Network Working Group
Internet-Draft
Intended status: Informational
Expires: April 28, 2011

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
October 25, 2010

Application-Layer Traffic Optimization (ALTO) Requirements
draft-ietf-alto-reqs-06.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 ALTO, which should be considered when specifying, assessing, or comparing protocols and implementations, and it solicits feedback and discussion.

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at http://www.ietf.org/ietf/1id-abstracts.txt.

The list of Internet-Draft Shadow Directories can be accessed at http://www.ietf.org/shadow.html.

This Internet-Draft will expire on April 28, 2011.

Copyright Notice

Table of Contents

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 operated by the network provider, 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, Overlay Network, Application protocol, Peer, P2P, Resource, Resource Identifier, Resource Provider, Resource Consumer, Resource Directory, Transport Address, ALTO Service, ALTO Server, ALTO Client, ALTO Client Protocol, ALTO Query, ALTO Reply, ALTO Transaction, Provisioning protocol, Inter ALTO-Server Protocol, Local Traffic, Peering Traffic, Transit Traffic.

Furthermore, the following additional terms will be used:

- Host Group Descriptor: Information used to describe the resource consumer which seeks ALTO guidance, or one or several candidate resource providers. 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. For a discussion, see Section 4.

- 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 the ALTO server and transmitted in the ALTO protocol. It may be evaluated according to the rating criteria.

- 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 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 of the abovementioned architecture:

- The ALTO client protocol, which is used for sending ALTO queries and ALTO replies between ALTO client and ALTO server.
- The discovery mechanism, which will be used by ALTO clients in order to find out where to send ALTO requests.
- The overall architecture, especially with respect to security and privacy issues.

Furthermore, this document describes the following data structures, which might be used in the ALTO client protocol:

- Host group descriptors, which are used to describe the location of a host in the network topology.
- Rating criteria, i. e., conditions that shall be evaluated in order to generate the ALTO guidance.

Requirements regarding other components are not considered in the current version of this document, but may be added later.

2.4. Sample Use Cases

The ALTO problem statement [RFC5693] presents a figure that gives a high-level overview of protocol interaction between ALTO elements. The following figures are somewhat more elaborated and extended versions of it, in order to give some non-normative examples of ALTO usage. It can also be seen that, in some use cases, some of the requirements presented in later sections are more relevant than in others.

Figure 1 shows an ALTO use case with a DHT-based P2P application. Using this distributed lookup mechanism, a peer can figure out which

other peers are candidate resource providers for a desired resource. Every peer software includes an ALTO client, in order to request and receive guidance on peer selection from the ALTO servers.

From an ALTO perspective this means that the ALTO servers will receive ALTO queries from a rather large number of different ALTO clients. The performance of many clients and their Internet connectivity may be rather limited and therefore, this puts certain restrictions on the amount of guiding data that can be sent to them. Furthermore, the privacy-sensitive IP addresses of the peers are visible to the (operators of the) ALTO servers, as these are also the source addresses of the ALTO query messages.

Figure 2 shows an ALTO use case with a P2P application that makes use of a centralized resource directory (in some specific P2P implementations called a "tracker"). In this scenario the ALTO servers receive queries only from few entities, i.e., the resource directories. As these resource directories must be powerful machines anyway, it may be reasonable to send large amounts of ALTO guidance data to them, which will be cached there. Furthermore, in this scenario it may be possible to hide the exact addresses of the peers from the ALTO server.

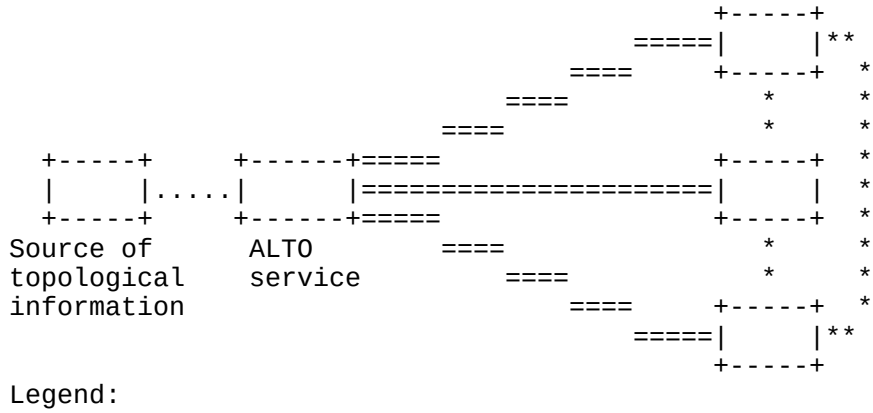

Legend:
=== ALTO client protocol
*** Application protocol
... Provisioning protocol

Figure 1: Overview of protocol interaction between ALTO elements, scenario without resource directory

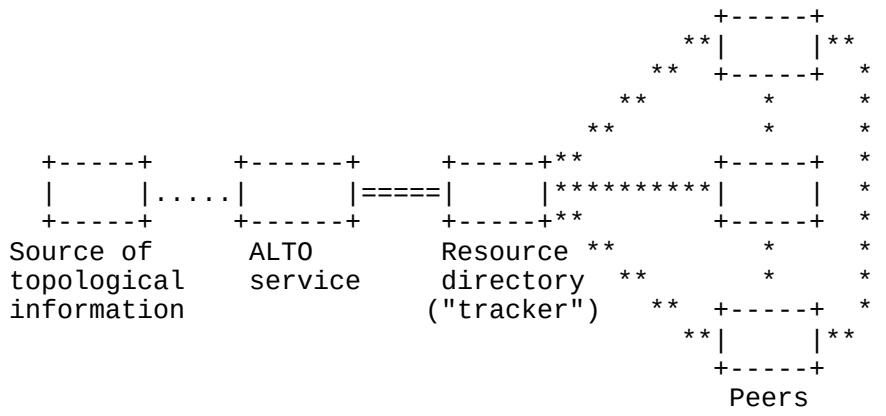

Figure 2: Overview of protocol interaction between ALTO elements, scenario with resource directory

3. ALTO Requirements

3.1. ALTO Client Protocol

3.1.1. General Requirements

REQ. ARv06-1: The ALTO service is provided by one or more ALTO servers. ALTO servers MUST implement the ALTO client protocol, for receiving ALTO queries from ALTO clients and for sending the corresponding ALTO replies.

REQ. ARv06-2: ALTO clients MUST implement the ALTO client protocol, for sending ALTO queries to ALTO servers and for receiving the corresponding ALTO replies.

REQ. ARv06-3: The format of the ALTO query message MUST allow the ALTO client to solicit guidance for selecting appropriate resource providers.

REQ. ARv06-4: The format of the ALTO reply message MUST allow the ALTO server to express its guidance for selecting appropriate resource providers.

REQ. ARv06-5: The detailed specification of a protocol is out of the scope of this document. However, any protocol specification that claims to implement the ALTO client protocol MUST be compliant to the requirements itemized in this document.

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 several rating criteria.

REQ. ARv06-6: The ALTO client protocol MUST support the usage of several different host group descriptor types.

REQ. ARv06-7: The ALTO client protocol specification MUST define a basic set of host group descriptor types, which MUST be supported by all implementations of the ALTO client protocol.

REQ. ARv06-8: The 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. It is also possible to specify a broader address range (i.e., a shorter prefix length) than the intended group of hosts

actually uses, in order to conceal their exact identity.

REQ. ARv06-9: The ALTO client protocol specification MUST define an appropriate procedure for adding new host group descriptor types, e.g., by establishing an IANA registry.

See Section 4 for a discussion of possible other host group descriptor types.

REQ. ARv06-10: ALTO clients and ALTO servers MUST clearly identify the type of each host group descriptor sent in ALTO queries or replies.

REQ. ARv06-11: 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. ARv06-12: Protocol functions for mapping other host group descriptor types to IPv4/IPv6 address prefixes SHOULD be designed and specified as part of the 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 the 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 downloaded as file via HTTP.

REQ. ARv06-13: The ALTO client protocol specification MUST define mechanisms, which can be used by the ALTO client and the ALTO server to indicate that a host group descriptor used by the other party is of an unsupported type, or that the indicated mapping mechanism could not be used.

3.1.3. Rating Criteria Support

REQ. ARv06-14: The ALTO client protocol MUST support the usage of several different rating criteria types.

REQ. ARv06-15: The ALTO client protocol specification MUST define a basic set of rating criteria types, which MUST be supported by all implementations of the ALTO client protocol.

REQ. ARv06-16: The ALTO client protocol specification MUST support the rating criteria type "relative operator’s preference." This is a relative measure, i.e., it is not associated with any unit of measurement. A higher rating according to this criterion indicates

that the application should prefer the respective candidate resource provider over others with lower ratings (if no other reasons speak against it, 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. ARv06-17: The ALTO client protocol specification MUST define an appropriate procedure for adding new rating criteria types, e.g., by establishing an IANA registry.

See Section 5 for a discussion of possible other rating criteria.

REQ. ARv06-18: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. ARv06-19:The ALTO reply message SHOULD allow the ALTO server to express which rating criteria have been considered when generating the reply.

REQ. ARv06-20: The 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:

- One mode of ALTO operation is that ALTO clients may be embedded directly in the resource consumer (e.g., peer of a DHT-based P2P application), which wants to access a resource.

- Another mode of operation is to perform ALTO queries indirectly, via resource directories (e.g., tracker of a P2P application), which may issue ALTO queries to solicit preference on potential resource providers, considering the respective resource consumer.

REQ. ARv06-21: The ALTO client protocol MUST support the mode of operation, in which the ALTO client is directly embedded in the resource consumer.

REQ. ARv06-22: The ALTO client protocol MUST support the mode of operation, in which the ALTO client is embedded in the resource

directory.

REQ. ARv06-23: The 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 IP access network.

REQ. ARv06-24: The 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:

- 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.

- 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. ARv06-25: The ALTO client protocol MUST support at least one of these two modes, either the target-aware or the target-independent query mode.

REQ. ARv06-26: The ALTO client protocol SHOULD support both the target-aware and the target-independent query mode.

REQ. ARv06-27: The ALTO client protocol SHOULD support lifetime attributes, to enable caching of recommendations at ALTO clients.

REQ. ARv06-28: The ALTO client protocol SHOULD specify an aging mechanism, which allows to give newer recommendations precedence over older ones.

REQ. ARv06-30: The ALTO client protocol SHOULD allow the ALTO server to add information about appropriate modes of re-use to its ALTO replies. Re-use may include redistributing an ALTO reply to other parties, as well as using the same ALTO information in a resource directory to improve the replies to different resource consumers, within the specified lifetime of the ALTO reply. The ALTO server SHOULD be able to express that

- no re-use should occur

- 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 reply, which is only a subset of the known actual "target audience", e.g., if required by operator policies

- 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

- 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

REQ. ARv06-31: The ALTO client protocol MUST support scenarios with the ALTO client 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. ARv06-32: The ALTO client protocol MUST include support for adding protocol extensions in a non-disruptive, backward-compatible way.

REQ. ARv06-33: The 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. ARv06-34: Any application designed to use ALTO MUST also work if no ALTO servers can be found or if no responses to ALTO queries are received, e.g., due to connectivity problems or overload situation.

REQ. ARv06-35: The ALTO client protocol MUST use TCP based transport.

REQ. ARv06-36: An ALTO server, which is operating close to its capacity limit, MUST be able to inform clients about its impending overload situation, and require them to throttle their query rate.

REQ. ARv06-37: An ALTO server, which is operating close to its capacity limit, MUST be able to inform clients about its impending overload situation, and redirect them to another ALTO server.

REQ. ARv06-38: An ALTO server, which is operating close to its capacity limit, MUST be able to inform clients about its impending overload situation, and terminate the conversation with the ALTO client.

REQ. ARv06-39: An ALTO server, which is operating close to its capacity limit, MUST be able to inform clients about its impending overload situation, and reject new conversation attempts.

3.2. ALTO Server Discovery

The ALTO client protocol is supported by one or several ALTO server discovery mechanisms, which will be used by ALTO clients in order to find out where to send ALTO requests.

REQ. ARv06-40: ALTO clients which are embedded in the resource consumer MUST be able to use the 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. ARv06-41: 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 the 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. ARv06-42: 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. ARv06-43: ALTO clients MUST be able to perform third-party ALTO server discovery, even if they are located behind a network address translator (NAT).

REQ. ARv06-44: 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. ARv06-45: The ALTO server discovery mechanism may be specified and provided using an existing protocol or mechanism, such as DNS, DHCP, or PPP based automatic configuration, etc. These candidate "base protocols" differ with respect to their availability in various access network architectures and their suitability for third-party queries. When evaluating different options this should be taken into account, in order to limit the total number of ALTO server discovery

mechanisms that have to be specified for supporting a reasonably wide range of deployment scenarios.

REQ. ARv06-46: The ALTO server discovery mechanism SHOULD be able to return the respective contact information for several ALTO servers.

REQ. ARv06-47: The ALTO server discovery mechanism SHOULD be able to indicate preferences for each returned ALTO server contact information.

3.3. Security and Privacy

REQ. ARv06-48: The ALTO client protocol MUST support mechanisms for the authentication of ALTO servers.

REQ. ARv06-49: The ALTO client protocol MUST support mechanisms for the authentication of ALTO clients.

REQ. ARv06-50: The ALTO client protocol MUST support different levels of detail in queries and responses, in order for the operator of an ALTO service to be able to control how much information (e.g., about the network topology) is disclosed.

REQ. ARv06-51: 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. ARv06-52: The 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. ARv06-53: The ALTO client protocol SHOULD be defined in a way, that the operator of one ALTO server cannot easily deduce the resource identifier (e.g., file name in P2P file sharing) which the resource consumer seeking ALTO guidance wants to access.

REQ. ARv06-54: The ALTO client protocol MUST include appropriate mechanisms to protect the ALTO service against DoS attacks.

4. Host Group Descriptors

Host group descriptors are used in the ALTO client protocol to describe the location of a host in the network topology. The ALTO client protocol specification defines a basic set of host group descriptor types, which have to be supported by all implementations, and an extension procedure for adding new descriptor types (see Section 3.1.2). The following list gives an overview on further host group descriptor types that have been proposed in the past, or which are in use by ALTO-related prototype implementations. This list is not intended as normative text. Instead, the only purpose of the following list is to document the descriptor types that have been proposed so far, and to solicit further feedback and discussion:

- Autonomous System (AS) number
- Protocol-specific group identifiers, which expand to a set of IP address ranges (CIDR) and/or AS numbers. In one specific solution proposal, these are called Partition ID (PID).

5. Rating Criteria

Rating criteria are used in the ALTO client protocol to express topology- or connectivity-related properties, which are evaluated in order to generate the ALTO guidance. The ALTO client protocol specification defines a basic set of rating criteria, which have to be supported by all implementations, and an extension procedure for adding new criteria (see Section 3.1.3). The following list gives an overview on further rating criteria that have been proposed in the past, or which are in use by ALTO-related prototype implementations. This list is not intended as normative text. Instead, the only purpose of the following list is to document the rating criteria that have been proposed so far, and to solicit further feedback and discussion:

5.1. Distance-related Rating Criteria

- Relative topological distance: relative means that a larger numerical value means greater distance, but it is up to the ALTO service how to compute the values, and the ALTO client will not be informed about the nature of the information. One way of generating this kind of information MAY be counting AS hops, but when querying this parameter, the ALTO client MUST NOT assume that the numbers actually are AS hops.

- Absolute topological distance, expressed in the number of traversed autonomous systems (AS).

- Absolute topological distance, expressed in the number of router hops (i.e., how much the TTL value of an IP packet will be decreased during transit).

- Absolute physical distance, based on knowledge of the approximate geolocation (continent, country) of an IP address.

5.2. Charging-related Rating Criteria

- Traffic volume caps, in case the Internet access of the resource consumer is not charged by "flat rate". For each candidate resource provider, the ALTO service could indicate the amount of data that may be transferred from/to this resource provider until a given point in time, and how much of this amount has already been consumed. Furthermore, it would have to be indicated how excess traffic would be handled (e.g., blocked, throttled, or charged separately at an indicated price). The interaction of several applications running on a host, out of which some use this criterion while others don’t, as well as the evaluation of this criterion in resource directories, which issue ALTO queries on

behalf of other peers, are for further study.

5.3. Performance-related Rating Criteria

The following rating criteria are subject to the remarks below.

- The minimum achievable throughput between the resource consumer and the candidate resource provider, which is considered useful by the application (only in ALTO queries), or
- An arbitrary upper bound for the throughput from/to the candidate resource provider (only in ALTO replies). This may be, but is not necessarily the provisioned access bandwidth of the candidate resource provider.
- The maximum round-trip time (RTT) between resource consumer and the candidate resource provider, which is acceptable for the application for useful communication with the candidate resource provider (only in ALTO queries), or
- An arbitrary lower bound for the RTT between resource consumer and the candidate resource provider (only in ALTO replies). This may be, for example, based on measurements of the propagation delay in a completely unloaded network.

The ALTO client MUST be aware, that with high probability, the actual performance values differ significantly from these upper and lower bounds. In particular, an ALTO client MUST NOT consider the "upper bound for throughput" parameter as a permission to send data at the indicated rate without using congestion control mechanisms.

The discrepancies are due to various reasons, including, but not limited to the facts that

- the ALTO service is not an admission control system
- the ALTO service may not know the instantaneous congestion status of the network
- the ALTO service may not know all link bandwidths, i.e., where the bottleneck really is, and there may be shared bottlenecks
- the ALTO service may not know whether the candidate peer itself is overloaded
- the ALTO service may not know whether the candidate peer throttles the bandwidth it devotes for the considered application

- the ALTO service may not know whether the candidate peer will throttle the data it sends to us (e.g., because of some fairness algorithm, such as tit-for-tat)

Because of these inaccuracies and the lack of complete, instantaneous state information, which are inherent to the ALTO service, the application must use other mechanisms (such as passive measurements on actual data transmissions) to assess the currently achievable throughput, and it MUST use appropriate congestion control mechanisms in order to avoid a congestion collapse. Nevertheless, these rating criteria may provide a useful shortcut for quickly excluding candidate resource providers from such probing, if it is known in advance that connectivity is in any case worse than what is considered the minimum useful value by the respective application.

5.4. Inappropriate Rating Criteria

Rating criteria that SHOULD NOT be defined for and used by the ALTO service include:

- Performance metrics that are closely related to the instantaneous congestion status. The definition of alternate approaches for congestion control is explicitly out of the scope of ALTO. Instead, other appropriate means, such as using TCP based transport, have to be used to avoid congestion.

6. 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.

7. Security Considerations

7.1. 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].

7.2. Classification of Information Disclosure Scenarios

The unwanted disclosure of information is one key concern related to ALTO. The following list gives a classification of information disclosure scenarios, which may be considered more or less critical by different parties:

- (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.

- (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.

- (3) Disclosure of ALTO server operator's data (e.g., network topology information) to an unauthorized third party. There are a couple of 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.

- (4) Disclosure of the application behavior to an unauthorized third party.

- (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 replies among each other (see also case 3c). By correlating the ALTO replies they could find out more information than intended to be disclosed by the ALTO server operator.

(1) may be addressed by the ALTO server operator choosing the level of detail of the information to be populated into the ALTO server. Furthermore, access control mechanisms for filtering ALTO replies 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) is addressed 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, (2) can be addressed by ALTO clients by obfuscating the identity of candidate resource consumers, e.g., by zeroing-out or randomizing the last few bits of the 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 won't 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 by the requirements presented in this document.

7.3. Security Requirements

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

8. References

8.1. Normative References

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

8.2. Informative References

[ALTO-charter] Marocco, E. and V. Gurbani, "Application-Layer Traffic Optimization (ALTO) Working Group Charter", February 2009.

[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

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

- Richard Alimi <richard.alimi@yale.edu>
- Zoran Despotovic <despotovic@docomolab-euro.com>
- Jason Livingood <Jason_Livingood@cable.comcast.com>
- Saverio Niccolini <saverio.niccolini@nw.neclab.eu>
- 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.

Appendix B. Acknowledgments

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

The authors would like to thank

- Vijay K. Gurbani <vkg@alcatel-lucent.com>

- Enrico Marocco <enrico.marocco@telecomitalia.it>

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

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.

Saverio Niccolini, Jan Seedorf, and Martin Stiemerling are partially supported by the NAPA-WINE project (Network-Aware P2P-TV Application over Wise Networks, http://www.napa-wine.org), a research project supported by the European Commission under its 7th Framework Program (contract no. 214412). 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 NAPA-WINE project or the European Commission.

Authors' Addresses

Sebastian Kiesel (editor)
University of Stuttgart Computing Center
Allmandring 30
Stuttgart 70550
Germany

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

Stefano Previdi
Cisco Systems, Inc.

Email: sprevidi@cisco.com

Martin Stiemerling
NEC Laboratories Europe/University of Goettingen

Email: martin.stiemerling@neclab.eu
URI: http://ietf.stiemerling.org

Richard Woundy
Comcast Corporation

Email: Richard_Woundy@cable.comcast.com

Yang Richard Yang
Yale University

Email: yry@cs.yale.edu