

ALTO WG
Internet-Draft
Intended status: Standards Track
Expires: April 29, 2010

R. Penno, Ed.
Juniper Networks
Y. Yang, Ed.
Yale University
October 26, 2009

ALTO Protocol
draft-penno-alto-protocol-04.txt

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 29, 2010.

Copyright Notice

Copyright (c) 2009 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 in effect on the date of publication of this document (<http://trustee.ietf.org/license-info>). Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

Abstract

Networking applications today already have access to a great amount of Inter-Provider network topology information. For example, views

of the Internet routing table are easily available at looking glass servers and entirely practical to be downloaded by clients. What is missing is knowledge of the underlying network topology from the ISP or Content Provider (henceforth referred as Provider) point of view. In other words, what an Provider prefers in terms of traffic optimization -- and a way to distribute it.

The ALTO Service provides information such as preferences of network resources with the goal of modifying network resource consumption patterns while maintaining or improving application performance. This document describes a protocol implementing the ALTO Service. While such service would primarily be provided by the network (i.e., the ISP), content providers and third parties could also operate this service. Applications that could use this service are those that have a choice in connection endpoints. Examples of such applications are peer-to-peer (P2P) and content delivery networks.

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](#) [1].

Table of Contents

1.	Introduction	5
1.1.	Background and Problem Statement	5
1.2.	Design History and Merged Proposals	5
1.3.	Solution Benefits	5
1.3.1.	Service Providers	6
1.3.2.	Applications	6
2.	Architecture	6
2.1.	Terminology	6
2.1.1.	Endpoint Address	6
2.1.2.	Network Location	7
2.2.	ALTO Service and Protocol Scope	7
3.	Protocol Structure	8
3.1.	Server Capability	9
3.2.	Services	9
3.2.1.	Map Service	9
3.2.2.	Map Filtering Service	10
3.2.3.	Endpoint Property Service	10
3.2.4.	Ranking Service	10
4.	Network Map	10
4.1.	PID	11
4.2.	Example Network Map	11
5.	Path Rating	12
5.1.	Path Cost	12
5.1.1.	Cost Type	12
5.1.2.	Cost Mode	12
5.2.	Path Rating Query	13
5.2.1.	Cost Map	13
5.2.2.	Ranking List	13
5.2.3.	Network Map and Cost Map Dependency	13
6.	Protocol Overview	14
6.1.	Design Approach	14
6.1.1.	Use of Existing Infrastructure	14
6.1.2.	ALTO Information Reuse and Redistribution	14
6.2.	Message Format	15
6.2.1.	Query Message	15
6.2.2.	Response Message	16
6.2.3.	Query and Response Body Encoding	16
7.	Protocol Messaging	16
7.1.	Client Processing	16
7.1.1.	General Processing	17
7.1.2.	General Error Conditions	17
7.2.	Server Processing	17
7.2.1.	Successful Responses	17
7.2.2.	General Error Conditions	17
7.2.3.	Caching Parameters	17
7.3.	ALTO Queries	18

7.3.1.	Server Capability	18
7.3.2.	Map Service	18
7.3.3.	Map Filtering Service	19
7.3.4.	Endpoint Property Service	21
7.3.5.	Ranking Service	22
8.	Use Cases	23
8.1.	ALTO Client Embedded in P2P Tracker	23
8.2.	ALTO Client Embedded in P2P Client: Numerical Costs	25
8.3.	ALTO Client Embedded in P2P Client: Ranking	26
9.	Discussions	27
9.1.	Discovery	27
9.2.	Network Address Translation Considerations	27
9.3.	Mapping IPs to ASNs	28
9.4.	Endpoint and Path Properties	28
9.5.	P2P Peer Selection	28
9.5.1.	Client-based Peer Selection	29
9.5.2.	Server-based Peer Selection	29
10.	IANA Considerations	29
11.	Security Considerations	29
11.1.	ISPs	29
11.2.	ALTO Clients	29
11.3.	ALTO Information	30
11.4.	ALTO Information Redistribution	30
11.5.	Denial of Service	31
12.	References	31
12.1.	Normative References	31
12.2.	Informative References	31
Appendix A.	Contributors	33
Appendix B.	Acknowledgements	35
Authors' Addresses	35

1. Introduction

1.1. Background and Problem Statement

Today, network information available to applications is mostly from the view of endhosts. There is no clear mechanism to convey information about the network's preferences to applications. By leveraging better network-provided information, applications have potential to become more network-efficient (e.g., reduce network resource consumption) and achieve better application performance (e.g., accelerated download rate). The ALTO Service intends to provide a simple way to convey network information to applications.

The goal of the protocol specified in this document is to provide a simple, unified protocol that meets the ALTO requirements [5], providing a migration path for Internet Service Providers (ISP), Content Providers, and clients that have deployed protocols with similar intentions (see below). This document is a work in progress and will be updated with further developments.

1.2. Design History and Merged Proposals

The protocol specified here consists of contributions from

- o P4P [6],[7];
- o ALTO Info-Export [8];
- o Query/Response [9],[10];
- o ATTP [ATTP].
- o Proxidor [19].

The people listed here should be viewed as co-authors of this document: Obi Akonjang, Richard Alimi, Saumitra M. Das, Syon Ding, Anja Feldmann, Doug Pasko, Reinaldo Penno, Laird Popkin, Stefano Previdi, Satish Raghunath, Stanislav Shalunov, Albert Tian, Yu-Shun Wang, Richard Woundy, Y. Richard Yang, David Zhang, and Yunfei Zhang. Due to the limit of 5 authors per draft, the complete list of authors were moved to the contributors section at this point.

1.3. Solution Benefits

The ALTO Service offers many benefits to both end-users (consumers of the service) and Internet Service Providers (providers of the service).

1.3.1. Service Providers

The ALTO Service enables ISPs to influence the neighborhood selection process of overlay networks to increase locality of traffic and also regain the ability to efficiently engineer traffic that traverses more expensive links such as backbone and transit links, thus allowing a better provisioning of the networking infrastructure.

1.3.2. Applications

Applications that use the ALTO Service can benefit in multiple ways. For example, they may no longer need to infer topology information, and some applications can reduce reliance on measuring path performance metrics themselves. They can take advantage of the ISP's knowledge to avoid bottlenecks and boost performance.

An example type of application is a Peer-to-Peer overlay where peer selection can be improved by including ALTO information in the selection process.

2. Architecture

Two key design objectives of the ALTO Protocol are simplicity and extensibility. At the same time, it introduces additional techniques to address potential scalability and privacy issues. Below we start with an introduction to the terminology. Then we define the overall architecture and how the ALTO Protocol fits into the architecture.

2.1. Terminology

We use the following terms defined in [11]: Application, Overlay Network, Peer, Resource, Resource Identifier, Resource Provider, Resource Consumer, Resource Directory, Transport Address, Host Location Attribute, ALTO Service, ALTO Server, ALTO Client, ALTO Query, ALTO Reply, ALTO Transaction, Local Traffic, Peering Traffic, Transit Traffic.

We also use the following additional terms: Endpoint Address and Network Location.

2.1.1. Endpoint Address

An endpoint address represents the communication address of an endpoint. An endpoint address can be network-attachment based (IP address) or network-attachment agnostic. Common forms of endpoint addresses include IP address, MAC address, overlay ID, and phone number.

2.1.2. Network Location

Network Location is a generic concept denoting a single endpoint or group of endpoints. Whenever we say Network Location, we refer to either a single endpoint or a group of endpoints.

2.2. ALTO Service and Protocol Scope

An ALTO Server conveys the network information from the perspective of a network region. We say that the ALTO Server presents its "my-Internet View" [[12](#)] of the network region. A network region in this context can be an Autonomous System, an ISP, perhaps a smaller region, or perhaps a set of ISPs; the details depend on the deployment scenario and discovery mechanism.

To better understand the ALTO Service and the role of the ALTO Protocol, we show in Figure 1 the overall system architecture. In this architecture, an ALTO Client uses ALTO Service Discovery to identify an appropriate ALTO Server; an ALTO Server prepares ALTO Information; and the ALTO Client requests available ALTO Information from the ALTO Server using the ALTO Protocol.

The ALTO Information provided by the ALTO Server can be updated dynamically based on network conditions, or can be seen as a policy which is updated at a larger time-scale.

More specifically, the ALTO Information provided by an ALTO Server may be influenced (at the operator's discretion) by other systems. Examples include (but are not limited to) static network configuration databases, dynamic network information, routing protocols, provisioning policies, and interfaces to outside parties. These components are shown in the figure for completeness but outside the scope of this specification.

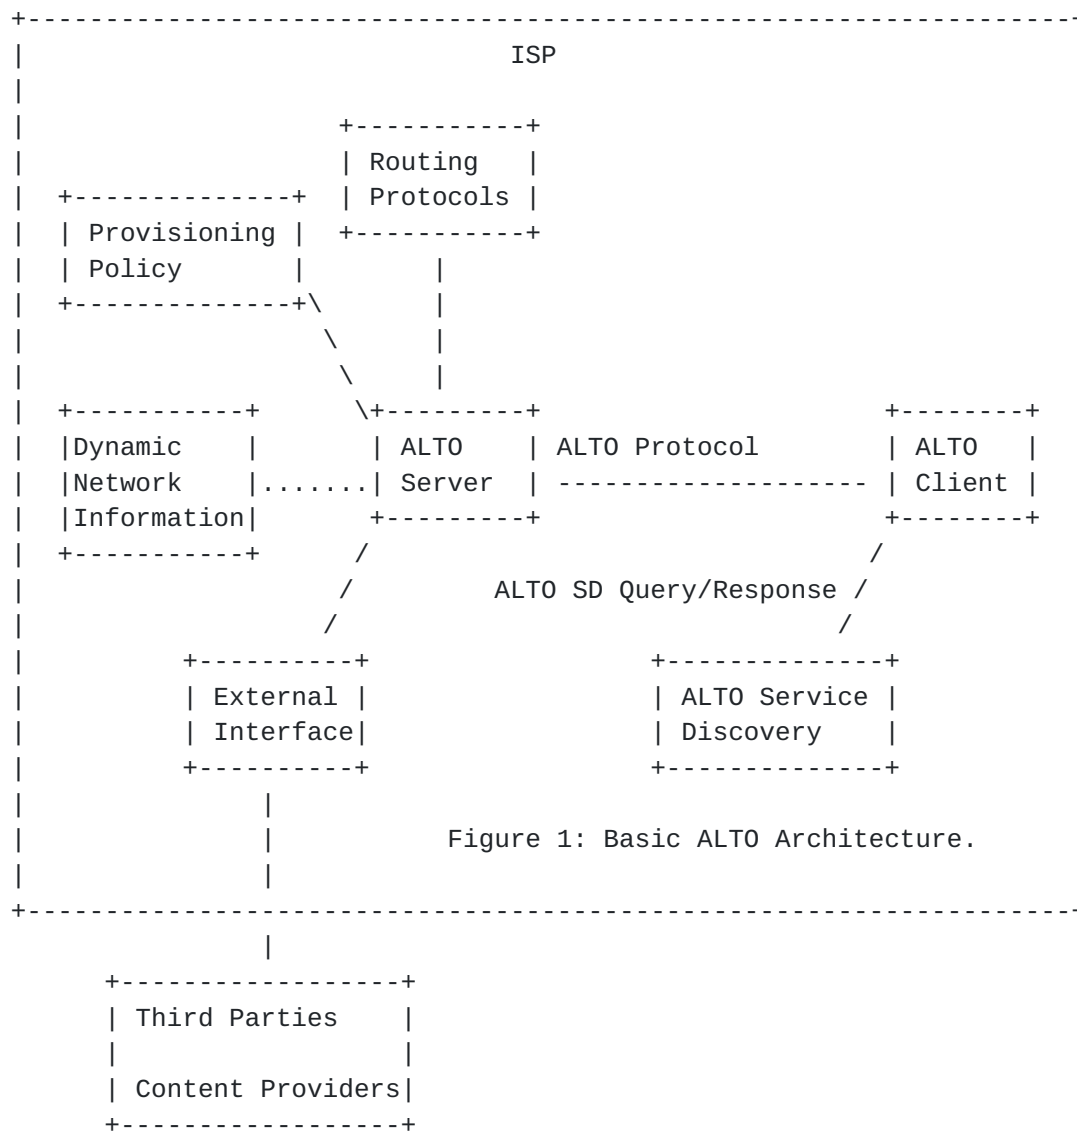


Figure 1: Basic ALTO Architecture.

ALTO Architecture

3. Protocol Structure

The ALTO Protocol uses a simple extensible framework to convey network information. In the general framework, the ALTO protocol will convey properties on both Endpoints and paths between network locations.

In this document, we focus on a particular endpoint property to denote the location of an endpoint and a particular path property called Path Rating to denote the ISP-defined cost of a path.

The ALTO Protocol is built on a common transport protocol, messaging

structure and encoding, and transaction model. The protocol is subdivided into services of related functionality. ALTO-Core provides the Map Service. Other services can provide additional functionality. There are three such services defined in this document: the Map Filtering Service, Endpoint Property Service, and Ranking Service. Additional services may be defined in the future in companion documents.

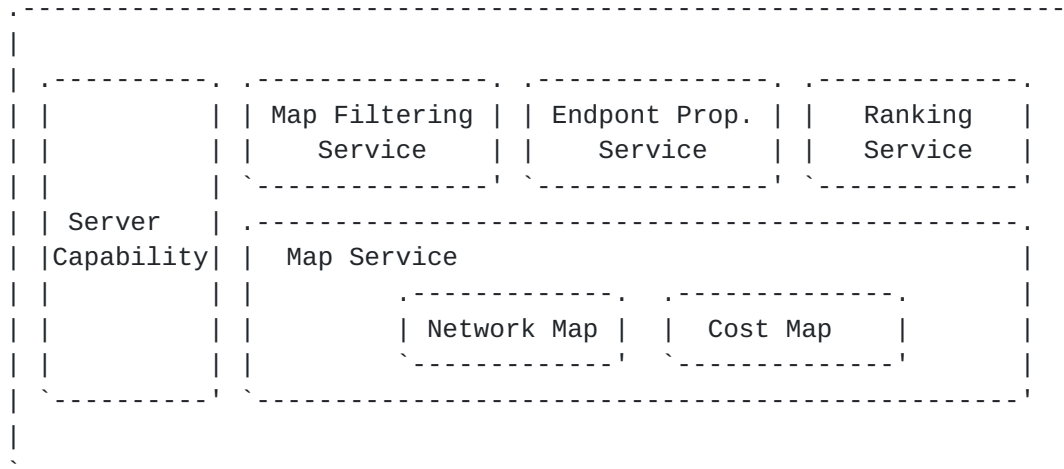


Figure 1: ALTO Protocol Structure

[3.1.](#) Server Capability

It lists the details on the information that can be provided by an ALTO Server. The configuration includes, for example, details about the operations and cost metrics supported by the ALTO Server. The capability document can be downloaded by ALTO Clients or provisioned into devices.

[3.2.](#) Services

[3.2.1.](#) Map Service

The Map Service provides batch information to ALTO Clients. Two maps are provided in this document. The Network Map (See [Section 4](#)) provides the full set of network location groupings defined by the ALTO Server and the endpoints contained with each grouping. The Cost Map (see [Section 5.2.1](#)) provides costs between the defined groupings.

These two maps can be thought of (and implemented as) as simple files with appropriate encoding provided by the ALTO Server.

3.2.2. Map Filtering Service

Resource constrained ALTO Clients may benefit from query results being filtered at the ALTO Server. This avoids an ALTO Client spending network bandwidth collecting results and performing client-side filtering. The Map Filtering Service allows ALTO Clients to query for ALTO Server maps based on additional parameters.

3.2.3. Endpoint Property Service

This service allows ALTO Clients to look up properties for individual endpoints. An example endpoint property is its network location (its grouping defined by the ALTO Server) or connectivity type (e.g., ADSL, Cable, or FioS).

3.2.4. Ranking Service

Some ALTO Clients may also benefit from querying for rankings and costs based on endpoints. The Ranking Service allows an ALTO Server to return either numerical costs or ordinal costs (rankings) for additional network locations types such as Endpoints.

4. Network Map

In this section, we give more detail on the particular endpoint property named PID. In the next section, we give more detail about the particular path property named Path Rating.

In reality many endpoints are very close to one another in terms of network connectivity, for example, endpoints on the same site of an enterprise. By treating a group of endpoints together as a single entity in ALTO, we can achieve much greater scalability without loosing any critical information.

The Network Location endpoint property allows an ALTO Server to group endpoints together to indicate their proximity. The resulting set of groupings is called the ALTO Network Map.

The Network Map may also be used to communicate simple preferences. For example, an ISP may prefer that endpoints associated with the same PoP (Point-of-Presence) in a P2P application communicate locally instead of communicating with endpoints in other PoPs.

Note that the definition of proximity varies depending on the granularity of the ALTO algorithm. In one deployment, endpoints on the same subnet may be considered close; while in another deployment, endpoints connected to the same PoP may be considered close.

4.1. PID

Each group can be identified by a provider-defined Network Location identifier called a PID. There can be many different ways of grouping the endpoints and assigning PIDs.

Thus, a PID is a identifier providing an indirect and network-agnostic way to specify a network aggregation. For example, a PID may be defined (by the ALTO service provider) to denote a subnet, a set of subnets, a metropolitan area, a PoP, an autonomous system, or a set of autonomous systems. Aggregation of endpoints into PIDs can indicate proximity. Also, aggregation can improve scalability. In particular, network preferences (costs) may be specified between PIDs, allowing cost information to be more compact and updated at a smaller time scale than the network aggregations themselves.

4.2. Example Network Map

Figure 2 illustrates an example Network Map. PIDs are used to identify network-agnostic aggregations.

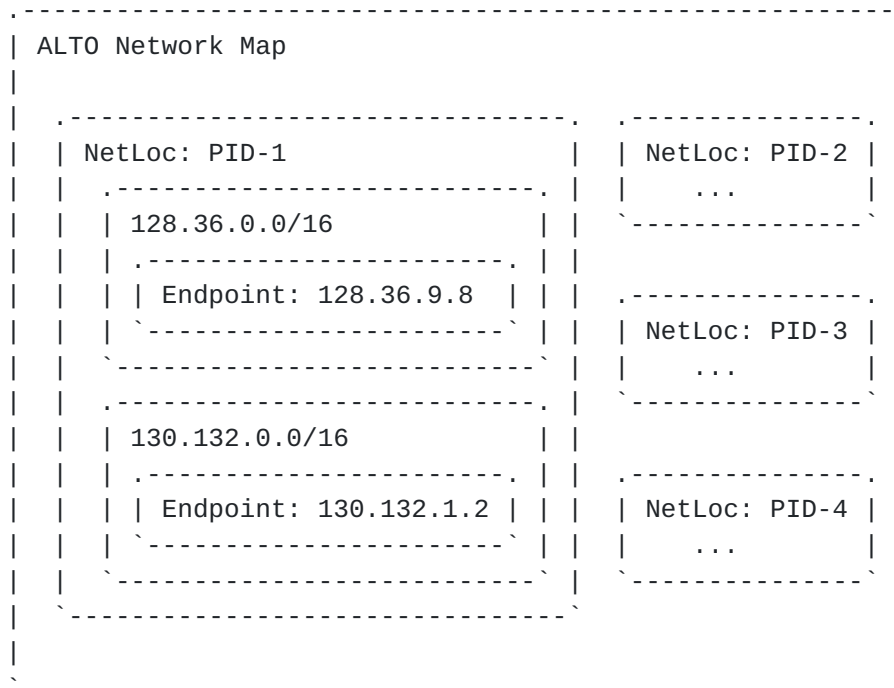


Figure 2: Example Network Map

5. Path Rating

In this section we define a particular path property named Path Rating.

5.1. Path Cost

Path Rating is based on Path Cost, which conveys the preference of an ALTO Server on communication among Network Locations. Path Costs have attributes:

- o Type: identifies what the costs represent;
- o Mode: identifies how the costs should be interpreted (numerical or ordinal interpretation).

5.1.1. Cost Type

The Type attribute indicates what the cost represents. For example, an ALTO Server could define costs representing air-miles, hop-counts, or generic routing costs.

Cost types are indicated in protocol messages as alphanumeric strings. An ALTO Server **MUST** at least define the routing cost type denoted by the string 'routingcost'.

Note that an ISP may internally compute routing cost using any method it chooses (including air-miles or hop-count).

If an ALTO Client requests a Cost Type that is not available, the ALTO Server responds with an error as specified in [Section 7.2.2.2](#).

5.1.2. Cost Mode

The Mode attribute indicates how costs should be interpreted. For example, an ALTO Server could return costs that should be interpreted as numerical values or ordinal rankings.

It is important to communicate such information to ALTO Clients, as certain operations may not be valid on certain costs returned by an ALTO Server. For example, it is possible for an ALTO Server to return a set of IP addresses with costs indicating a ranking of the IP addresses. Arithmetic operations, such as summation, that would make sense for numerical values, do not make sense for ordinal rankings. ALTO Clients may want to handle such costs differently.

Cost Modes are indicated in protocol messages as alphanumeric strings. An ALTO Server **MUST** at least define the modes 'numerical'

and 'ordinal'.

If an ALTO Client requests a cost Mode that is not supported, the ALTO Server MUST reply with costs having Mode either 'numerical' or 'ordinal'. Thus, an ALTO Server must implement at least one of 'numerical' or 'ordinal' Costs, but it may choose which to support. ALTO Clients may choose how to handle such situations. Two particular possibilities are to use the returned costs as-is (e.g., treat numerical costs as ordinal rankings) or ignore the ALTO information altogether.

5.2. Path Rating Query

The Path Rating Query consists of a Cost Type, a Cost Mode, a list of Source Network Locations (note that a Network Location can be an endpoint address or a PID), and a list of Destination Network Locations.

Specifically, assume that a Path Rating query has a list of multiple Source Network Locations, say [Src_1, Src_2, ..., Src_m], and a list of multiple Destination Network Locations, say [Dst_1, Dst_2, ..., Dst_n], then the ALTO Server will compute the Path Cost for each communicating pair (i.e., Src_1 -> Dst_1, ..., Src_1 -> Dst_n, ..., Src_m -> Dst_1, ..., Src_m -> Dst_n).

5.2.1. Cost Map

We refer to the Response containing the $m \times n$ entries as a Cost Map.

If the Cost Type is ordinal, the ranking of each communicating pair is relative to the $m \times n$ entries.

5.2.2. Ranking List

If there is a single Source Network Location, we also say that the response is a Ranking List.

5.2.3. Network Map and Cost Map Dependency

Note that if a Path Rating query contains any PID in the list of Source Network Locations or the list of Destination Network Locations, we say that the particular Path Rating is generated based on a particular Network Map. Version Tags are introduced to ensure that ALTO Clients are able to use consistent information even though the information is provided in two maps. One advantage of separating ALTO information into a Network Map and a Cost Map is that the two components can be updated at different time scales. For example, Network Maps may be stable for a longer time while Cost Maps may be

updated to reflect dynamic network conditions.

6. Protocol Overview

6.1. Design Approach

The ALTO Protocol design uses a RESTful interface with the goal of leveraging current HTTP [2] [3] implementations and infrastructure. ALTO messages are denoted with HTTP Content-Type "application/alto". Message encodings use a structured, text-based format. The exact encoding will be documented in a future revision, as the current focus is overall protocol architecture and operations.

These design decisions make the protocol easier to understand and debug, and allows for flexible ALTO Server implementation strategies. More importantly, however, this enables use of existing implementations and infrastructure, and allows for simple caching and redistribution of ALTO information to increase scalability.

6.1.1. Use of Existing Infrastructure

An important design consideration for the ALTO Protocol is easy integration with existing applications and infrastructure. As outlined above, HTTP is a natural choice. In particular, this ALTO Protocol design leverages:

- o the huge installed base of infrastructure, including HTTP caches,
- o mature software implementations,
- o the fact that many P2P clients already have an embedded HTTP client, and
- o authentication and encryption mechanisms in HTTP and SSL.

6.1.2. ALTO Information Reuse and Redistribution

ALTO information may be useful to a large number of applications and users. Distributing ALTO information must be efficient and not become a bottleneck. Therefore, the ALTO Protocol specified in this document integrates with existing HTTP caching infrastructure to allow reuse of ALTO information by ALTO Clients and reduce load on ALTO servers. ALTO information may also be cached or redistributed using application-dependent mechanisms, such as P2P DHTs or P2P file-sharing.

For example, a full Network Map may be reused by all ALTO Clients

using the ALTO Server.

6.2. Message Format

The ALTO Protocol uses a RESTful design operating over HTTP. Both Query and Response follow the standard format for HTTP Request and Response messages [2] [3]. This section provides an overview of the components of a Query message sent from an ALTO Client to an ALTO Server, as well as the components of a Response message returned by an ALTO Server. Note that if caching or redistribution is used, the Response message may be returned from another (possibly third-party) entity. Reuse and Redistribution is further discussed in [Section 11.4](#).

6.2.1. Query Message

A Query message is generated by an ALTO Client and sent to an ALTO Server. The ALTO Protocol employs the following components of the HTTP request message:

Method: Indicates operation requested by the ALTO Client (along with URI Path).

URI Path: Indicates the operation requested by the ALTO Client (along with Method).

URI Query String Parameters: Indicates parameters for the requested operation. Note that in the messaging specification in [Section 7](#), we abbreviate these as 'URI QS Params'. Order of query string parameters is not specified. Some parameters are restricted in how many times they appear. We use the notation 'min..max' to denote the minimum and maximum times they may appear, where 'max' may be '*' to denote unbounded. If no parameters are defined for a particular ALTO request message, the query string MUST be empty (and there MUST be no '?' included in the URI Path).

Headers: Indicates encoding of the Body. If the HTTP body of the request is non-empty, the Content-Type header MUST be set to "application/alto".

Body: Indicates additional request parameters that are not concisely representable as Query String parameters. If no Body is defined for a particular ALTO request message, the HTTP request body MUST be empty.

6.2.2. Response Message

A Response message is generated by an ALTO Server, which corresponds to a particular Query message. The ALTO Protocol employs the following components of the HTTP Response message:

Status Code: Indicates either success or an error condition.

Headers: Indicates encoding of the Body and caching directives. If the HTTP body of the response is non-empty, the Content-Type header MUST be set to "application/alto"

Body: Contains data requested by the ALTO Client.

6.2.3. Query and Response Body Encoding

When the Body of a Query or Response message is not empty, it MUST contain a properly-encoded message. This section will be updated to include common requirements when an encoding format is decided.

7. Protocol Messaging

This section specifies client and server processing, as well as messages in the ALTO Protocol. Details common to ALTO Server processing of all messages is first discussed, followed by details of the individual messages.

Note that the primary focus of the current draft is the architecture and protocol operations, and only the structure of messages is specified (without actual message encoding). Additionally, the following details have been omitted for clarity:

- o protocol version number
- o transaction ID
- o map version tags
- o HTTP URL encoding

This section will be updated as revisions are made to protocol details and encodings.

7.1. Client Processing

[7.1.1.](#) General Processing

The protocol is structured in such a way that independent of the query type there are a set of general processing steps. The ALTO Client selects a specific ALTO Server to communicate with and establishes a TCP connection. The ALTO protocol on top of this TCP connection MAY be secured through SSL/TLS to implement server and/or client authentication. HTTP Basic or Digest authentication MAY additionally be used. The client then creates and sends a query message, which MUST be constructed as specified in [Section 7.3](#). The ALTO Client MUST additionally follow any HTTP encoding rules as well as TCP transport considerations.

[7.1.2.](#) General Error Conditions

In the case the client does not receive an appropriate response from the server it can choose another server to communicate or fall back to perform peer selection without the use of ALTO information.

[7.2.](#) Server Processing

[7.2.1.](#) Successful Responses

If a Query message is successfully processed an an ALTO response is generated, the HTTP status code in the response MUST be set to 200.

[7.2.2.](#) General Error Conditions

This section specifies ALTO Server behavior when it receives a Query message that cannot be processed due to a problem with processing the Query message itself.

[7.2.2.1.](#) Invalid Query Format

If any component of the Query message is formatted incorrectly (i.e., it does not follow the formats in [Section 7.3](#)), the ALTO Server MUST return HTTP Status Code 400.

[7.2.2.2.](#) Unsupported Query

If an ALTO Server does not support the operation indicated in the Query message, the ALTO Server MUST return HTTP Status Code 501.

[7.2.3.](#) Caching Parameters

If the response generated by the ALTO Server is cachable, the ALTO Server MAY include relevant HTTP headers, enabling it to be cached by existing HTTP Caches or the ALTO Client itself.

[7.3.](#) ALTO Queries

[7.3.1.](#) Server Capability

The Server Capability query allows an ALTO Client to determine the available operations of a particular ALTO Server

This query **MUST** be supported.

[7.3.1.1.](#) Query

Method : 'GET'
URI Path : '/capability'

[7.3.1.2.](#) Example Response Structure

```
{
  "instance-name": "alto.example.com",
  "uri": "http:\\\\alto.example.com:6671",
  "cost": [
    {"type": "latency", "units": "ms"},
    {"type": "pDistance", "units": "scalar"},
    {"type": "bandwidth", "units": "kbps"}
  ],
  "constraint-support": false
}
```

[7.3.2.](#) Map Service

The Map Service provides batch information to ALTO Clients in the form of two maps: a Network Map and Cost Map. All queries in the Map Service **MUST** be supported.

[7.3.2.1.](#) Network Map

The full Network Map is an instantiation of the Reverse Property Lookup where the ALTO Server lists for each PID, the network locations (endpoints) within the PID.

[7.3.2.1.1.](#) Query

Method : 'GET'
URI Path : '/prop/pid/map'

[7.3.2.1.2.](#) Example Response Structure

```
{
  "PID1" : [
    "128.36.1.0/24",
    "132.130.1.0/24",
    "132.130.2.0/24"
  ],
  "PID2" : [ "130.132.3.0/24" ],
  "PID3" : [ "0.0.0.0/0" ]
}
```

[7.3.2.2.](#) Cost Map

The full Cost Map is an instantiation of the Path Rating Query where the ALTO Server lists for each pair of source/destination PID, the Path Cost from the source to destination.

ALTO Server provides costs using the default Cost Type ('routingcost') and default Cost Mode ('numerical').

[7.3.2.2.1.](#) Query

```
Method      : 'GET'
URI Path    : '/cost/pid/map'
```

[7.3.2.2.2.](#) Example Response Structure

```
{
  "Type": "routingcost",
  "Mode": "numerical",
  "Map" : {
    "PID1": { "PID1": 1, "PID2": 5 , "PID3": 10 },
    "PID2": { "PID1": 5 , "PID2": 1 , "PID3": 15 },
    "PID3": { "PID1": 20, "PID2": 15, "PID3": 1  }
  }
}
```

[7.3.3.](#) Map Filtering Service

The Map Filtering Service allows ALTO Clients to specify filtering criteria to return a subset of the full maps available in the Map Service.

The services defined in this section are OPTIONAL.

7.3.3.1. Network Map

ALTO Clients can query for a subset of the full network map (see [Section 7.3.2.1](#)).

7.3.3.1.1. Query

Method : 'POST'
URI Path : '/prop/pid/filter'

7.3.3.1.2. Example Request Structure

POST /prop/pid/filter ...

["PID1", "PID2"]

7.3.3.1.3. Example Response Structure

```
{
  "PID1" : [
    "128.36.1.0/24",
    "132.130.1.0/24",
    "132.130.2.0/24"
  ],
  "PID2" : [ "130.132.3.0/24" ],
  "PID3" : [ "0.0.0.0/0" ]
}
```

7.3.3.2. Cost Map

ALTO Clients can query for the Cost Map (see [Section 7.3.2.2](#)) based on additional parameters.

7.3.3.2.1. Query

Method : 'POST'
URI Path : '/cost/pid/filter'
URI QS Params : 'type=[costtype]' (multiplicity: 0..1)
 'mode=[costmode]' (multiplicity: 0..1)
 'constraint=[constraint]' (multiplicity: 0..*)

The 'constraint' parameter is optional and is to be used only if the ALTO service supports it. It allows a client to specify a target numerical cost. The constraint contains two entities: (1) an operator either 'gt' for greater than, 'lt' for less than or 'eq' for equal to with 10 percent on either side, (2) a target numerical cost. The numerical cost is a number that **MUST** be defined in the units specified in the ALTO service configuration document obtained

from ALTO service discovery. If multiple 'constraint' parameters are specified, the ALTO Server assumes they are related to each other with a logical AND.

If the query does not specify the 'type' and 'mode' query string parameters, then the server assumes the type to be 'routingcost' and the mode to be 'numerical'. A Query MUST contain no more than one 'type' parameter, and no more than one 'mode' parameter.

The request body MAY specify a list of source PIDs, and a list of destination PIDs. If a list is empty, it is interpreted by the ALTO Server as the full set of PIDs. The ALTO Server returns costs between each pair of source/destination PID.

7.3.3.2.2. Example Request Structure

```
POST /cost/pid/filter

{
  "src": [ "PID1" ],
  "dst": [ "PID1", "PID2", "PID3" ]
}
```

7.3.3.2.3. Example Response Structure

```
{
  "Type": "routingcost",
  "Mode": "numerical",
  "Map" : {
    "PID1": { "PID1" : 1, "PID2": 5 , "PID3": 10 },
  }
}
```

7.3.4. Endpoint Property Service

The Endpoint Property Lookup query allows an ALTO Client to query for properties of Endpoints known to the ALTO Server. If the ALTO Server provides the Endpoint Property Service, the ALTO Server MUST define at least the 'pid' property for Endpoints. Additional supported properties can be defined in the Server Capability response.

The services defined in this section are OPTIONAL.

7.3.4.1. Query

```
Method       : 'POST'
URI Path     : '/endpoint/m'
URI QS Params : 'prop=[propertyname]' (multiplicity: 1..*)
```


The request body includes a list of Endpoints for which the property value should be returned. If the request body is empty, the ALTO Server implicitly assumes that the request contains a single-element list with the Endpoint address of the requesting client.

Also note that the 'prop' parameter may be specified multiple times to query for multiple properties simultaneously. For example, the query string could be 'prop=pid&prop=bandwidth'.

7.3.4.2. Example Request Structure

```
POST /endpoint/m?prop=pid ...  
  
[ "ipv4:128.36.1.34" ]
```

7.3.4.3. Example Response Structure

```
{  
  "ipv4:128.36.1.34" : { "pid": "PID1" }  
}
```

7.3.5. Ranking Service

The Ranking Service allows ALTO Clients to supply lists of endpoints to an ALTO Server. The ALTO Server replies with costs (numerical or ordinal) amongst the endpoints.

In particular, this service allows lists of Endpoint addresses to be ranked (ordered) by an ALTO Server.

The services defined in this section are OPTIONAL.

7.3.5.1. Ranking Query

7.3.5.1.1. Query

```
Method       : 'POST'  
URI Path     : '/cost/endpoint/ranking'  
URI QS Params : 'type=[costtype]'          (multiplicity: 0..1)  
               'mode=[costmode]'          (multiplicity: 0..1)  
               'constraint=[constraint]'    (multiplicity: 0..*)
```

The request body includes a list of source and destination endpoints that should be assigned a cost by the ALTO Server. The 'type', 'mode', and 'constraint' parameters behave as specified in [Section 7.3.3.2](#).

The request body MUST specify a list of source Endpoints, and a list

of destination Endpoints. If the list of source Endpoints is empty (or it is not included), the ALTO Server MUST treat it as if it contained the Endpoint address of the requesting client. The list of destination Endpoints MUST NOT be empty. The ALTO Server returns costs between each pair of source/destination Endpoint.

7.3.5.1.2. Example Request Structure

```
POST /cost/endpoint/ranking?mode=ordinal ...
```

```
{
  "src": "ipv4:128.30.24.2"
  "dst": [
    "ipv4:128.30.24.89",
    "ipv4:12.32.67.3",
    "ipv4:130.132.33.4"
  ]
}
```

7.3.5.1.3. Example Response Structure

```
{
  "Type": "routingcost",
  "Mode": "ordinal",
  "Ranking" : {
    "ipv4:128.30.24.2": {
      "ipv4:128.30.24.89" : 1,
      "ipv4:130.132.33.4" : 2,
      "ipv4:12.32.67.3"  : 3
    }
  }
}
```

8. Use Cases

The sections below depict typical use cases.

8.1. ALTO Client Embedded in P2P Tracker

Many P2P currently-deployed P2P systems use a Tracker to manage swarms and perform peer selection. P2P trackers may currently use a variety of information to perform peer selection to meet application-specific goals. By acting as an ALTO Client, an P2P tracker can use ALTO information as an additional information source to enable more network-efficient traffic patterns and improve application performance.

A particular requirement of many P2P trackers is that they must handle a large number of P2P clients. A P2P tracker can obtain and locally store ALTO information (the Network Map and Cost Map) from the ISPs containing the P2P clients, and benefit from the same aggregation of network locations done by ALTO Servers.

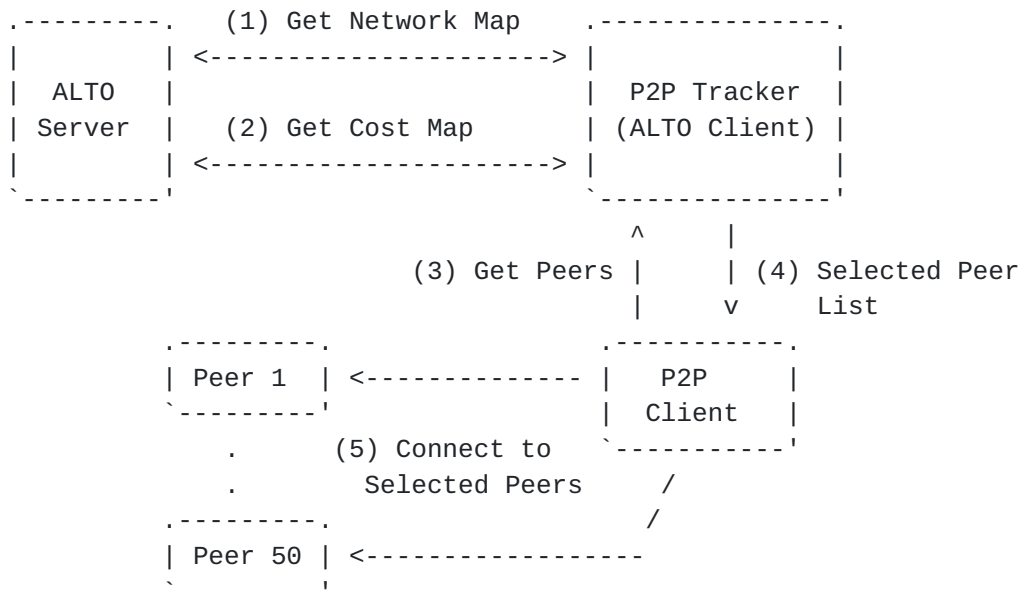


Figure 3: ALTO Client Embedded in P2P Tracker

Figure 3 shows an example use case where a P2P tracker is an ALTO Client and applies ALTO information when selecting peers for its P2P clients. The example proceeds as follows:

1. The P2P Tracker requests the Network Map covering all PIDs from the ALTO Server using the Reverse Property Lookup query. The Network Map includes the IP prefixes contained in each PID, allowing the P2P tracker to locally map P2P clients into a PIDs.
2. The P2P Tracker requests the Cost Map amongst all PIDs from the ALTO Server.
3. A P2P Client joins the swarm, and requests a peer list from the P2P Tracker.
4. The P2P Tracker returns a peer list to the P2P client. The returned peer list is computed based on the Network Map and Cost Map returned by the ALTO Server, and possibly other information sources. Note that it is possible that a tracker may use only the Network Map to implement hierarchical peer selection by preferring peers within the same PID and ISP.

5. The P2P Client connects to the selected peers.

Note that the P2P tracker may provide peer lists to P2P clients distributed across multiple ISPs. In such a case, the P2P tracker may communicate with multiple ALTO Servers.

8.2. ALTO Client Embedded in P2P Client: Numerical Costs

P2P clients may also utilize ALTO information themselves when selecting from available peers. It is important to note that not all P2P systems use a P2P tracker for peer discovery and selection. Furthermore, even when a P2P tracker is used, the P2P clients may rely on other sources, such as peer exchange and DHTs, to discover peers.

When an P2P Client uses ALTO information, it typically queries only the ALTO Server servicing its own ISP. The my-Internet view provided by its ISP's ALTO Server can include preferences to all potential peers.

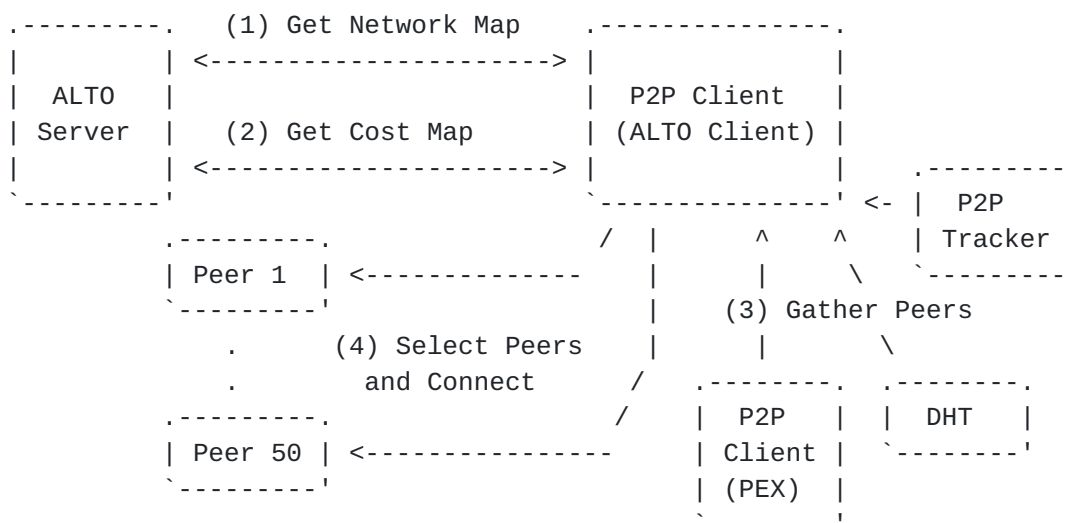


Figure 4: ALTO Client Embedded in P2P Client

Figure 4 shows an example use case where a P2P Client locally applies ALTO information to select peers. The use case proceeds as follows:

1. The P2P Client requests the Network Map covering all PIDs from the ALTO Server servicing its own ISP.
2. The P2P Client requests the Cost Map amongst all PIDs from the ALTO Server. The Cost Map by default specifies numerical costs.

3. The P2P Client discovers peers from sources such as Peer Exchange (PEX) from other P2P Clients, Distributed Hash Tables (DHT), and P2P Trackers.
4. The P2P Client uses ALTO information as part of the algorithm for selecting new peers, and connects to the selected peers.

8.3. ALTO Client Embedded in P2P Client: Ranking

It is also possible for a P2P Client to offload the selection and ranking process to an ALTO Server. In this use case, the ALTO Client gathers a list of known peers in the swarm, and asks the ALTO Server to rank them.

As in the use case using numerical costs, the P2P Client typically only queries the ALTO Server servicing its own ISP.

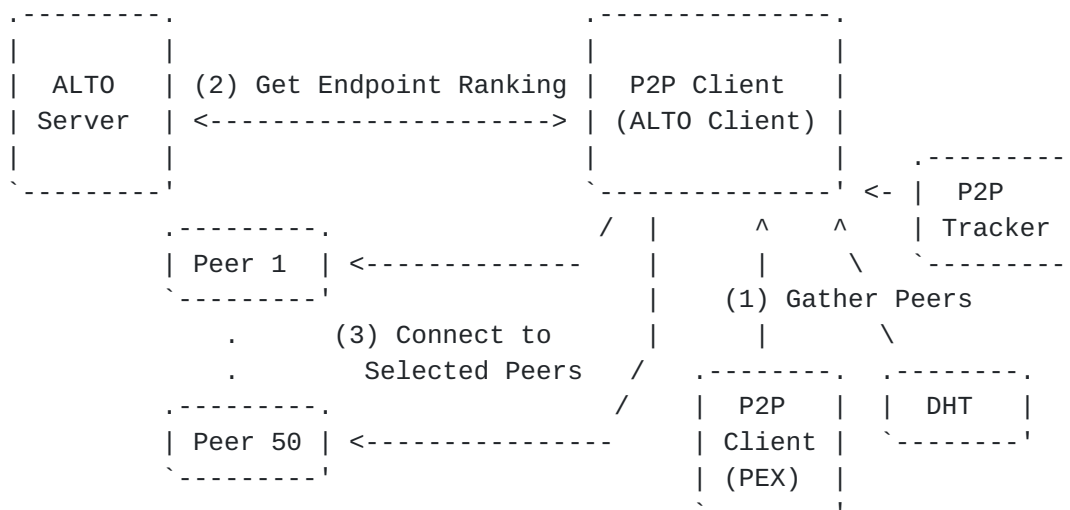


Figure 5: ALTO Client Embedded in P2P Client: Ranking

Figure 5 shows an example of this scenario. The use case proceeds as follows:

1. The P2P Client discovers peers from sources such as Peer Exchange (PEX) from other P2P Clients, Distributed Hash Tables (DHT), and P2P Trackers.
2. The P2P Client queries the ALTO Server's Ranking Service, including discovered peers as the set of Destination Endpoints, and indicates the 'ordinal' Cost Mode. The response indicates the ranking of the candidate peers.

3. The P2P Client connects to the peers in the order specified in the ranking.

9. Discussions

9.1. Discovery

The particular mechanism by which an ALTO Client discovers its ALTO Server is an important component to the ALTO architecture and numerous techniques have been discussed [13] [14]. However, the discovery mechanism is out of scope for this document.

Some ISPs have proposed the possibility of delegation, in which an ISP provides information for customer networks which do not wish to run Portal Servers themselves. A consideration for delegation is that customer networks may wish to explicitly configure such delegation.

9.2. Network Address Translation Considerations

At this day and age of NAT v4<->v4, v4<->v6 [15], and possibly v6<->v6[16], a protocol should strive to be NAT friendly and minimize carrying IP addresses in the payload, or provide a mode of operation where the source IP address provide the information necessary to the server.

The protocol specified in this document provides a mode of operation where the source NL-ID is computed by the ALTO Server (via the Endpoint Property Lookup interface) from the source IP address found in the ALTO Client query packets. This is similar to how some P2P Trackers (e.g., BitTorrent Trackers - see "Tracker HTTP/HTTPS Protocol" in [17]).

The ALTO client SHOULD use the Session Traversal Utilities for NAT (STUN) [4] to determine a public IP address to use as a source NL-ID. If using this method, the host MUST use the "Binding Request" message and the resulting "XOR-MAPPED-ADDRESS" parameter that is returned in the response. Using STUN requires cooperation from a publicly accessible STUN server. Thus, the ALTO client also requires configuration information that identifies the STUN server, or a domain name that can be used for STUN server discovery. To be selected for this purpose, the STUN server needs to provide the public reflexive transport address of the host.

9.3. Mapping IPs to ASNs

It may be desired for the ALTO Protocol to provide ALTO information including ASNs. Thus, ALTO Clients may need to identify the ASN for a Resource Provider to determine the cost to that Resource Provider.

Applications can already map IPs to ASNs using information from a BGP Looking Glass. To do so, they must download a file of about 1.5MB when compressed (as of October 2008, with all information not needed for IP to ASN mapping removed) and periodically (perhaps monthly) refresh it.

Alternatively, Reverse Property Lookup query defined in this document could be extended to map ASNs into a set of IP prefixes. The mappings provided by the ISP would be both smaller and more authoritative.

For simplicity of implementation, it's highly desirable that clients only have to implement exactly one mechanism of mapping IPs to ASNs.

9.4. Endpoint and Path Properties

An ALTO Server could make available many properties about Endpoints beyond their network location or grouping. For example, connection type, geographical location, and others may be useful to applications. The current draft focuses on network location and grouping, but the protocol may be extended to handle other Endpoint properties.

9.5. P2P Peer Selection

This section discusses possible approaches to peer selection using ALTO information (Network Location Identifiers and associated Costs) from an ALTO Server. Specifically, the application must select which peers to use based on this and other sources of information. With this in mind, the usage of ALTO Costs is intentionally flexible, because:

Different applications may use the information differently. For example, an application that connects to just one address may have a different algorithm for selecting it than an application that connects to many.

Though initial experiments have been conducted [[18](#)], more investigation is needed to identify other methods.

In addition, the application might account for robustness, perhaps using randomized exploration to determine if it performs better

without ALTO information.

9.5.1. Client-based Peer Selection

One possibility is for peer selection using ALTO costs to be done entirely by a P2P client. The following are some techniques have been proposed and/or used:

- o Prefer network locations with lower ordinal rankings (i.e., higher priority) [[19](#)] [[8](#)].
- o Optimistically unchoking low-cost peers with higher probability [[8](#)].

9.5.2. Server-based Peer Selection

Another possibility is for ALTO costs to be used by an Application Tracker (e.g., BitTorrent Tracker) when returning peer lists. The following are techniques that have been proposed and/or used:

- o Using bandwidth matching (e.g., at an Application Tracker) and choosing solution (within bound of optimal) with minimal network cost [[18](#)].

10. IANA Considerations

This document request the registration of a new media type:
"application/alto"

11. Security Considerations

11.1. ISPs

ISPs must be cognizant of the network topology and provisioning information provided through ALTO Interfaces. ISPs should evaluate how much information is revealed and the associated risks. In particular, providing overly fine-grained information may make it easier for attackers to infer network topology. On the other hand, revealing overly coarse-grained information may not provide benefits to network efficiency or performance improvements to ALTO Clients.

11.2. ALTO Clients

Applications using the information must be cognizant of the possibility that the information is malformed or incorrect. Even when it is correct, its use might harm the performance. When an

application concludes that it would get better performance disregarding the ALTO information, the decision to discontinue the use of ALTO information is likely best left to the user.

ALTO Clients should also be cognizant of revealing Network Location Identifiers (IP addresses or fine-grained PIDs) to the ALTO Server, as doing so may allow the ALTO Server to infer communication patterns. One possibility is for the ALTO Client to only rely on Network Map for PIDs and Cost Map amongst PIDs to avoid passing IP addresses of their peers to the ALTO Server.

11.3. ALTO Information

An ALTO Server may optionally use authentication and encryption to protect ALTO information. SSL/TLS can provide encryption as well as authentication of the client and server. HTTP Basic or Digest authentication can provide authentication of the client (combined with SSL/TLS, it can additionally provide encryption and authentication of the server).

ISPs should be cognizant that encryption only protects ALTO information until it is decrypted by the intended ALTO Client. Digital Rights Management (DRM) techniques and legal agreements protecting ALTO information are outside of the scope of this document.

11.4. ALTO Information Redistribution

It is possible for applications to redistribute ALTO information to improve scalability. Even with such a distribution scheme, ALTO Clients obtaining ALTO information must be able to validate the received ALTO information to ensure that it was actually generated by the correct ALTO Server. Further, to prevent the ALTO Server from being a target of attack, the verification scheme must not require ALTO Clients to contact the ALTO Server to validate every set of information.

Note that the redistribution scheme must additionally handle details such as ensuring ALTO Clients retrieve ALTO information from the correct ALTO Server. See [20] and [21] for further discussion. Details of a particular redistribution scheme are outside the scope of this document.

To fulfill these requirements, ALTO Information meant to be redistributable contains a digital signature which includes a hash of the ALTO information signed by the ALTO Server's private key. The corresponding public key should either be part of the ALTO information itself, or it could be included in the server capability

response. The public key SHOULD include the hostname of the ALTO Server and it SHOULD be signed by a trusted authority.

11.5. Denial of Service

ISPs should be cognizant of the workload at the ALTO Server generated by certain ALTO Queries, such as certain queries to the Map Filtering Service and Ranking Service. The Map Service allows ALTO Servers to pre-generate maps that can be useful to many ALTO Clients.

12. References

12.1. Normative References

- [1] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [2] Berners-Lee, T., Fielding, R., and H. Nielsen, "Hypertext Transfer Protocol -- HTTP/1.0", [RFC 1945](#), May 1996.
- [3] 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.
- [4] Rosenberg, J., Mahy, R., Matthews, P., and D. Wing, "Session Traversal Utilities for (NAT) (STUN)", [draft-ietf-behave-rfc3489bis-18](#) (work in progress), July 2008.

12.2. Informative References

- [5] Kiesel, S., Popkin, L., Previdi, S., Woundy, R., and Y. Yang, "Application-Layer Traffic Optimization (ALTO) Requirements", [draft-kiesel-alto-reqs-01](#) (work in progress), November 2008.
- [6] Alimi, R., Pasko, D., Popkin, L., Wang, Y., and Y. Yang, "P4P: Provider Portal for P2P Applications", [draft-p4p-framework-00](#) (work in progress), November 2008.
- [7] Wang, Y., Alimi, R., Pasko, D., Popkin, L., and Y. Yang, "P4P Protocol Specification", [draft-wang-alto-p4p-specification-00](#) (work in progress), March 2009.
- [8] Shalunov, S., Penno, R., and R. Woundy, "ALTO Information Export Service", [draft-shalunov-alto-infoexport-00](#) (work in progress), October 2008.
- [9] Das, S. and V. Narayanan, "A Client to Service Query Response

- Protocol for ALTO", [draft-saumitra-alto-queryresponse-00](#) (work in progress), March 2009.
- [10] Das, S., Narayanan, V., and L. Dondeti, "ALTO: A Multi Dimensional Peer Selection Problem", [draft-saumitra-alto-multi-ps-00](#) (work in progress), October 2008.
- [11] Seedorf, J. and E. Burger, "Application-Layer Traffic Optimization (ALTO) Problem Statement", [draft-marocco-alto-problem-statement-04](#) (work in progress), February 2009.
- [12] Yang, Y., Popkin, L., Penno, R., and S. Shalunov, "An Architecture of ALTO for P2P Applications", [draft-yang-alto-architecture-00](#) (work in progress), March 2009.
- [13] Garcia, G., Tomsu, M., and Y. Wang, "ALTO Discovery Protocols", [draft-wang-alto-discovery-00](#) (work in progress), March 2009.
- [14] Song, H., Even, R., Pascual, V., and Y. Zhang, "Application-Layer Traffic Optimization (ALTO): Discover ALTO Servers", [draft-song-alto-server-discovery-00](#) (work in progress), March 2009.
- [15] Baker, F., Li, X., and C. Bao, "Framework for IPv4/IPv6 Translation", [draft-baker-behave-v4v6-framework-02](#) (work in progress), February 2009.
- [16] Wasserman, M. and F. Baker, "IPv6-to-IPv6 Network Address Translation (NAT66)", [draft-mrw-behave-nat66-02](#) (work in progress), March 2009.
- [17] "Bittorrent Protocol Specification v1.0", <http://wiki.theory.org/BitTorrentSpecification>, 2009.
- [18] H. Xie, YR. Yang, A. Krishnamurthy, Y. Liu, and A. Silberschatz., "P4P: Provider Portal for (P2P) Applications", In SIGCOMM 2008.
- [19] Akonjang, O., Feldmann, A., Previdi, S., Davie, B., and D. Saucez, "The PROXIDOR Service", [draft-akonjang-alto-proxidior-00](#) (work in progress), March 2009.
- [20] Yingjie, G., Alimi, R., and R. Even, "ALTO information redistribution", [draft-gu-alto-redistribution-00](#) (work in progress), October 2009.

- [21] Stiemerling, M., "ALTO Information Redistribution Considered Harmful", [draft-stiemerling-alto-info-redist-00](#) (work in progress), August 2009.

Appendix A. Contributors

The people listed here should be viewed as co-authors of the document. Due to the limit of 5 authors per draft the co-authors were moved to the contributors section at this point.

Obi Akonjang

DT Labs/TU Berlin/

Email: obi@net.t-labs.tu-berlin.de

Richard Alimi

Yale University

Email: richard.alimi@yale.edu

Saumitra M. Das

Qualcomm Inc.

Email: saumitra@qualcomm.com

Syon Ding

China Telecom

Email: syding@chinatelecom.com

Doug Pasko

Verizon

EMail: pasko@verizon.com

Laird Popkin

Pando Networks

EMail: laird@pando.com

Stefano Previdi

Cisco

EMail: sprevidi@cisco.com

Satish Raghunath

Juniper Networks

satishr@juniper.net

Stanislav Shalunov

BitTorrent

EMail: shalunov@bittorrent.com

Albert Tian

Ericsson/Redback

EMail: alberttian@gmail.com

Yu-Shun Wang

Microsoft Corp.

yu-shun.wang@microsoft.com

Richard Woundy

Comcast

Richard_Woundy@cable.comcast.com

David Zhang

PPLive

davidzhang@pplive.com

Yunfei Zhang

China Mobile

zhangyunfei@chinamobile.com

Appendix B. Acknowledgements

We would like to thank the following additional people who were involved in the projects that contributed to this merged document: Alex Gerber (AT&T), Chris Griffiths (Comcast), Ramit Hora (Pando Networks), Arvind Krishnamurthy (University of Washington), Marty Lafferty (DCIA), Erran Li (Bell Labs), Jin Li (Microsoft), Y. Grace Liu (IBM Watson), Jason Livingood (Comcast), Michael Merritt (AT&T), Ingmar Poesse (DT Labs/TU Berlin), James Royalty (Pando Networks), Damien Saucez (UCL) Thomas Scholl (AT&T), Emilio Sepulveda (Telefonica), Avi Silberschatz (Yale University), Hassan Sipra (Bell Canada), Georgios Smaragdakis (DT Labs/TU Berlin), Haibin Song (Huawei), Oliver Spatscheck (AT&T), See-Mong Tang (Microsoft), Jia Wang (AT&T), Hao Wang (Yale University), Ye Wang (Yale University), Haiyong Xie (Yale University).

Authors' Addresses

Reinaldo Penno (editor)
Juniper Networks
1194 N Mathilda Avenue
Sunnyvale, CA
USA

Email: rpenno@juniper.net

Y. Richard Yang (editor)
Yale University

Email: yry@cs.yale.edu

