

ALTO
Internet-Draft
Intended status: Standards Track
Expires: September 22, 2013

S. Kiesel
University of Stuttgart
M. Stiemerling
NEC Europe Ltd.
N. Schwan
Stuttgart, Germany
M. Scharf
Alcatel-Lucent Bell Labs
H. Song
Huawei
March 21, 2013

ALTO Server Discovery
draft-ietf-alto-server-discovery-08

Abstract

The goal of Application-Layer Traffic Optimization (ALTO) is to provide guidance to applications that have to select one or several hosts from a set of candidates capable of providing a desired resource. ALTO is realized by a client-server protocol. Before an ALTO client can ask for guidance it needs to discover one or more ALTO servers that can provide suitable guidance.

This document specifies a procedure for resource consumer initiated ALTO server discovery, which can be used if the ALTO client is embedded in the resource consumer.

Terminology and Requirements Language

This document makes use of the ALTO terminology defined in [[RFC5693](#)].

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

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 22, 2013.

Copyright Notice

Copyright (c) 2013 IETF Trust and the persons identified as the document authors. All rights reserved.

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

Table of Contents

1.	Introduction	4
2.	ALTO Server Discovery Procedure Overview	5
3.	ALTO Server Discovery Procedure Specification	6
3.1.	Step 1: Retrieving the Domain Name	6
3.1.1.	Step 1, Option 1: User input	6
3.1.2.	Step 1, Option 2: DHCP	6
3.2.	Step 2: U-NAPTR Resolution	7
4.	Deployment Considerations	8
4.1.	Issues with Home Gateways	8
4.2.	Issues with Multihoming, Mobility and Changing IP Addresses	8
5.	IANA Considerations	10
6.	Security Considerations	11
6.1.	General Security Considerations	11
6.2.	Security Considerations for U-NAPTR	11
7.	References	13
7.1.	Normative References	13
7.2.	Informative References	13
Appendix A.	Contributors List and Acknowledgments	15
	Authors' Addresses	16

1. Introduction

The goal of Application-Layer Traffic Optimization (ALTO) is to provide guidance to applications that have to select one or several hosts from a set of candidates capable of providing a desired resource [[RFC5693](#)]. ALTO is realized by a client-server protocol; see requirement AR-1 in [[RFC6708](#)]. Before an ALTO client can ask for guidance it needs to discover one or more ALTO servers that can provide suitable guidance.

This document specifies a procedure for resource consumer initiated ALTO server discovery, which can be used if the ALTO client is embedded in the resource consumer. In other words, this document tries to meet requirement AR-32 in [[RFC6708](#)] while AR-33 is out of scope. A different approach, which tries to meet requirement AR-33, i.e., third-party ALTO server discovery, is addressed in [[I-D.kist-alto-3pdisc](#)].

The ALTO protocol specification [[I-D.ietf-alto-protocol](#)] is based on HTTP and expects the discovery procedure to yield an HTTP(S) URI. Therefore, this procedure is based on U-NAPTR [[RFC4848](#)]. It tries to directly find one or more ALTO server(s) that can give suitable guidance to the ALTO client. Other schemes, such as discovering a random ALTO server (which might not be able to give suitable guidance to the client in question) and asking it to redirect the client to a better server, are not considered in this document.

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

2. ALTO Server Discovery Procedure Overview

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 ALTO server discovery procedure is performed in two steps:

1. One or - in case of multiple interfaces and/or IPv4/v6 dual stack operation - more DNS domain names are yielded, either by manual input or by means of DHCP.
2. These DNS domain names are used for U-NAPTR lookups yielding one or more URIs. Further DNS lookups may be necessary to determine the ALTO server's IP address(es).

The primary means for retrieving the DNS domain name is DHCP. However, there may be situations where DHCP is not available or does not return a suitable value. Furthermore, there might be situations in which the user wishes to override the value that could be retrieved from DHCP. In these situations, manual input may be used.

Typically, but not necessarily, the DNS domain name is the domain name in which the client is located, i.e., a PTR lookup on the client's IP address would yield a similar name. However, due to the widespread use of network address translation (NAT), trying to determine the DNS domain name through a PTR lookup on an interface's IP address is not recommended for resource consumer initiated ALTO server discovery.

3. ALTO Server Discovery Procedure Specification

As already outlined in [Section 2](#) the ALTO server discovery procedure is performed in two steps, which will be specified in [Section 3.1](#) and [Section 3.2](#), respectively.

3.1. Step 1: Retrieving the Domain Name

3.1.1. Step 1, Option 1: User input

A user may want to use an ALTO service instance provided by an entity that is not the operator of the underlying IP network. Therefore, we allow the user to specify a DNS domain name, for example in a configuration file option. An example domain name is:

```
my-alternative-alto-provider.example.org
```

In case no ALTO-specific NAPTR records are found, we consider the discovery process based on user input as failed. A client MAY try to continue with DHCP (see below). If DHCP-based discovery succeeds the software SHOULD inform the user that the user input has been ignored and replaced by information retrieved from the network.

3.1.2. Step 1, Option 2: DHCP

As a second option network operators may configure the domain name to be used for service discovery within an access network using DHCP.

[RFC 5986](#) [[RFC5986](#)] defines DHCP IPv4 and IPv6 access network domain name options to identify a domain name that is suitable for service discovery within the access network. [RFC 2132](#) [[RFC2132](#)] defines the DHCP IPv4 domain name option. While this option is less suitable, it still may be useful if the [RFC 5986](#) option is not available.

For IPv6, the ALTO server discovery procedure MUST try to retrieve DHCP option 57 (OPTION_V6_ACCESS_DOMAIN). If no such option can be retrieved the procedure fails for this interface. For IPv4, the ALTO server discovery procedure MUST try to retrieve DHCP option 213 (OPTION_V4_ACCESS_DOMAIN). If no such option can be retrieved, the procedure SHOULD try to retrieve option 15 (Domain Name). If neither option can be retrieved the procedure fails for this interface. If a result can be retrieved it will be used as an input for the next step (U-NAPTR resolution). One example result could be:

```
example.net
```


3.2. Step 2: U-NAPTR Resolution

The first step of the ALTO server discovery procedure (see [Section 3.1](#)) yielded one or - in case of multiple interfaces and/or IPv4/v6 dual stack operation - several domain names, which will be used as U-NAPTR/DDDS (URI-Enabled NAPTR/Dynamic Delegation Discovery Service) [[RFC4848](#)] application unique strings. An example is:

example.net

In the second step, the ALTO Server discovery procedure uses a U-NAPTR [[RFC4848](#)] lookup with the "ALTO" Application Service Tag and either the "http" or the "https" Application Protocol Tag to obtain one or more URIs (indicating protocol, host and possibly path elements) for the ALTO server's Information Resource Directory. In this document, only the HTTP and HTTPS URI 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 result can be any valid HTTP(S) URI.

The following two U-NAPTR resource records can be used for mapping "example.net" to the HTTPS URI `https://altoserver.example.net/secure/directory` or the HTTP URI `http://altoserver.example.net/directory`, with the former being preferred.

example.net.

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

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

If no ALTO-specific U-NAPTR records can be retrieved, the discovery procedure fails for this domain name (and the corresponding interface and IP protocol version). If further domain names yielded by Step 1 are known, the discovery procedure may perform the corresponding U-NAPTR lookups immediately. However, before retrying a lookup that has failed, a client MUST wait a time period that is appropriate for the encountered error (NXDOMAIN, timeout, etc.).

4. Deployment Considerations

4.1. Issues with Home Gateways

[Section 3.1.2](#) describes the usage of a DHCP option. It enables the network operator of the network, in which the ALTO client is located, to provide a DNS domain name. 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 unknown 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 domain names which "override" the DNS domain name provided by the network operator. For instance, a host behind a home gateway may receive a DNS domain name ".local" instead of "example.net". In general, this domain name is not usable for the server discovery procedure, unless a DNS server in the home gateway resolves the corresponding NAPTR lookup correctly, e.g., by means of a DNS split horizon approach.

4.2. Issues with Multihoming, Mobility and Changing IP Addresses

If the user decides to enter the DNS domain name manually, only one set of ALTO servers will be discovered, irrespectively of multihoming and mobility. Particularly in mobile scenarios this can lead to undesirable results.

The DHCP-based discovery method can discover different sets of ALTO servers for each interface and address family (i.e., IPv4/v6). In general, if a client wishes to communicate using one of its interfaces and using a specific IP address family, it SHOULD query the ALTO server(s) that have been discovered for this specific interface and address family. Selecting an interface and IP address family, as well as comparing results returned from different ALTO servers, is out of the 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 re-run the ALTO server discovery procedure without relying on earlier gained information.

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.

When using Virtual Private Network (VPN) connections there is usually no DHCP. The user has to enter the DNS domain name manually. For good optimization results, a DNS domain name corresponding to the VPN concentrator, not corresponding to the user's current location, has to be entered. Similar considerations apply for Mobile IP.

5. IANA Considerations

IANA is requested to register the following U-NAPTR [[RFC4848](#)] application service tag for ALTO:

Application Service Tag: ALTO

Intended usage: see [[RFC5693](#)] or: "The goal of Application-Layer Traffic Optimization (ALTO) is to provide guidance to applications that have to select one or several hosts from a set of candidates capable of providing a desired resource."

Defining Publication: The specification contained within this document

Contact information: The authors of this document

Author/Change controller: The IESG

Interoperability considerations: No interoperability issues are known or expected. This tag is to be registered specifically for ALTO, which is a new application without any legacy deployments.

Security considerations: see [Section 6](#) and in particular [Section 6.2](#) of this document.

Related publications: This document specifies a procedure for discovering an HTTP or HTTPS URI of an ALTO server. HTTP is specified in [[RFC2616](#)] and HTTPS is specified in [[RFC2818](#)]. The HTTP(S)-based ALTO protocol is specified in [[I-D.ietf-alto-protocol](#)].

Application Protocol Tag: This document specifies how to use the application service tag "ALTO" with the application protocol tags "http" (defining publication: [[RFC2616](#)] and "https" (defining publication: [[RFC2818](#)]), which have already been registered in the respective IANA registry. Therefore, IANA is not requested by this document to register any new application protocol tag.

6. Security Considerations

6.1. General Security Considerations

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 subsequently correspond to a deployment scenario without ALTO guidance.

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, said ALTO server may provide suboptimal 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 anomalies. 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](#)].

6.2. Security Considerations 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.

7. References

7.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC2132] Alexander, S. and R. Droms, "DHCP Options and BOOTP Vendor Extensions", [RFC 2132](#), March 1997.
- [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.
- [RFC4848] Daigle, L., "Domain-Based Application Service Location Using URIs and the Dynamic Delegation Discovery Service (DDDS)", [RFC 4848](#), April 2007.
- [RFC5986] Thomson, M. and J. Winterbottom, "Discovering the Local Location Information Server (LIS)", [RFC 5986](#), September 2010.

7.2. Informative References

- [I-D.ietf-alto-deployments] Stiemerling, M., Kiesel, S., and S. Previdi, "ALTO Deployment Considerations", [draft-ietf-alto-deployments-06](#) (work in progress), February 2013.
- [I-D.ietf-alto-protocol] Alimi, R., Penno, R., and Y. Yang, "ALTO Protocol", [draft-ietf-alto-protocol-14](#) (work in progress), February 2013.
- [I-D.kist-alto-3pdisc] Kiesel, S., Krause, K., and M. Stiemerling, "Third-Party ALTO Server Discovery (3pdisc)", [draft-kist-alto-3pdisc-02](#) (work in progress), February 2013.
- [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.
- [RFC4033] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "DNS Security Introduction and Requirements", [RFC 4033](#), March 2005.

- [RFC5693] Seedorf, J. and E. Burger, "Application-Layer Traffic Optimization (ALTO) Problem Statement", [RFC 5693](#), October 2009.
- [RFC6418] Blanchet, M. and P. Seite, "Multiple Interfaces and Provisioning Domains Problem Statement", [RFC 6418](#), November 2011.
- [RFC6708] Kiesel, S., Previdi, S., Stiemerling, M., Woundy, R., and Y. Yang, "Application-Layer Traffic Optimization (ALTO) Requirements", [RFC 6708](#), September 2012.

Appendix A. Contributors List and Acknowledgments

The initial version of this document was co-authored by Marco Tomsu.

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.

Olafur Gudmundsson provided an excellent DNS expert review on an earlier version of this document.

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.

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 Stiernerling is partially supported by the CHANGE project (<http://www.change-project.eu>), a research project supported by the European Commission under its 7th Framework Program (contract no. 257422). 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 CHANGE 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 Stiernerling
NEC Laboratories Europe
Kurfuerstenanlage 36
Heidelberg 69115
Germany

Phone: +49 6221 4342 113

Email: martin.stiernerling@neclab.eu

URI: <http://ietf.stiernerling.org>

Nico Schwan
Stuttgart, Germany

Email: ietf@nico-schwan.de

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

