1)  ISP my-isp.net assigns an IP address a.b.c.d to its customer

(2)  The customer decides that reverse mapping for a.b.c.d should be
     whatever.customerdomain.com

(3)  If the customer wants to support ALTO, he has to ask the ISP for
     the URI of the ISP's ALTO server which can give guidance to
     a.b.c.d.  Assume that ISP replies it is
     http://altoserver.my-isp.net

(4)  The customer establishes a U-NAPTR entry for his domain

    customerdomain.com.    IN NAPTR 200  10   "u"    "ALTO:http"
         "!.*!http://altoserver.my-isp.net!"  ""

## 5.3.  Large Customers

For very large customers with multiple upstream connections we assume
that they have their very own traffic optimization policies and thus
run their own ALTO server anyway.  In this case they need to manage
their DNS entries accordingly.

6.  **IANA Considerations**

   This document registers the following U-NAPTR application service
   tag:


      Application Service Tag:  ALTO

      Defining Publication:  The specification contained within this
         document.

   This document registers the following U-NAPTR application protocol
   tags:

   o  Application Protocol Tag: http

      Defining Publication: RFC 2616 [RFC2616]

   o  Application Protocol Tag: https

      Defining Publication: RFC 2818 [RFC2818]

## 7. Security Considerations

### 7.1. General

This is still to be done in later revision of this draft, as the draft evolves heavily right now.

### 7.2. For U-NAPTR

The address of an ALTO server is usually well-known within an access network; therefore, interception of messages does not introduce any specific concerns.

The primary attack against the methods described in this document is one that would lead to impersonation of a ALTO server since a device does not necessarily have a prior relationship with a ALTO server.

An attacker could attempt to compromise ALTO discovery at any of three stages:

1.  providing a falsified domain name to be used as input to U-NAPTR

2.  altering the DNS records used in U-NAPTR resolution

3.  impersonation of the ALTO

This document focuses on the U-NAPTR resolution process and hence this section discusses the security considerations related to the DNS handling.  The security aspects of obtaining the domain name that is used for input to the U-NAPTR process is described in respective documents, such as [I-D.ietf-geopriv-lis-discovery].

The domain name that is used to authenticated the ALTO server is the domain name in the URI that is the result of the U-NAPTR resolution. Therefore, if an attacker were able to modify or spoof any of the DNS records used in the DDDS resolution, this URI could be replaced by an invalid URI.  The application of DNS security (DNSSEC) [RFC4033] provides a means to limit attacks that rely on modification of the DNS records used in U-NAPTR resolution.  Security considerations specific to U-NAPTR are described in more detail in [RFC4848].

An "https:" URI is authenticated using the method described in Section 3.1 of [RFC2818].  The domain name used for this authentication is the domain name in the URI resulting from U-NAPTR resolution, not the input domain name as in [RFC3958].  Using the domain name in the URI is more compatible with existing HTTP client software, which authenticate servers based on the domain name in the URI.

An ALTO server that is identified by an "http:" URI cannot be authenticated.  If an "http:" URI is the product of the ALTO discovery, this leaves devices vulnerable to several attacks.  Lower layer protections, such as layer 2 traffic separation might be used to provide some guarantees.

8.  Open Issues

   Here are a few open issues to be clarified:

   Handling of reverse DNS lookups for IPv6:  Refer to [RFC4472] for a
      discussion about the issues.

   Missing reverse DNS entries for an IP address:  There may be cases
      where the reverse DNS lookup does not yield any result.  However,
      this will leave the ALTO client with no choice, other than giving
      up.  This needs better documentation.

   How to handled multiple results:  For instance, a host behind a NAT
      that yields an ALTO server in the private IP address domain and
      one in the public IP address domain.  Whom to ask?

   Suffix Issues  Document issues with suffix information provided by
      DHCP or by other means.  For instance, a host behind a NAT may
      have a configured DNS suffix ".local".  This suffix is not usuable
      for the server discovery procedure.

9.  **Contributors**

   Hannes Tschofenig provided the initial input to the U-NAPTR solution
   part.  Hannes and Martin Thomson provided excellent feedback and
   input to the server discovery.

   The authors would also like to thank the following persons for their
   contribution to this document or predecessors:

      Roni Even

      Gustavo Garcia

      Yu-Shun Wang

      Victor Pascual

      Richard Alimi

      Yunfei Zhang

      Y. Richard Yang

      Xingfeng Jiang

      Jay Gu

      Ning Zong

      David Bryan

      Enrico Marocco

## 10. Conclusion

This document describes a general ALTO server discovery process and discusses how the process can be applied in different deployment scenarios, including the resouce consumer discovery as well as the third party discovery.

## 11.  References

### 11.1.  Normative References

[RFC2616]  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.

[RFC2818]  Rescorla, E., "HTTP Over TLS", RFC 2818, May 2000.

[RFC3958]  Daigle, L. and A. Newton, "Domain-Based Application
           Service Location Using SRV RRs and the Dynamic Delegation
           Discovery Service (DDDS)", RFC 3958, January 2005.

[RFC4033]  Arends, R., Austein, R., Larson, M., Massey, D., and S.
           Rose, "DNS Security Introduction and Requirements",
           RFC 4033, March 2005.

[RFC5389]  Rosenberg, J., Mahy, R., Matthews, P., and D. Wing,
           "Session Traversal Utilities for NAT (STUN)", RFC 5389,
           October 2008.

### 11.2.  Informative References

[I-D.ietf-alto-protocol]
           Alimi, R., Penno, R., and Y. Yang, "ALTO Protocol",
           draft-ietf-alto-protocol-04 (work in progress), May 2010.

[I-D.ietf-alto-reqs]
           Kiesel, S., Previdi, S., Stiemerling, M., Woundy, R., and
           Y. Yang, "Application-Layer Traffic Optimization (ALTO)
           Requirements", draft-ietf-alto-reqs-08 (work in progress),
           March 2011.

[I-D.ietf-geopriv-lis-discovery]
           Thomson, M. and J. Winterbottom, "Discovering the Local
           Location Information Server (LIS)",
           draft-ietf-geopriv-lis-discovery-15 (work in progress),
           March 2010.

[I-D.song-alto-server-discovery]
           Yongchao, S., Tomsu, M., Garcia, G., Wang, Y., and V.
           Avila, "ALTO Service Discovery",
           draft-song-alto-server-discovery-03 (work in progress),
           July 2010.

[I-D.stiemerling-alto-deployments]
           Stiemerling, M. and S. Kiesel, "ALTO Deployment

                 Considerations", draft-stiemerling-alto-deployments-03
                 (work in progress), June 2010.

   [RFC4472]  Durand, A., Ihren, J., and P. Savola, "Operational
                 Considerations and Issues with IPv6 DNS", RFC 4472,
                 April 2006.

   [RFC4848]  Daigle, L., "Domain-Based Application Service Location
                 Using URIs and the Dynamic Delegation Discovery Service
                 (DDDS)", RFC 4848, April 2007.

   [RFC5693]  Seedorf, J. and E. Burger, "Application-Layer Traffic
                 Optimization (ALTO) Problem Statement", RFC 5693,
                 October 2009.

   [RFC5986]  Thomson, M. and J. Winterbottom, "Discovering the Local
                 Location Information Server (LIS)", RFC 5986,
                 September 2010.

   [bep24]    Harrison, D., "Tracker Returns External IP",
                 BEP http://bittorrent.org/beps/bep_0024.html.

## Appendix A.  Acknowledgments

The authors would like to thank Haibin Song, Richard Alimi, and Roni
Even for fruitful discussions during the 75th IETF meeting.

Hannes Tschofenig provided the initial input to the U-NAPTR solution
part.  Hannes and Martin Thomson provided excellent feedback and
input to the server discovery.

Marco Tomsu and Nico Schwan are partially supported by the ENVISION
project (http://www.envision-project.org), a research project
supported by the European Commission under its 7th Framework Program
(contract no. 248565).  The views and conclusions contained herein
are those of the authors and should not be interpreted as necessarily
representing the official policies or endorsements, either expressed
or implied, of the ENVISION project or the European Commission.

Michael Scharf is supported by the German-Lab project
(http://www.german-lab.de) funded by the German Federal Ministry of
Education and Research (BMBF).

Martin Stiemerling is partially supported by the COAST project
(COntent Aware Searching, retrieval and sTreaming,
http://www.coast-fp7.eu), a research project supported by the
European Commission under its 7th Framework Program (contract no.
248036).  The views and conclusions contained herein are those of the
authors and should not be interpreted as necessarily representing the
official policies or endorsements, either expressed or implied, of
the COAST project or the European Commission.

Authors' Addresses

    Sebastian Kiesel
    University of Stuttgart Computing Center
    Allmandring 30
    Stuttgart  70550
    Germany

    Email: ietf-alto@skiesel.de
    URI:   http://www.rus.uni-stuttgart.de/nks/


    Martin Stiemerling
    NEC Laboratories Europe
    Kurfuerstenanlage 36
    Heidelberg  69115
    Germany

    Phone: +49 6221 4342 113
    Email: martin.stiemerling@neclab.eu
    URI:   http://ietf.stiemerling.org


    Nico Schwan
    Alcatel-Lucent Bell Labs
    Lorenzstrasse 10
    Stuttgart  70435
    Germany

    Email: nico.schwan@alcatel-lucent.com
    URI:   www.alcatel-lucent.com/bell-labs


    Michael Scharf
    Alcatel-Lucent Bell Labs
    Lorenzstrasse 10
    Stuttgart  70435
    Germany

    Email: michael.scharf@alcatel-lucent.com
    URI:   www.alcatel-lucent.com/bell-labs

Marco Tomsu
Alcatel-Lucent
Lorenzstrasse 10
Stuttgart  70435
Germany

Email: marco.tomsu@alcatel-lucent.com
URI:   www.alcatel-lucent.com/bell-labs


Haibin Song
Huawei

Email: melodysong@huawei.com