

Network Working Group
Internet-Draft
Intended status: Informational
Expires: January 12, 2012

S. Kiesel, Ed.
University of Stuttgart
S. Previdi
Cisco Systems, Inc.
M. Stiemerling
NEC Europe Ltd.
R. Woundy
Comcast Corporation
Y R. Yang
Yale University
July 11, 2011

Application-Layer Traffic Optimization (ALTO) Requirements
draft-ietf-alto-reqs-11.txt

Abstract

Many Internet applications are used to access resources, such as pieces of information or server processes, which are available in several equivalent replicas on different hosts. This includes, but is not limited to, peer-to-peer file sharing applications. 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. This guidance shall be based on parameters that affect performance and efficiency of the data transmission between the hosts, e.g., the topological distance. The ultimate goal is to improve performance (or Quality of Experience) in the application while reducing resource consumption in the underlying network infrastructure.

This document enumerates requirements for specifying, assessing, or comparing protocols and implementations.

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 12, 2012.

Copyright Notice

Copyright (c) 2011 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. Terminology and Architectural Framework	5
2.1. Requirements Notation	5
2.2. ALTO Terminology	5
2.3. Architectural Framework for ALTO	6
3. ALTO Requirements	7
3.1. ALTO Client Protocol	7
3.1.1. General Requirements	7
3.1.2. Host Group Descriptor Support	7
3.1.3. Rating Criteria Support	8
3.1.4. Placement of Entities and Timing of Transactions	9
3.1.5. Protocol Extensibility	12
3.1.6. Error Handling and Overload Protection	12
3.2. ALTO Server Discovery	13
3.3. Security and Privacy	14
4. IANA Considerations	16
5. Security Considerations	17
5.1. High-level security considerations	17
5.2. Information Disclosure Scenarios	17
5.2.1. Classification of Information Disclosure Scenarios	17
5.2.2. Discussion of Information Disclosure Scenarios	18
5.3. Security Requirements	19
6. References	20
6.1. Normative References	20
6.2. Informative References	20
Appendix A. Contributors List and Acknowledgments	21
Authors' Addresses	22

1. Introduction

The motivation for Application-Layer Traffic Optimization (ALTO) is described in the ALTO problem statement [RFC5693].

The goal of ALTO is to provide information which can help peer-to-peer (P2P) applications to make better decisions with respect to peer selection. However, ALTO may be useful for non-P2P applications as well. For example, clients of client-server applications may use information provided by ALTO to select one of several servers or information replicas. As another example, ALTO information could be used to select a media relay needed for NAT traversal. The goal of these informed decisions is to improve performance (or Quality of Experience) in the application while reducing resource consumption in the underlying network infrastructure.

Usually, it would be difficult or even impossible for application entities to acquire this information by other mechanisms (e.g., using measurements between the peers of a P2P overlay), because of complexity or because it is based on network topology information, network operational costs, or network policies, which the respective network provider does not want to disclose in detail.

The logical entities that provide the ALTO service do not take part in the actual user data transport, i.e., they do not implement functions for relaying user data. They may be placed on various kinds of physical nodes, e.g., on dedicated servers, as auxiliary processes in routers, on "trackers" or "super peers" of a P2P application, etc.

2. Terminology and Architectural Framework

2.1. Requirements Notation

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 [RFC2119].

2.2. ALTO Terminology

This document uses the following ALTO-related terms, which are defined in [RFC5693]:

Application, Peer, P2P, Resource, Resource Identifier, Resource Provider, Resource Consumer, Transport Address, Overlay Network, Resource Directory, ALTO Service, ALTO Server, ALTO Client, ALTO Query, ALTO Response, ALTO Transaction, Local Traffic, Peering Traffic, Transit Traffic, Application protocol, ALTO Client Protocol, Provisioning protocol.

Furthermore, the following additional terms will be used:

- o Host Group Descriptor: Information used to describe one or more Internet hosts (such as the resource consumer which seeks ALTO guidance, or one or more candidate resource providers) and their location within the network topology. This can be, for example, a single IP address, an address prefix or address range that contains the host(s), or an autonomous system (AS) number. Different options may provide different levels of detail. Depending on the system architecture, this may have implications on the quality of the guidance ALTO is able to provide, on whether recommendations can be aggregated, and on how much privacy-sensitive information about users might be disclosed to additional parties.
- o Host Characteristics Attribute: Properties of a host (other than the host group descriptor), in particular related to its attachment to the network. This information may be stored in an ALTO server and transmitted via an ALTO protocol. It may be evaluated according to the rating criteria.
- o Rating Criterion: The condition or relation that defines the "better" in "better-than-random peer selection", which is the ultimate goal of ALTO. Examples may include "host's Internet access is not subject to volume based charging (flat rate)" or "low topological distance". Some rating criteria, such as "low topological distance", need to include a reference point, i. e., "low topological distance from a given resource consumer", which

can be described by means of a host group descriptor.

2.3. Architectural Framework for ALTO

There are various architectural options for how ALTO could be implemented, and specifying or mandating one specific architecture is out of the scope of this document.

The ALTO Working Group Charter [ALTO-charter] itemizes several key components, which shall be elaborated and specified by the ALTO Working Group. The ALTO problem statement [RFC5693] defines a terminology (see Section 2.2) and presents a figure that gives a high-level overview of protocol interaction between ALTO elements.

This document itemizes requirements for the following components and information elements that are part of the above-mentioned architecture:

- o An ALTO client protocol, which is used for sending ALTO queries and ALTO responses between ALTO client and ALTO server.
- o The discovery mechanism, which will be used by ALTO clients in order to find out where to send ALTO requests.
- o The overall architecture, especially with respect to security and privacy issues.
- o Host group descriptors, which are used to describe the location of a host in the network topology.
- o Rating criteria, i. e., conditions or relations that shall be evaluated in order to generate the ALTO guidance.

3. ALTO Requirements

[*** Note to the RFC editor: before publication as an RFC, please remove the draft version number from the requirements numbering, i.e., change ARv11-1 to AR-1, and so on. Furthermore, remove this note. ***]

3.1. ALTO Client Protocol

3.1.1. General Requirements

REQ. ARv11-1: The ALTO service is provided by one or more ALTO servers. ALTO servers MUST implement an ALTO client protocol, for receiving ALTO queries from ALTO clients and for sending the corresponding ALTO responses.

REQ. ARv11-2: ALTO clients MUST implement an ALTO client protocol, for sending ALTO queries to ALTO servers and for receiving the corresponding ALTO responses.

REQ. ARv11-3: The format of the ALTO query message MUST allow the ALTO client to solicit guidance for selecting appropriate resource providers.

REQ. ARv11-4: The format of the ALTO response message MUST allow the ALTO server to express its guidance for selecting appropriate resource providers.

The detailed specification of an ALTO client protocol is out of the scope of this document. However, this document enumerates requirements for ALTO, to be considered when specifying, assessing, or comparing protocols and implementations.

3.1.2. Host Group Descriptor Support

The ALTO guidance is based on the evaluation of several resource providers or groups of resource providers, which are characterized by means of host group descriptors, considering one or more rating criteria.

REQ. ARv11-5: An ALTO client protocol MUST support the host group descriptor types "IPv4 address prefix" and "IPv6 address prefix". They can be used to specify the IP address of one host, or an IP address range (in CIDR notation), which contains all hosts in question.

REQ. ARv11-6: An ALTO client protocol MUST be extensible to enable support of other host group descriptor types in future. An ALTO

client protocol specification MUST define an appropriate procedure for adding new host group descriptor types, e.g., by establishing an IANA registry.

REQ. ARv11-7: ALTO clients and ALTO servers MUST clearly identify the type of each host group descriptor sent in ALTO queries or responses.

REQ. ARv11-8: For host group descriptor types other than "IPv4 address prefix" and "IPv6 address prefix", the host group descriptor type identification MUST be supplemented by a reference to a facility, which can be used to translate host group descriptors of that type to IPv4/IPv6 address prefixes, e.g., by means of a mapping table or an algorithm.

REQ. ARv11-9: Protocol functions for mapping other host group descriptor types to IPv4/IPv6 address prefixes SHOULD be designed and specified as part of an ALTO client protocol, and the corresponding address mapping information SHOULD be made available by the same entity that wants to use these host group descriptors within an ALTO client protocol. However, an ALTO server or an ALTO client MAY also send a reference to an external mapping facility, e.g., a translation table to be obtained via an alternative mechanism.

REQ. ARv11-10: An ALTO client protocol specification MUST define mechanisms, which can be used by the ALTO server to indicate that a host group descriptor used by the ALTO client is of an unsupported type, or that the indicated mapping mechanism could not be used.

REQ. ARv11-11: An ALTO client protocol specification MUST define mechanisms, which can be used by the ALTO client to indicate that a host group descriptor used by the ALTO server is of an unsupported type, or that the indicated mapping mechanism could not be used.

3.1.3. Rating Criteria Support

REQ. ARv11-12: An ALTO client protocol specification MUST define a rating criterion that can be used to express and evaluate the "relative operator's preference." This is a relative measure, i.e., it is not associated with any unit of measurement. A more-preferred rating according to this criterion indicates that the application should prefer the respective candidate resource provider over others with less-preferred ratings (unless information from non-ALTO sources suggests a different choice, such as transmission attempts suggesting that the path is currently congested). The operator of the ALTO server does not have to disclose how and based on which data the ratings are actually computed. Examples could be: cost for peering or transit traffic, traffic engineering inside the network, and other

policies.

REQ. ARv11-13: An ALTO client protocol MUST be extensible to enable support of other rating criteria types in future. An ALTO client protocol specification MUST define an appropriate procedure for adding new rating criteria types, e.g., by establishing an IANA registry.

REQ. ARv11-14: ALTO client protocol specifications MUST NOT define rating criteria closely related to the instantaneous network congestion state, whose primary aim is to serve an alternative to established congestion control strategies, such as using TCP-based transport.

One design assumption for ALTO is that it is acceptable that the host characteristics attributes, which are stored and processed in the ALTO servers for giving the guidance, are updated rather infrequently. Typical update intervals may be several orders of magnitude longer than the typical network-layer packet round-trip time (RTT). Therefore, ALTO cannot be a replacement for TCP-like congestion control mechanisms. The definition of alternate approaches for congestion control is explicitly a non-goal for the ALTO working group [ALTO-charter].

REQ. ARv11-15: Applications using ALTO guidance MUST NOT rely on the ALTO guidance to avoid causing network congestion. Instead, applications MUST use other appropriate means, such as TCP based transport, to avoid causing excessive congestion.

REQ. ARv11-16: The ALTO query message SHOULD allow the ALTO client to express which rating criteria should be considered, as well as their relative relevance for the specific application that will eventually make use of the guidance.

REQ. ARv11-17: The ALTO response message SHOULD allow the ALTO server to express which rating criteria have been considered when generating the response.

REQ. ARv11-18: An ALTO client protocol specification MUST define mechanisms, which can be used by the ALTO client and the ALTO server to indicate that a rating criteria used by the other party is of an unsupported type.

3.1.4. Placement of Entities and Timing of Transactions

With respect to the placement of ALTO clients, several modes of operation exist:

- o One mode of ALTO operation is that an ALTO client may be 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 order to access the desired resource. For example, an ALTO client could be integrated into the peer of a P2P application that uses a distributed algorithm such as "query flooding" for resource discovery.
- o Another mode of operation is to integrate the ALTO client into a third party such as a resource directory, which may issue ALTO queries to solicit preference on potential resource providers, considering the respective resource consumer. For example, an ALTO client could be integrated into the tracker of a tracker-based P2P application, in order to request ALTO guidance on behalf of the peers contacting the tracker.

REQ. ARv11-19: An ALTO client protocol MUST support the mode of operation in which the ALTO client is directly embedded in the resource consumer.

REQ. ARv11-20: An ALTO client protocol MUST support the mode of operation in which the ALTO client is embedded in a third party, which performs queries on behalf of resource consumers.

REQ. ARv11-21: An ALTO client protocol MUST be designed in a way that the ALTO service can be provided by an entity which is not the operator of the underlying IP network.

REQ. ARv11-22: An ALTO client protocol MUST be designed in a way that different instances of the ALTO service operated by different providers can coexist.

With respect to the timing of ALTO queries, several modes of operation exist:

- o In target-aware query mode, an ALTO client performs the ALTO query when the desired resource and a set of candidate resource providers are already known, i. e., after DHT lookups, queries to the resource directory, etc.
- o In target-independent query mode, ALTO queries are performed in advance or periodically, in order to receive comprehensive, "target-independent" guidance, which will be cached locally and evaluated later, when a resource is to be accessed.

REQ. ARv11-23: An ALTO client protocol MUST support at least one of these two modes, either the target-aware or the target-independent query mode.

REQ. ARv11-24: An ALTO client protocol SHOULD support both the target-aware and the target-independent query mode.

REQ. ARv11-25: An ALTO client protocol SHOULD support version numbering, TTL (time-to-live) attributes, and/or similar mechanisms in ALTO transactions, in order to enable time validity checking for caching, and to enable comparisons of multiple recommendations obtained through redistribution.

REQ. ARv11-26: An ALTO client protocol SHOULD allow the ALTO server to add information about appropriate modes of re-use to its ALTO responses. Re-use may include redistributing an ALTO response to other parties, as well as using the same ALTO information in a resource directory to improve the responses to different resource consumers, within the specified lifetime of the ALTO response. The ALTO server SHOULD be able to express that

- o no re-use should occur
- o re-use is appropriate for a specific "target audience", i.e., a set of resource consumers explicitly defined by a list of host group descriptors. The ALTO server MAY specify a "target audience" in the ALTO response, which is only a subset of the known actual "target audience", e.g., if required by operator policies
- o re-use is appropriate for any resource consumer that would send (or cause a third party sending on behalf of it) the same ALTO query (i.e., with the same query parameters, except for the resource consumer ID, if applicable) to this ALTO server
- o re-use is appropriate for any resource consumer that would send (or cause a third party sending on behalf of it) the same ALTO query (i.e., with the same query parameters, except for the resource consumer ID, if applicable) to any other ALTO server, which was discovered (using an ALTO discovery mechanism) together with this ALTO server
- o re-use is appropriate for any resource consumer that would send (or cause a third party sending on behalf of it) the same ALTO query (i.e., with the same query parameters, except for the resource consumer ID, if applicable) to any ALTO server in the whole network

REQ. ARv11-27: An ALTO client protocol MUST support the exchange of ALTO transactions even if the ALTO client is located in the private address realm behind a network address translator (NAT). There are different types of NAT, see [RFC4787] and [RFC5382].

3.1.5. Protocol Extensibility

REQ. ARv11-28: An ALTO client protocol MUST include support for adding protocol extensions in a non-disruptive, backward-compatible way.

REQ. ARv11-29: An ALTO client protocol MUST include protocol versioning support, in order to clearly distinguish between incompatible versions of the protocol.

3.1.6. Error Handling and Overload Protection

REQ. ARv11-30: An ALTO client protocol MUST use TCP based transport.

REQ. ARv11-31: An ALTO client protocol specification MUST specify mechanisms, or detail how to leverage appropriate mechanisms provided by underlying protocol layers, which can be used by an ALTO server to inform clients about an impending or occurring overload situation, and require them to throttle their query rate.

In particular, as a simple way of achieving some basic form of throttling, an ALTO server MAY answer ALTO queries with a "Retry After: {point in time | time delta}" message. This "Retry After" MAY be sent as part of the ALTO reply together with the requested guiding information, or as a standalone (error) message not giving the requested guidance.

REQ. ARv11-32: An ALTO client protocol specification MUST specify mechanisms, or detail how to leverage appropriate mechanisms provided by underlying protocol layers, which can be used by an ALTO server to inform clients about an impending or occurring overload situation, and redirect them to another ALTO server.

REQ. ARv11-33: An ALTO client protocol specification MUST specify mechanisms, or detail how to leverage appropriate mechanisms provided by underlying protocol layers, which can be used by an ALTO server to inform clients about an impending or occurring overload situation, and terminate the conversation with the ALTO client.

REQ. ARv11-34: An ALTO client protocol specification MUST specify mechanisms, or detail how to leverage appropriate mechanisms provided by underlying protocol layers, which can be used by an ALTO server to inform clients about its inability to answer queries due to technical problems or system maintenance, and advise them to retry after an indicated point in time or after an indicated period of time has elapsed.

REQ. ARv11-35: An ALTO client protocol specification MUST specify

mechanisms, or detail how to leverage appropriate mechanisms provided by underlying protocol layers, which can be used by an ALTO server to inform clients about its inability to answer queries due to technical problems or system maintenance, and redirect them to another ALTO server.

REQ. ARv11-36: An ALTO client protocol specification MUST specify mechanisms, or detail how to leverage appropriate mechanisms provided by underlying protocol layers, which can be used by an ALTO server to inform clients about its inability to answer queries due to technical problems or system maintenance, and terminate the conversation with the ALTO client.

Note: The existence of the above-mentioned protocol mechanisms does not imply that an ALTO server must use them when facing an overload, technical problem, or maintenance situation, respectively. Some servers may be unable to use them in that situation, or they may prefer to simply refuse the connection or not to send any answer at all.

3.2. ALTO Server Discovery

An ALTO client protocol is supported by one or more ALTO server discovery mechanisms, which may be used by ALTO clients in order to determine one or more ALTO servers, to which ALTO requests can be sent. This section enumerates requirements for an ALTO client, as well as general requirements to be fulfilled by the ALTO server discovery mechanisms.

REQ. ARv11-37: ALTO clients which are embedded in the resource consumer MUST be able to use an ALTO server discovery mechanism, in order to find one or several ALTO servers that can provide ALTO guidance suitable for the resource consumer. This mode of operation is called "resource consumer initiated ALTO server discovery".

REQ. ARv11-38: ALTO clients which are embedded in a resource directory and perform third-party ALTO queries on behalf of a remote resource consumer MUST be able to use an ALTO server discovery mechanism, in order to find one or several ALTO servers that can provide ALTO guidance suitable for the respective resource consumer. This mode of operation is called "third-party ALTO server discovery".

REQ. ARv11-39: ALTO clients MUST be able to perform resource consumer initiated ALTO server discovery, even if they are located behind a network address translator (NAT).

REQ. ARv11-40: ALTO clients MUST be able to perform third-party ALTO server discovery, even if they are located behind a network address

translator (NAT).

REQ. ARv11-41: ALTO clients MUST be able to perform third-party ALTO server discovery, even if the resource consumer, on behalf of which the ALTO query will be sent, is located behind a network address translator (NAT).

REQ. ARv11-42: ALTO server discovery mechanisms SHOULD leverage an existing protocol or mechanism, such as DNS, DHCP, or PPP based automatic configuration, etc. A single mechanism with a broad spectrum of applicability SHOULD be preferred over several different mechanisms with narrower scopes.

REQ. ARv11-43: Every ALTO server discovery mechanism SHOULD be able to return the respective contact information for multiple ALTO servers.

REQ. ARv11-44: Every ALTO server discovery mechanism SHOULD be able to indicate preferences for each returned ALTO server contact information.

3.3. Security and Privacy

REQ. ARv11-45: An ALTO client protocol specification MUST specify mechanisms for the authentication of ALTO servers, or how to leverage appropriate mechanisms provided by underlying protocol layers.

REQ. ARv11-46: An ALTO client protocol specification MUST specify mechanisms for the authentication of ALTO clients, or how to leverage appropriate mechanisms provided by underlying protocol layers.

REQ. ARv11-47: An ALTO client protocol specification MUST specify mechanisms for the encryption of messages, or how to leverage appropriate mechanisms provided by underlying protocol layers.

REQ. ARv11-48: The operator of an ALTO server MUST NOT assume that an ALTO client will implement mechanisms or comply with rules that limit the ALTO client's ability to redistribute information retrieved from the ALTO server to third parties.

REQ. ARv11-49: An ALTO client protocol MUST support different levels of detail in queries and responses, in order to protect the privacy of users, to ensure that the operators of ALTO servers and other users of the same application cannot derive sensitive information.

REQ. ARv11-50: An ALTO client protocol MAY include mechanisms that can be used by the ALTO client when requesting guidance to specify the resource (e.g., content identifiers) it wants to access. An ALTO

server MUST provide adequate guidance even if the ALTO client prefers not to specify the desired resource (e.g., keeps the data field empty). The mechanism MUST be designed in a way that the operator of the ALTO server cannot easily deduce the resource identifier (e.g., file name in P2P file sharing) if the ALTO client prefers not to specify it.

REQ. ARv11-51: An ALTO client protocol specification MUST specify appropriate mechanisms for protecting the ALTO service against DoS attacks, or how to leverage appropriate mechanisms provided by underlying protocol layers.

4. IANA Considerations

This requirements document does not mandate any immediate IANA actions. However, such IANA considerations may arise from future ALTO specification documents which try to meet the requirements given here.

5. Security Considerations

5.1. High-level security considerations

High-level security considerations for the ALTO service can be found in the "Security Considerations" section of the ALTO problem statement document [RFC5693].

5.2. Information Disclosure Scenarios

The unwanted disclosure of information is one key concern related to ALTO. This section presents a classification and discussion of information disclosure scenarios and potential countermeasures.

5.2.1. Classification of Information Disclosure Scenarios

- o (1) Excess disclosure of ALTO server operator's data to an authorized ALTO client. The operator of an ALTO server has to feed information, such as tables mapping host group descriptors to host characteristics attributes, into the server, thereby enabling it to give guidance to ALTO clients. Some operators might consider the full set of this information confidential (e.g., a detailed map of the operator's network topology), and might want to disclose only a subset of it or somehow obfuscated information to an ALTO client.
- o (2) Disclosure of the application behavior to the ALTO server. The operator of an ALTO server could infer the application behavior (e.g., content identifiers in P2P file sharing applications, or lists of resource providers that are considered for establishing a connection) from the ALTO queries sent by an ALTO client.
- o (3) Disclosure of ALTO server operator's data (e.g., network topology information) to an unauthorized third party. There are a three sub-cases here:
 - * (3a) An ALTO server sends the information directly to an unauthorized ALTO client.
 - * (3b) An unauthorized party snoops on the data transmission from the ALTO server to an authorized ALTO client.
 - * (3c) An authorized ALTO client knowingly forwards the information it had received from the ALTO server to an unauthorized party.

- o (4) Disclosure of the application behavior to an unauthorized third party.
- o (5) Excess retrieval of ALTO server operator's data by collaborating ALTO clients. Several authorized ALTO clients could ask an ALTO server for guidance, and redistribute the responses among each other (see also case 3c). By correlating the ALTO responses they could find out more information than intended to be disclosed by the ALTO server operator.

5.2.2. Discussion of Information Disclosure Scenarios

Scenario (1) may be addressed by the ALTO server operator choosing the level of detail of the information to be populated into the ALTO server and returned in the responses. For example, by specifying a broader address range (i.e., a shorter prefix length) than a group of hosts in question actually uses, an ALTO server operator may control to some extent how much information about the network topology is disclosed. Furthermore, access control mechanisms for filtering ALTO responses according to the authenticated ALTO client identity might be installed in the ALTO server, although this might not be effective given the lack of efficient mechanisms for addressing (3c) and (5), see below.

(2) can and needs to be addressed in several ways: If the ALTO client is embedded in the resource consumer, the resource consumer's IP address (or the "public" IP address of the outermost NAT in front of the resource consumer) is disclosed to the ALTO server as a matter of principle, because it is in the source address fields of the IP headers. By using a proxy, the disclosure of source addresses to the ALTO server can be avoided at the cost of disclosing them to said proxy. If, in contrast, the ALTO client is embedded in a third party (e.g., a resource directory) which issues ALTO requests on behalf of resource consumers, it is possible to hide the exact addresses of the resource consumers from the ALTO server, e.g., by zeroing-out or randomizing the last few bits of IP addresses. However, there is the potential side effect of yielding inaccurate results.

The disclosure of candidate resource providers' addresses to the ALTO server can be avoided by allowing ALTO clients to use the target-independent query mode. In this mode of operation, guiding information (e.g., "maps") is retrieved from the ALTO server and used entirely locally by the ALTO client, i.e., without sending host location attributes of candidate resource providers to the ALTO server. In the target-aware query mode, this issue can be addressed by ALTO clients through obfuscating the identity of candidate resource consumers, e.g., by specifying a broader address range (i.e., a shorter prefix length) than a group of hosts in question

actually uses, or by zeroing-out or randomizing the last few bits of IP addresses. However, there is the potential side effect of yielding inaccurate results.

(3a), (3b), and (4) may be addressed by authentication, access control, and encryption schemes for the ALTO client protocol. However, deployment of encryption schemes might not be effective given the lack of efficient mechanisms for addressing (3c) and (5), see below.

Straightforward authentication and encryption schemes will not help solving (3c) and (5), and there is no other simple and efficient mechanism known. The cost of complex approaches, e.g., based on digital rights management (DRM), might easily outweigh the benefits of the whole ALTO solution, and therefore they are not considered as a viable solution. That is, ALTO server operators must be aware that (3c) and (5) cannot be prevented from happening, and therefore they should feed only such data into an ALTO server, which they do not consider sensitive with respect to (3c) and (5).

These insights are reflected in the requirements in this document.

5.3. Security Requirements

For a set of specific security requirements please refer to Section 3.3 of this document.

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.

6.2. Informative References

- [ALTO-charter]
Marocco, E. and V. Gurbani, "Application-Layer Traffic Optimization (ALTO) Working Group Charter (<http://tools.ietf.org/wg/alto/charters?item=charter-alto-2011-04-28.txt>)", April 2011.
- [RFC4787] Audet, F. and C. Jennings, "Network Address Translation (NAT) Behavioral Requirements for Unicast UDP", BCP 127, RFC 4787, January 2007.
- [RFC5382] Guha, S., Biswas, K., Ford, B., Sivakumar, S., and P. Srisuresh, "NAT Behavioral Requirements for TCP", BCP 142, RFC 5382, October 2008.
- [RFC5693] Seedorf, J. and E. Burger, "Application-Layer Traffic Optimization (ALTO) Problem Statement", RFC 5693, October 2009.

Appendix A. Contributors List and Acknowledgments

The initial version of this document was co-authored by Laird Popkin.

The authors would like to thank

- o Vijay K. Gurbani <vkg@alcatel-lucent.com>
- o Enrico Marocco <enrico.marocco@telecomitalia.it>

for fostering discussions that lead to the creation of this document, and for giving valuable comments on it.

The authors were supported by the following people, who have contributed to this document:

- o Richard Alimi <ralimi@google.com>
- o Zoran Despotovic <despotovic@docomolab-euro.com>
- o Jason Livingood <Jason_Livingood@cable.comcast.com>
- o Saverio Niccolini <saverio.niccolini@nw.neclab.eu>
- o Michael Scharf <michael.scharf@alcatel-lucent.com>
- o Nico Schwan <nico.schwan@alcatel-lucent.com>
- o Jan Seedorf <jan.seedorf@nw.neclab.eu>

The authors would like to thank the members of the P2PI and ALTO mailing lists for their feedback.

Laird Popkin and Y. Richard Yang are grateful to the many contributions made by the members of the P4P working group and Yale Laboratory of Networked Systems. The P4P working group is hosted by DCIA.

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 (editor)
University of Stuttgart Computing Center
Networks and Communication Systems Department
Allmandring 30
70550 Stuttgart
Germany

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

Stefano Previdi
Cisco Systems, Inc.

Email: sprevidi@cisco.com

Martin Stiernerling
NEC Laboratories Europe

Email: martin.stiernerling@neclab.eu
URI: <http://ietf.stiernerling.org>

Richard Woundy
Comcast Corporation

Email: Richard_Woundy@cable.comcast.com

Yang Richard Yang
Yale University

Email: yry@cs.yale.edu

