ALTO
Internet-Draft
Intended status: Standards Track
Expires: June 21, 2013

S. Kiesel
University of Stuttgart
R. Penno
Cisco Systems

December 18, 2012

ALTO Server Discovery based on well-known IP Address
draft-kiesel-alto-ip-based-srv-disc-00

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.

This document establishes a well-known IP address for the ALTO service and specifies how ALTO clients embedded in the resource consumer can use it to access the ALTO service.

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

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 June 21, 2013.

Copyright Notice

Table of Contents

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]. An HTTP based ALTO client protocol is specified in [I-D.ietf-alto-protocol].

Before an ALTO client can ask for guidance it needs to discover one or more ALTO servers that can provide suitable guidance. Several algorithms have been specified that produce a suitable HTTP URI for a given ALTO client (i.e., the URI may vary for different clients or different points of network attachment, etc.). These approaches are based on user input or DHCP [I-D.ietf-alto-server-discovery], a "reverse DNS" (PTR) lookup [I-D.kist-alto-3pdisc], or redirection within the application protocol [I-D.kiesel-alto-alto4alto]. However, each of this approaches has technical or operational issues that will hinder the fast deployment of ALTO.

This document follows a different approach: it establishes a well-known address for the ALTO service (TBD: this approach could easily be generalized in order to discover other services as well. But this is for further study). All ALTO clients seeking ALTO guidance are expected to send requests to this address. It is then the duty of "the network" to direct the query to a suitable server. This (re-)directing could be done on several layers, e.g., by resolving a well-known DNS domain name to different IP addresses (DNS split horizon), or by routing IP packets with the well-known IP address to different servers. This document follows the second option, as ALTO is closely related to IP routing and routing costs.

This document specifies a procedure that 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. Note that AR-20 mandates that "an ALTO client protocol must be designed in a way that the ALTO service can be provided by an entity that is not the operator of the underlying IP network." Though not violating said requirement, the procedure specified here is not helpful to fulfill it.

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].

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

2. ALTO Server Discovery based on well-known IP Address

2.1. Well-Known ALTO Server Discovery IP Address (WkAsdIPa)

IANA is requested to register a single IPv4 address 192.0.0.X (TBD) and a single IPv6 address 2001::XXXX (TBD) within the respective Special Purpose Address Registries as the well-known IP anycast addresses for the ALTO service. These addresses are called WkAsdIPa (well-known ALTO server discovery IP address(es)) in this document.

2.2. Well-Known ALTO Server Discovery URIs (WkAsdURI)

The Well-Known ALTO Server Discovery URIs (WkAsdURI) are formed using the HTTP or HTTPS protocol identifier, the WkAsdIPa in their literal forms (for literal IPv6 addresses in URIs see [RFC2732]), and a constant suffix. That is, there are four WkAsdURIs (TBD: replace XXX with real value assigned by IANA):

http://192.0.0.X/alto

https://192.0.0.X/alto

http://[2001::XXXX]/alto

https://[2001::XXXX]/alto

2.3. ALTO Discovery Client behavior

ALTO Clients that need to discover an ALTO server use the HTTP GET method [RFC2616] to access one WkAsdURI, e.g. GET http://192.0.0.X/alto . They MUST be prepared to receive an HTTP 307 temporary redirect to the ALTO server's Information Resource Directory URI (Sec. 6.7 of [I-D.ietf-alto-protocol]).

2.4. ALTO Discovery Server behavior

ALTO discovery servers MUST listen on the WkAsdIPa on the HTTP and HTTPS ports for incoming HTTP(S) requests. They MUST answer GET requests to WkAsdURI using the 307 (Temporary Redirect) status code and redirect to an ALTO server's Information Resource Directory URI.

The ALTO discovery server MAY consider the client's address and other information when generating the reply, in order to redirect to different ALTO servers depending on the client's identity or location within the network topology.

The Information Resource Directory itself MUST NOT reside on a WkAsdIPa, and it MUST not reside on an URI that resolves via DNS to a

WkAsdIPa. After issuing the 307 status code ALTO discovery serves MUST close the HTTP(S) connection.

Rationale for the requirements in the previous paragraph: The goal is to keep the TCP connection to the WkAsdIPa as short as possible. When using anycast routing, IP packets belonging to an established TCP connection could be diverted to another ALTO discovery server due to state changes in the routing protocol or due to scheduled maintenance. Keeping the connection duration as short as possible reduces the risk of stalled or aborted connections. A UDP based lookup using one query and one reply packet would further reduce that risk. However, there seems not to be a well-standardized candidate protocol and studies have shown that short-lived TCP connections work well enough with anycast routing (http://www.nanog.org/meetings/nanog37/presentations/matt.levine.pdf).

TBD: do we need some URI such as http://192.0.0.X/discovery-server-identity in order to be able to identify the (misbehaving) discovery server that currently serves us?

TBD: how should the ALTO discovery server handle GET requests to other URIs or other HTTP methods?

TBD: should the discovery server always redirect http requests to the http URI of the information resource Directory and redirect https always to https? Or are there other reasonable scenarios?

3. Deployment Considerations

Network operators have to install one or more ALTO discovery servers as specified above. Depending on the the network deployment scenario they may use IP routing tables, HTTP proxies with URI rewriting, or other suitable mechanisms to direct GET-requests for a WkAsdURI to one of these servers.

[TBD: explain in more detail] This works fine even with cascaded access routers with NATs. After each router hop the operator may decide whether to handle the discovery requests, e.g., using a static routing table entry, or whether let them flow "automatically" towards the internet backbones using the default routing table entry.

TBD: what happens if an operator does not deploy these scheme? Requests could be dropped at administrative borders, or there could be one or several "public" default discovery servers.

[TBD: explain in more detail] The advantage of this scheme is that it does not need support in home gateways, which would harm quick deployment. This scheme also doesn't need new interfaces between the operating system and applications, e.g., for passing DHCP options from the operating system to the application.

4. IANA Considerations

4.1. Registration of IPv4 Special Purpose Address

IANA is requested to register a single IPv4 address in the IANA IPv4 Special Purpose Address Registry [RFC5736].

[RFC5736] itemizes some information to be recorded for all designations:

1. The designated address prefix.

Prefix: TBD by IANA. Prefix length: /32

2. The RFC that called for the IANA address designation.

This document.

3. The date the designation was made.

TBD.

4. The date the use designation is to be terminated (if specified as a limited-use designation).

Unlimited. No termination date.

5. The nature of the purpose of the designated address (e.g., unicast experiment or protocol service anycast).

protocol service anycast.

6. For experimental unicast applications and otherwise as appropriate, the registry will also identify the entity and related contact details to whom the address designation has been made.

N/A.

7. The registry will also note, for each designation, the intended routing scope of the address, indicating whether the address is intended to be routable only in scoped, local, or private contexts, or whether the address prefix is intended to be routed globally.

Typically used within a network operator's network domain, but in principle globally routable.

8. The date in the IANA registry is the date of the IANA action, i.e., the day IANA records the allocation.

TBD.

4.2. Registration of IPv6 Special Purpose Address

IANA is requested to register a single IPv6 address in the IANA IPv6 Special Purpose Address Block [RFC4773].

[RFC4773] itemizes some information to be recorded for all designations:

1. The designated address prefix.

Prefix: TBD by IANA. Prefix length: /128

2. The RFC that called for the IANA address designation.

This document.

3. The date the designation was made.

TBD.

4. The date the use designation is to be terminated (if specified as a limited-use designation).

Unlimited. No termination date.

5. The nature of the purpose of the designated address (e.g., unicast experiment or protocol service anycast).

protocol service anycast.

6. For experimental unicast applications and otherwise as appropriate, the registry will also identify the entity and related contact details to whom the address designation has been made.

N/A.

7. The registry will also note, for each designation, the intended routing scope of the address, indicating whether the address is intended to be routable only in scoped, local, or private contexts, or whether the address prefix is intended to be routed globally.

Typically used within a network operator's network domain, but in principle globally routable.

8. The date in the IANA registry is the date of the IANA action, i.e., the day IANA records the allocation.

TBD.

5. Security Considerations

TBD

Issue: how to deal with TLS certificates for HTTPS?

TBD: rules for filtering route at administrative boundaries

6. References

6.1. 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.

[RFC2732] Hinden, R., Carpenter, B., and L. Masinter, "Format for Literal IPv6 Addresses in URL's", RFC 2732, December 1999.

[RFC4773] Huston, G., "Administration of the IANA Special Purpose IPv6 Address Block", RFC 4773, December 2006.

[RFC5736] Huston, G., Cotton, M., and L. Vegoda, "IANA IPv4 Special Purpose Address Registry", RFC 5736, January 2010.

6.2. Informative References

[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] Alimi, R., Penno, R., and Y. Yang, "ALTO Protocol", draft-ietf-alto-protocol-13 (work in progress), September 2012.

[I-D.ietf-alto-server-discovery] Kiesel, S., Stiemerling, M., Schwan, N., Scharf, M., and S. Yongchao, "ALTO Server Discovery", draft-ietf-alto-server-discovery-06 (work in progress), November 2012.

[I-D.kiesel-alto-alto4alto] Kiesel, S., "Using ALTO for ALTO server selection", draft-kiesel-alto-alto4alto-00 (work in progress), July 2010.

[I-D.kist-alto-3pdisc] Kiesel, S. and M. Stiemerling, "Third-Party ALTO Server Discovery (3pdisc)", draft-kist-alto-3pdisc-01 (work in progress), October 2012.

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

[RFC6708] Kiesel, S., Previdi, S., Stiemerling, M., Woundy, R., and Y. Yang, "Application-Layer Traffic Optimization (ALTO) Requirements", RFC 6708, September 2012.

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/

Reinaldo Penno
Cisco Systems
170 West Tasman Dr
San Jose CA
USA

Email: repenno@cisco.com