ALTO
Internet-Draft
Intended status: Standards Track
Expires: September 8, 2012

S. Kiesel
University of Stuttgart
M. Stiemerling
NEC Europe Ltd.
N. Schwan
M. Scharf
Alcatel-Lucent Bell Labs
H. Song
Huawei

March 7, 2012

ALTO Server Discovery
draft-ietf-alto-server-discovery-03

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 deployment scenarios.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

Status of this Memo

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. Hence, ALTO clients need to know the contact information of ALTO servers, which can provide appropriate guidance for a given resource consumer. Typically the closer an ALTO server is to a resource consumer the more accurate guidance it can provide. Thus a design objective is to automatically discover an ALTO server topologically close to the resource consumer, if available. Redirecting an ALTO client from one ALTO server to another, potentially closer, ALTO server raises several issues. First ALTO servers by definition provide Network Maps for the whole IP address space and thus can provide each client with a potentially useful answer. Second ALTO servers are deployed independently and are thus not necessarily aware of each other. The contact information of the ALTO server is thus retrieved by invoking the ALTO discover procedure defined in this document.

The ALTO protocol specification [I-D.ietf-alto-protocol] is based on HTTP. Therefore, it expects that the ALTO discovery procedure yields the HTTP(S) URI of the ALTO server's Information Resource Directory, which gives further information about the capabilities and services provided by that ALTO server. Further (DNS) lookups may be necessary in order to find out the ALTO server's IP address.

There are various architectural options where to place the ALTO client and the ALTO server discovery procedure:

- One option is that the ALTO client and the ALTO server discovery procedure are embedded directly in the resource consumer, i.e., the application protocol entity that will eventually initiate data transmission to/from the selected resource provider(s). In this case, the ALTO server discovery procedure might be able to interact with the user (i.e., prompt for a host name). Furthermore, it may use services such as DHCP, which are only available within the access network to which the resource consumer is connected.

- Another option is to integrate the ALTO client and the ALTO server discovery procedure into a third party such as a resource directory ("peer-to-peer tracker"), which issues ALTO queries on behalf of various resource consumers. This third party may reside

in a different part of the network (administrative domain) than the resource consumer. It may occur that said third party whishes to issue ALTO queries on behalf of a resouce consumer, but all it knows about the resource consumer is the source IP address of messages originating from it (i.e., the resource consumer's IP address or the "public" IP address of the outermost NAT in front of the resource consumer). This IP address will be the only input parameter to the ALTO server discovery procedure, which will have to find an ALTO server that can give appropriate guidance for that resource consumer.

A more detailed discussion of various options where to place the funcional entities comprising the overall ALTO architecture can be found in [I-D.ietf-alto-deployments].

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 dependencies for ALTO. We propose a schema which employs the U-NAPTR mechanism [RFC4848] to determine the URI of the ALTO server and where mutliple input methods to the U-NAPTR process can be used. U-NAPTR is used because the discovery mechanism must return an URI, and thus other discovery mechanisms are not applicable (e. g., DNS SRV records).

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

1.1. 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 peer-to-peer (P2P) application in order retrieve the point of contact of the ALTO Service.

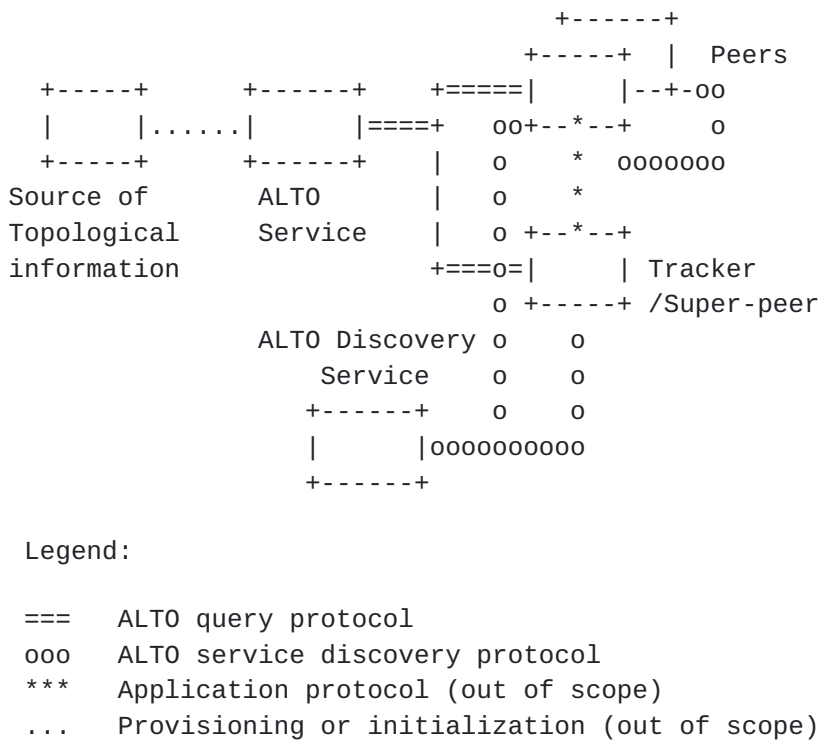

Figure 1: ALTO Discovery Overview

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.

1.1.1. 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 resource consumer to allow peers a selection of peers after retrieving the peer list from the application tracker. Another option is that the resource 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 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.

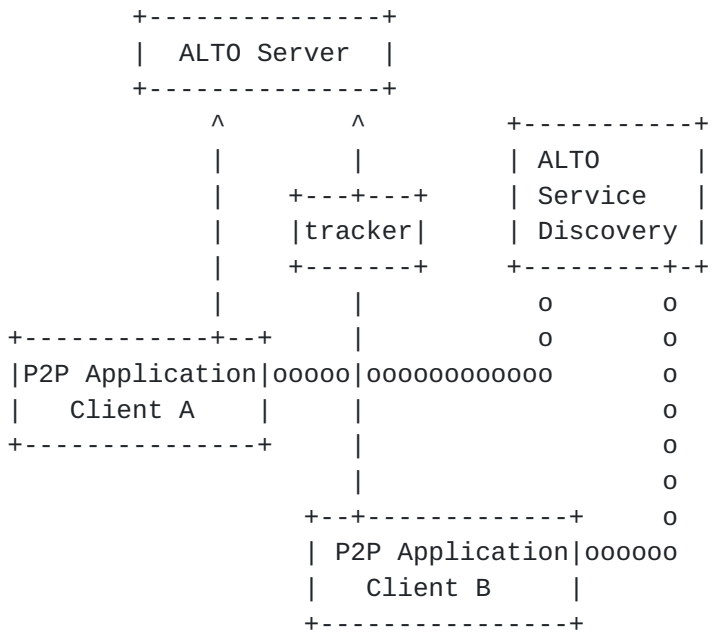

Figure 2: Resource Consumer ALTO Server Discovery (Example)

1.1.2. 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 discovery 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 Client1 (1), the tracker has not cached yet the mapping between Client1's network operator and its ALTO server address, so it uses the ALTO Discovery Service to determine the address of the ALTO server in that operator's domain (2). Then the tracker interacts with ALTO Server1 (3)(4) on behalf of Client1 (to get the network map and cost map), finally, the ranked list is sent back to Client1 (5). For Client2, the tracker has cached the mapping between Client2's network operator and ALTO Server2's address, so it does not need to perform the discovery process (which are the labels (a),(b), (c), and (d)). If the application tracker already has the network map and cost map from ALTO Server2, then it does not need to query the ALTO

Server for network map and cost map frequently.

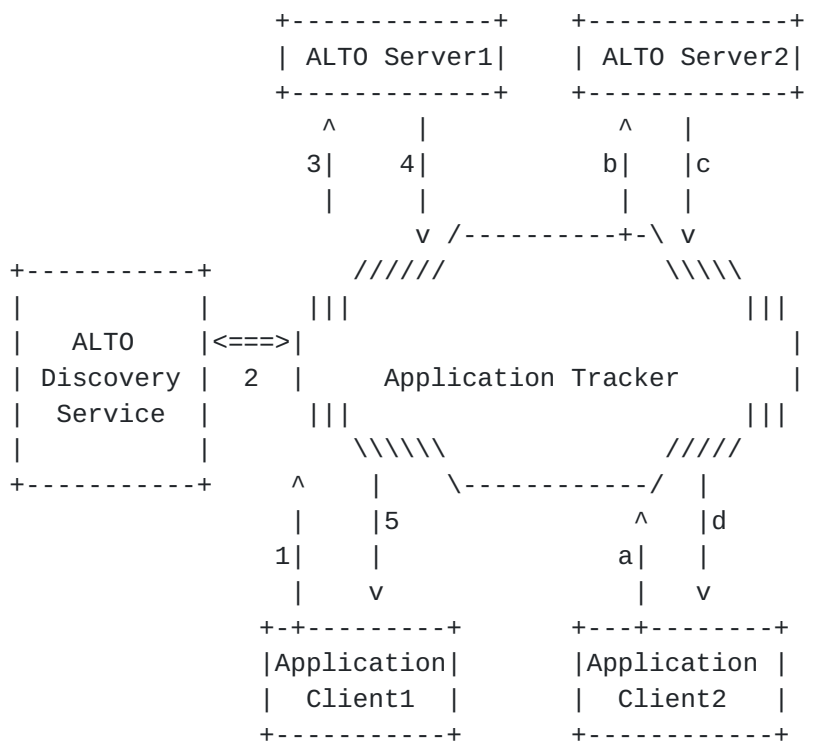

Figure 3: Third Party ALTO Server Discovery (Example)

1.2. Pre-Conditions

The whole document assumes certain pre-conditions, in particular:

- 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.

- 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.

- 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.

- There are several challenges with DNS on hosts with multiple interfaces [RFC6418], which can affect the ALTO server discovery. If the DNS resolution is performed on the wrong interface, it can return an ALTO server that could provide sub-optimal or wrong guidance. Finding the best ALTO server for multi-interfaced hosts is outside the scope of this document.

- The discovery procedure may need information about the public IP address and thus have to discover NATs. Details of NAT discovery are not discussed in this memo.

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.

It should be noted that there is no silver bullet solution to the ALTO server discovery, as there too many deployment scenarios in the server discovery space.

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

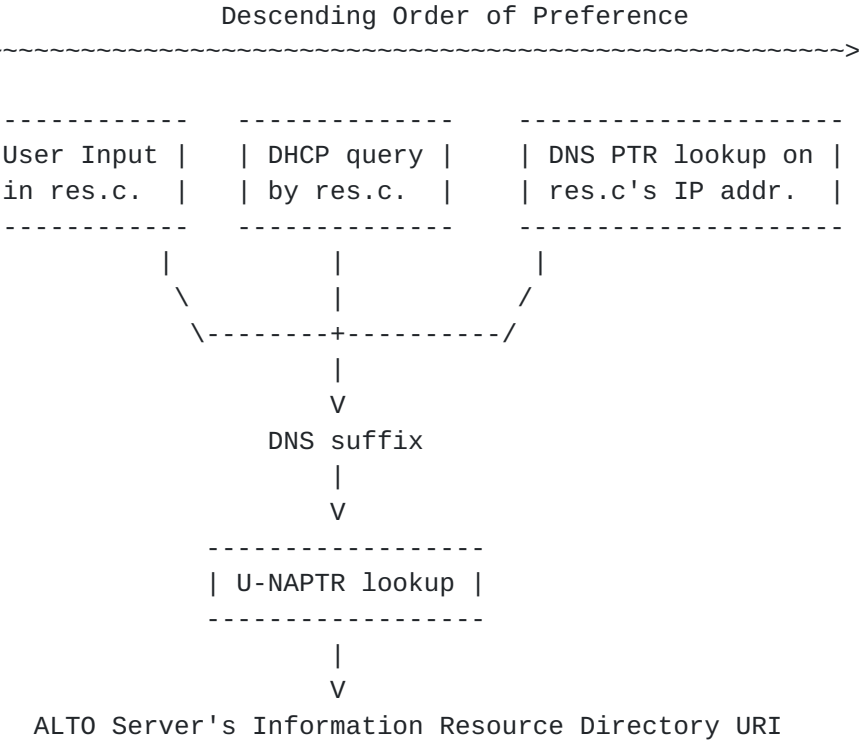

Figure 4: Protocol Overview

Figure 4 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, as are listed in descending order of preference. A client SHOULD use the first DNS suffix determined and MAY try other methods in case the U-NAPTR lookup failed.

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. This input MAY also origin from a web page where the user downloads the configuration, which is loaded as user input, or obtained by other discovery methods.

DHCP: A network provider MAY provide the DNS suffix through a DHCP option.

Reverse DNS: The DNS system MAY 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. It must be noted that the Reverse DNS lookup results in significant operational and deployment challanges and it is thus RECOMMENDED to avoid that method if possible.

This discovery method SHOULD be repeated if the resource consumer moves, i. e., if its IP address changes.

Instead of using the standard ALTO server discovery method, applications MAY also use own methods to discover an ALTO server. This variant is outside the scope of this document.

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 is required as input. Thus the section is devided into two parts: Section 3.2 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 3.1.

3.1. Retrieving the Domain Name

The U-NAPTR resolution process requires a domain name as input. The algorithm that SHOULD be 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]. Determining the correct IP address out of multiple options strongly depends on the deployment scenario but is out of scope for this document, although we discuss it to some extend in Section 4. For the following examples we assume that the IP address of the resource consumer is a.b.c.d.

3.1.1. 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 SHOULD be 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 3.2. This is useful in case there are multiple ALTO servers deployed, and the ALTO client should be redirected to the ALTO server closest to the client based on the IP address.

Multiple lookups with different domain names MAY be necessary to complete the U-NAPTR resolution process. If there is no response for a lookup i. e., if no ALTO NAPTR records are found, the domain name is shortened by the IP address part for the succeeding lookup, as for example

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

myaltoprovider.org.

In case not ALTO NAPTR records are found for both lookups we consider the discovery process based on user input as failed. A client MAY try one of the other options.

3.1.2. Option 2: DHCP

As a second option network operators MAY 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

3.1.3. 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 MAY be used for the final U-NAPTR lookup Section 3.2. If there is no response to the lookup the domain name ii MAY be 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.

3.2. U-NAPTR Resolution

The ALTO protocol specification [I-D.ietf-alto-protocol] expects that the ALTO discovery procedure yields the HTTP(S) URI of the ALTO server's Information Resource Directory, which gives further information about the capabilities and services provided by that ALTO server. The first step of the ALTO server discovery procedure (see Section 3.1) yielded an U-NAPTR/DDDS (URI-Enabled NAPTR/Dynamic Delegation Discovery Service) [RFC4848] application unique strings, in the form of a DNS name. An example is "example.com".

In the second step, the ALTO Server discovery procedure needs to use the U-NAPTR [RFC4848] specification described below to obtain a URI (indicating host and protocol) for the ALTO server's Information Resource Directory. In this document, only the HTTP and HTTPS URL schemes are defined, as the ALTO protocol specification defines the access over both protocols, but no other [I-D.ietf-alto-protocol]. Note that the HTTP URL can be any valid HTTP(s) 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/directory or the HTTP URL http://altoserver.example.com/directory, with the former being preferred.

```
example.com.

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

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

There is a potential that retrieveing the domain name or the U-NAPTR lookup itself does not yield to a result, i.e. no ALTO NAPTR record is found. In this case the discovery procedure failed for this IP address. It is RECOMMENDED that clients give up the discovery process and wait a period of time before repeating the procedure. Clients MAY repeat the discovery procedure for a different IP address instantaneously.

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 executes 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 3.1.2). Finally the client MAY try to retrieve the domain name by the last option, the DNS reverse lookup on its IP address as described in Section 3.1.3.

A client can have several candidate IP addresses that it may use for the discovery process. For example if it is located behind a NAT, a private and a public IP address may be used for the discovery process. It depends on the deployment scenario which of the IP addresses is the correct one. Thus it is out-of-scope of this document to specify how exactly the client finds the right IP address. However in the following we list methods that may be used in order to determine these candidate IP addresses. Generally in P2P environments peers already have implemented mechanisms for NAT-traversal. This includes proprietary solutions to determine a peer's public IP address, for example by asking a neighbour peer about its record of the own IP address. Non-proprietary solutions that are favorable include the Session Traversal Utilities for NAT (STUN) [RFC5986] protocol to determine the public address. If the client is behind a residential gateway another option may be to use Universal Plug and Play (UPnP) [UPnP-IGD-WANIPConnection1] or the NAT Port Mapping Protocol (NAT-PMP) [I-D.cheshire-nat-pmp].

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 3.2.

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 3.1.3. 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.

5.1. Reverse DNS Lookup

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.

5.1.1. 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.

5.1.2. 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!"  ""
```

5.1.3. 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.

5.2. Operational Considerations

A service discovery based on reverse DNS lookup results in several limitations:

First, there is no established unique way of maintaining the DNS tree, and there are different practices in different networks. Furthermore, it is possible that a lookup fails or that the returned value is not valid. For instance, it can point to a different domain. As a result, users of the ALTO discovery mechanism must be able to deal with failures of the reverse DNS lookup and react accordingly. As the reverse DNS lookup is the least preferred variant, failure of this discovery mechanism may imply that not ALTO server can be discovered and ALTO guidance is thus not available.

Second, determining a domain name from IP addresses by tree climbing is problematic, in particular for IPv6. [RFC4472] discusses the issues for IPv6.

Third, populating a DNS name space what looks like a reverse tree is a significant administrative DNS overhead.

Finally it must be emphasized that any tree walking procedure raises several issues. In this draft we therefore intentionally allow only one step for shortening domain names for any of the specified methods. Nevertheless, implementers of this specification SHOULD

consider skipping this step. For instance, there are different possible reasons why an ALTO NAPTR record cannot be found, including a timeout or a lack of name or record, and the discovery client needs heuristics to deal with all of them.

5.3. DHCP option for DNS Suffix

Section 3.1.2 describes the usage of a DHCP option which allows the network operator of the network where the ALTO client is attached to, to provide a DNS suffix. However, this assumes that this particular DHCP option is correctly passed from the DHCP server to the actual host with the ALTO client, and that the particular host understands this DHCP option. This memo assumes the client to be able to understand the proposed DHCP option, otherwise there is no further use of the DHCP option, but the client has to use the other proposed mechanisms.

There are well-known issues with the handling of DHCP options in home gateways. One issue is that unkown DHCP options are not passed through some home gateways, effectively eliminating the DHCP option.

Another well-known issues is the usage of home gateway specific DNS suffixes which "override" the DNS suffix provided by the network operator. For instance, a host behind a home gateway may receive a DNS suffix ".local" instead of "example.com". This suffix is not usuable for the server discovery procedure.

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:

- Application Protocol Tag: http

  Defining Publication: RFC 2616 [RFC2616]

- Application Protocol Tag: https

  Defining Publication: RFC 2818 [RFC2818]

7. Security Considerations

7.1. General

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

There are two different failures for the ALTO server discovery, which can both be caused by malicious attacks or by configuration problems, e. g., in case of DNS configuration errors or multi-homed hosts.

First, the discovery might not be able to discover an ALTO server, even if a suitable ALTO server exists. In that case, ALTO guidance will not be used. The resulting application performance and traffic distribution will correspond to a deployment scenario without ALTO guidance. But given that users cannot rely on the availability of an ALTO server, this results in no significant additional security risk.

Second, the discovery procedure may discover a sub-optimal or wrong ALTO server. Such an ALTO server may either not be able to provide information for a given resource consumer (e. g., behind a NAT), thus rendering the ALTO service useless. Alternatively, it may provide sub-optimal or forged information. In the latter case, attackers could try to use ALTO to affect the traffic distribution or the performance of applications. Users may then observe performance problems, and network operators could detect traffic anormalities. A potential counter-measure is to disable the use of the ALTO service.

Security issues of ALTO in general and potential solutions are also discussed in [I-D.ietf-alto-protocol].

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 an ALTO server since a device does not necessarily have a prior relationship with an 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 server

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 [RFC5986].

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 was 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. 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.

9. References

9.1. Normative References

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

[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.

[RFC6418] Blanchet, M. and P. Seite, "Multiple Interfaces and Provisioning Domains Problem Statement", RFC 6418, November 2011.

9.2. Informative References

[I-D.cheshire-nat-pmp] Cheshire, S., "NAT Port Mapping Protocol (NAT-PMP)", draft-cheshire-nat-pmp-03 (work in progress), April 2008.

[I-D.ietf-alto-deployments] Stiemerling, M. and S. Kiesel, "ALTO Deployment Considerations", draft-ietf-alto-deployments-03 (work in progress), November 2011.

[I-D.ietf-alto-protocol] Penno, R., Alimi, R., and Y. Yang, "ALTO Protocol", draft-ietf-alto-protocol-10 (work in progress), October 2011.

[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-11 (work in progress), July 2011.

[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.

[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, November 2001.

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

Appendix A. Contributors List and Acknowledgments

The initial version of this document was co-authored by Marco Tomsu <marco.tomsu@alcatel-lucent.com>.

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 its predecessors: Richard Alimi, David Bryan, Roni Even, Gustavo Garcia, Jay Gu, Xingfeng Jiang, Enrico Marocco, Victor Pascual, Y. Richard Yang, Yu-Shun Wang, Yunfei Zhang, Ning Zong.

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

Haibin Song
Huawei

Email: melodysong@huawei.com