ALTO
Internet-Draft
Intended status: Standards Track
Expires: January 2, 2015

S. Kiesel
University of Stuttgart
R. Penno
Cisco Systems

July 1, 2014

Application-Layer Traffic Optimization (ALTO) Anycast Address
draft-kiesel-alto-ip-based-srv-disc-03

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.

Terminology and Requirements Language

This document makes use of the ALTO terminology defined in RFC 5693 [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 January 2, 2015.

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]. The ALTO requirements are itemized in [RFC6708]. The ALTO protocol [RFC7285] is a client-server protocol, which uses HTTP [RFC7230] for message transport.

Before an ALTO client can ask for guidance, it needs to discover one or more ALTO servers that can provide suitable guidance. Several procedures 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 [RFC7286], 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 to be used as application-layer anycast address. 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. ALTO Anycast IP Address (AAIPA)

IANA is requested to register (see Section 4) a single IPv4 address 192.0.0.X (TBD) and a single IPv6 address 2001:YYYY::ZZZZ (TBD) within the respective Special Purpose Address Registries as the well-known IP anycast addresses for the ALTO service. These addresses are called AAIPA (ALTO Anycast IP Address(es)) in this document.

2.2. ALTO Anycast Uniform Resource Identificator (AAURI)

The ALTO Anycast Uniform Resource Identificators (AAURIs) are formed using the HTTP or HTTPS protocol identifier, the AAIPA in their literal forms (for literal IPv6 addresses in URIs see [RFC2732]), and a constant suffix. That is, there are four AAURIs (TBD: replace X, Y, Z with real values assigned by IANA):

http://192.0.0.X/alto

https://192.0.0.X/alto

http://[2001:YYYY::ZZZZ]/alto

https://[2001:YYYY::ZZZZ]/alto

2.3. ALTO Anycast Client Behavior

ALTO Clients that need to discover an ALTO server use the HTTP GET method [RFC7231] to access one AAURI, 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. 9 of [RFC7285]).

For hosts equipped with multiple interfaces and/or using IPv4/v6 dual stack, this discovery method might yield different Information Resource Directory URIs for each interface and address familly (i.e., IPv4/v6). In general, if a client wishes to communicate using one of its interfaces and using a specific IP address familiy, it SHOULD use this interface and the IP address associated with this interface to access the AAURI of the corresponding IP 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.

TBD: rules for retrying (timers, etc.) in case of failure.

TBD: rules for caching discovery results.

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.

2.4. ALTO Anycast Server Behavior

ALTO anycast servers MUST listen on the IPv4 and/or IPv6 AAIPA(s) on the HTTPS ports for incoming HTTPS requests and they SHOULD listen on these AAIPA(S) on the HTTP port for incoming HTTP requests. They MUST answer GET requests to AAURI using the 307 (Temporary Redirect) status code and redirect to an ALTO server's Information Resource Directory URI.

The Information Resource Directory itself MUST NOT reside on a AAIPA, and it MUST NOT reside on an URI that resolves via DNS to a AAIPA. After issuing the 307 status code ALTO anycast servers MUST close the HTTP(S) connection.

> Rationale for the requirements in the previous paragraph: The goal is to keep the TCP connection to the AAIPA as short as possible. When using anycast routing, IP packets belonging to an established TCP connection could be diverted to another ALTO anycast 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 packet and one reply packet (e.g., based on httpu) would eliminate that risk. However, there seems not to be a well-standardized candidate protocol and studies [Levine2006] suggest that short-lived TCP connections work well enough with anycast routing.

An ALTO anycast server MUST redirect an HTTPS request for an HTTPS AAURI to an HTTPS IRD URI. It MAY redirect an HTTP request for an HTTP AAURI to an HTTP IRD URI, but it MAY also redirect it to an HTTPS IRD URI.

The ALTO anycast 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.

TBD: do we need some URI such as http://192.0.0.X/server-identity in

order to be able to identify the (misbehaving) ALTO anycast server that currently serves us?

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

3. Deployment Considerations

Network operators have to install one or more ALTO anycast 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 AAURI 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. As an alternative, there could be "public" ALTO anycast servers to answer all queries that had not been answered in the respective originating access network. These servers could use the third-party ALTO server discovery procedure [I-D.kist-alto-3pdisc] to find the redirection target based on the client's IP address.

[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.

[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.

[RFC7230] Fielding, R. and J. Reschke, "Hypertext Transfer Protocol (HTTP/1.1): Message Syntax and Routing", RFC 7230, June 2014.

[RFC7231] Fielding, R. and J. Reschke, "Hypertext Transfer Protocol (HTTP/1.1): Semantics and Content", RFC 7231, June 2014.

6.2. Informative References

[I-D.ietf-alto-deployments]
Stiemerling, M., Kiesel, S., Previdi, S., and M. Scharf, "ALTO Deployment Considerations", draft-ietf-alto-deployments-09 (work in progress), February 2014.

[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., Krause, K., and M. Stiemerling, "Third-Party ALTO Server Discovery (3pdisc)", draft-kist-alto-3pdisc-05 (work in progress), January 2014.

[Levine2006]
Levine, M., Lyon, B., and T. Underwood, "TCP Anycast - Don't believe the FUD. Operational experience with TCP and Anycast.", Presentation at NANOG37 http://www.nanog.org/meetings/nanog37/presentations/matt.levine.pdf, June 2006.

[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.

[RFC7285] Alimi, R., Penno, R., and Y. Yang, "Application-Layer Traffic Optimization (ALTO) Protocol", RFC 7285, June 2014.

[RFC7286] Kiesel, S., Stiemerling, M., Schwan, N., Scharf, M., and H. Song, "Application-Layer Traffic Optimization (ALTO) Server Discovery", RFC 7286, June 2014.

Authors' Addresses

Sebastian Kiesel
University of Stuttgart Information 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