ALTO WG
Internet-Draft
Intended status: Standards Track
Expires: September 15, 2011

R. Alimi, Ed.
Google
R. Penno, Ed.
Juniper Networks
Y. Yang, Ed.
Yale University
March 14, 2011

ALTO Protocol
draft-ietf-alto-protocol-07.txt

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 a 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].

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 September 15, 2011.

Copyright Notice

Table of Contents

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 the 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 this document is to specify a simple and unified protocol that meets the ALTO requirements [14] while 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

- P4P [15], [16], [17];
- ALTO Info-Export [18];
- Query/Response [19], [20];
- ATTP [ATTP];
- Proxidor [21].

See Appendix A for a list of people that have contributed significantly to this effort and the projects and proposals listed above.

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 peer selection process in distributed applications in order to increase locality of traffic,

improve user-experience, amongst others. It also helps ISPs to efficiently manage traffic that traverses more expensive links such as transit and backup 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. This section first introduces the terminology, and then defines the ALTO architecture and the ALTO Protocol's place in the overall architecture.

2.1. Terminology

We use the following terms defined in [22]: 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, Autonomous System Number (ASN), 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.

Each Endpoint Address has an associated Address Type, which indicates

both its syntax and semantics.

2.1.2. ASN

An Autonomous System Number.

2.1.3. Network Location

Network Location is a generic term denoting a single endpoint or group of endpoints.

2.1.4. ALTO Information

ALTO Information is a generic term referring to the network information sent by an ALTO Server.

2.1.5. ALTO Information Base

Internal representation of the ALTO Information maintained by the ALTO Server. Note that the structure of this internal representation is not defined by this document.

2.2. ALTO Service and Protocol Scope

An ALTO Server conveys the network information from the perspective of a network region; the ALTO Server presents its "my-Internet View" [23] of the network region. A network region in this context can be an Autonomous System, an ISP, or perhaps a smaller region or 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 Server prepares ALTO Information; an ALTO Client uses ALTO Service Discovery to identify an appropriate ALTO Server; 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.

Note that it may also be possible for ALTO Servers to exchange network information with other ALTO Servers (either within the same administrative domain or another administrative domain with the consent of both parties) in order to adjust exported ALTO Information. Such a protocol is also outside the scope of this specification.

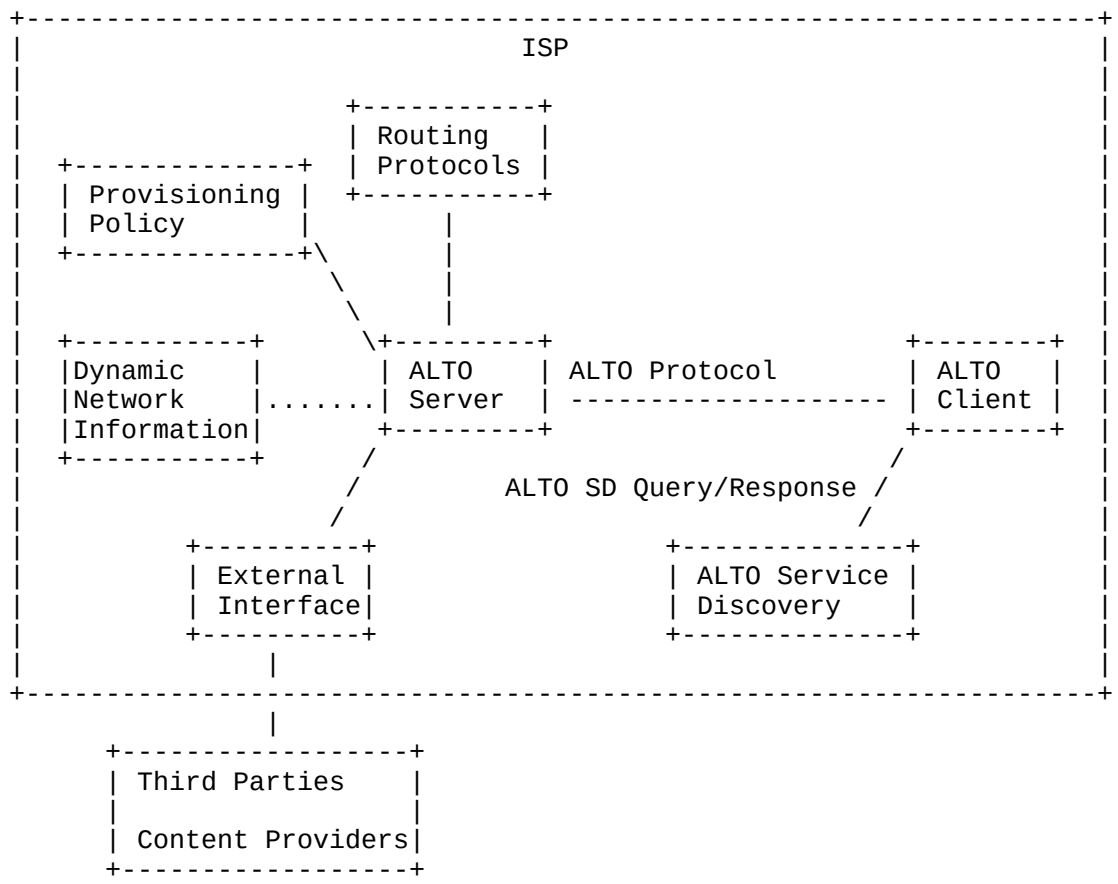

Figure 1: Basic ALTO Architecture.

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 Network Locations and the paths

between Network Locations.

In this document, we focus on a particular Endpoint property to denote the location of an endpoint, and provider-defined costs for paths between pairs of Network Locations.

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 Server Information Service and the Map Service to provide ALTO Information. Other ALTO Information services can provide additional functionality. There are three such services defined in this document: the Map Filtering Service, Endpoint Property Service, and Endpoint Cost Service. Additional services may be defined in companion documents. Note that functionality offered in different services are not totally non-overlapping (e.g., the Map Service and Map Filtering Service).

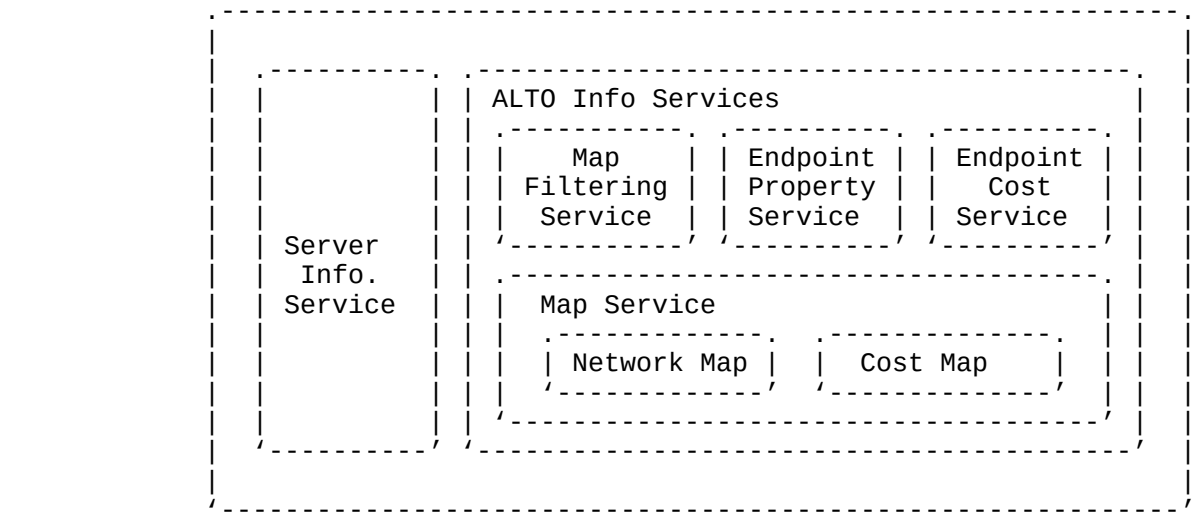

Figure 2: ALTO Protocol Structure

3.1. Server Information Service

The Server (Capability) Information Service lists the details on the information that can be provided by an ALTO Server and perhaps other ALTO Servers maintained by the network provider. The configuration includes, for example, details about the operations and cost metrics supported by the ALTO Server and other related ALTO Servers that may be usable by an ALTO Client. The capability document can be

downloaded by ALTO Clients. The capability information could also be provisioned to devices, but care must be taken to update it appropriately.

3.2. ALTO Information Services

Multiple, distinct services are defined to allow ALTO Clients to query ALTO Information from an ALTO Server. The ALTO Server internally maintains an ALTO Information Base that encodes the network provider's preferences. The ALTO Information Base encodes the Network Locations defined by the ALTO Server (and their corresponding properties), as well as the provider-defined costs between pairs of Network Locations.

3.2.1. Map Service

The Map Service provides batch information to ALTO Clients in the form of Network Map and Cost Map. 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) 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 or CPU collecting results and performing client-side filtering. The Map Filtering Service allows ALTO Clients to query for the ALTO Server Network Map and Cost Map 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. Endpoint Cost Service

Some ALTO Clients may also benefit from querying for costs and rankings based on endpoints. The Endpoint Cost Service allows an ALTO Server to return either numerical costs or ordinal costs (rankings) directly amongst Endpoints.

4. Network Map

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 losing 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 definition of proximity varies depending on the granularity of the ALTO information configured by the provider. 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.

As used in this document, the Network Map refers to the syntax and semantics of the information distributed by the ALTO Server. This document does not discuss the internal representation of this data structure within the ALTO Server.

4.1. PID

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

A PID is an identifier that provides 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 and can improve scalability. In particular, network preferences (costs) may be specified between PIDs, allowing cost information to be more compactly represented and updated at a faster time scale than the network aggregations themselves.

Using PIDs, the Network Map may also be used to communicate simple preferences with only minimal information from the Cost Map. 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. The ISP may aggregate endhosts within a PoP into a single PID in the Network Map. The Cost Map may be encoded to indicate that peering within the same PID is preferred; for example, cost(PID_i, PID_i) == c* and cost(PID_i, PID_j) > c* for i != j. Section 5 provides further

details about Cost Map structure.

4.2. Endpoint Addresses

Communicating endpoints may have many types of addresses, such as IP addresses, MAC addresses, or overlay IDs. The current specification only considers IP addresses.

4.2.1. IP Addresses

The endpoints aggregated into a PID are denoted by a list of IP prefixes. When either an ALTO Client or ALTO Server needs to determine which PID in a Network Map contains a particular IP address, longest-prefix matching MUST be used.

A Network Map MUST define a PID for each possible address in the IP address space for all of the address types contained in the map. A RECOMMENDED way to satisfy this property is to define a PID that contains the 0.0.0.0/0 prefix for IPv4 or ::/0 (for IPv6).

4.3. Example Network Map

Figure 3 illustrates an example Network Map. PIDs are used to identify network-agnostic aggregations.

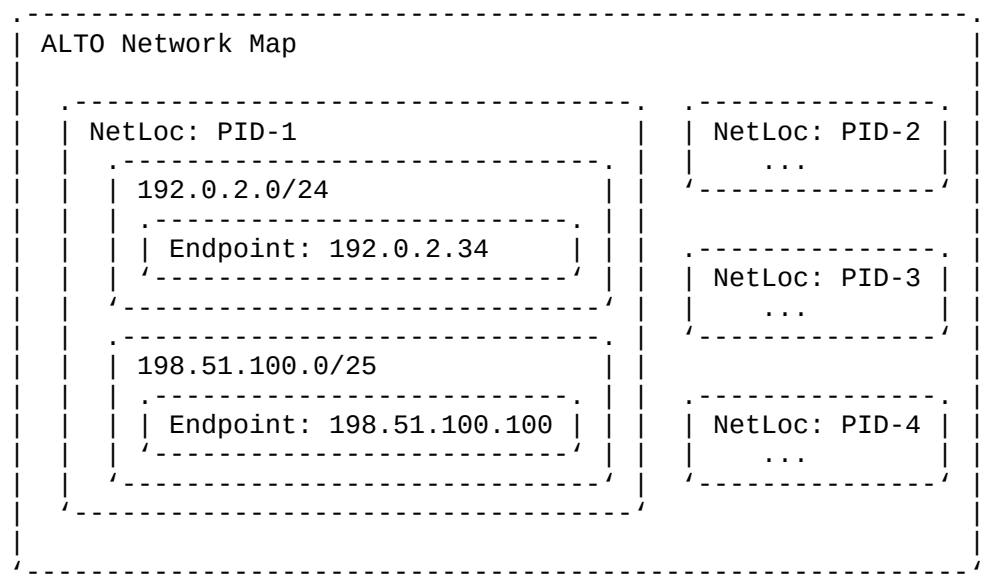

Figure 3: Example Network Map

5. Cost Map

An ALTO Server indicates preferences amongst network locations in the form of Path Costs. Path Costs are generic costs and can be internally computed by a network provider according to its own needs.

An ALTO Cost Map defines Path Costs pairwise amongst sets of source and destination Network Locations.

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.

As used in this document, the Cost Map refers to the syntax and semantics of the information distributed by the ALTO Server. This document does not discuss the internal representation of this data structure within the ALTO Server.

5.1. Cost Attributes

Path Costs have attributes:

- Type: identifies what the costs represent;
- Mode: identifies how the costs should be interpreted.

Certain queries for Cost Maps allow the ALTO Client to indicate the desired Type and Mode.

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 strings.

5.1.1.1. Cost Type: routingcost

An ALTO Server MUST define the ’routingcost’ Cost Type.

This Cost Type conveys a generic measure for the cost of routing traffic from a source to a destination. Lower values indicate a higher preference for traffic to be sent from a source to a destination.

Note that an ISP may internally compute routing cost using any method

it chooses (e.g., air-miles or hop-count) as long as it conforms to these semantics.

5.1.2. Cost Mode

The Mode attribute indicates how costs should be interpreted. Specifically, the Mode attribute indicates whether returned costs 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 handle such costs differently.

Cost Modes are indicated in protocol messages as strings.

An ALTO Server MUST support at least one of 'numerical' and 'ordinal' costs. ALTO Clients SHOULD be cognizant of operations when a desired cost mode is not supported. For example, an ALTO Client desiring numerical costs may adjust behavior if only the ordinal Cost Mode is available. Alternatively, an ALTO Client desiring ordinal costs may construct ordinal costs given numerical values if only the numerical Cost Mode is available.

5.1.2.1. Cost Mode: numerical

This Cost Mode is indicated by the string 'numerical'. This mode indicates that it is safe to perform numerical operations (e.g. summation) on the returned costs.

5.1.2.2. Cost Mode: ordinal

This Cost Mode is indicated by the string 'ordinal'. This mode indicates that the costs values to a set of Destination Network Locations from a particular Source Network Location are a ranking, with lower values indicating a higher preference.

It is important to note that the values in the Cost Map provided with the ordinal Cost Mode are not necessarily the actual cost known to the ALTO Server.

5.2. Cost Map Structure

A query for a Cost Map either explicitly or implicitly includes a list of Source Network Locations and a list of Destination Network

Locations. (Recall that a Network Location can be an endpoint address or a PID.)

Specifically, assume that a 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].

The ALTO Server will return 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). We refer to this structure as a Cost Map.

If the Cost Mode is 'ordinal', the Path Cost of each communicating pair is relative to the m*n entries.

5.3. Network Map and Cost Map Dependency

If a Cost Map contains PIDs in the list of Source Network Locations or the list of Destination Network Locations, the Path Costs are generated based on a particular Network Map (which defines the PIDs). Version Tags are introduced to ensure that ALTO Clients are able to use consistent information even though the information is provided in two maps.

A Version Tag is an opaque string associated with a Network Map maintained by the ALTO Server. When the Network Map changes, the Version Tag SHOULD also be changed. (Thus, the Version Tag is defined similarly to HTTP's ETag.) Possibilities for generating a Version Tag include the last-modified timestamp for the Network Map, or a hash of its contents.

A Network Map distributed by the ALTO Server includes its Version Tag. A Cost Map referring to PIDs also includes the Version Tag of the Network Map on which it is based.

6. Protocol Design Overview

The ALTO Protocol design uses a REST-like interface with the goal of leveraging current HTTP [2] [3] implementations and infrastructure, as well as familiarity with existing REST-like services in popular use. ALTO messages use JSON [4] to encode message bodies.

This document currently specifies both services and message encoding in a descriptive fashion. Care is taken to make descriptions precise and unambiguous, but it still lacks benefits of automatic tooling that exists for certain encoding formats.

Standards such as WSDL 2.0 and WADL are capable of describing available interfaces. JSON Schema [24] allows message encodings to be specified precisely and messages may be verified against the schema. It is not yet clear whether such an approach should be taken in this document.

Benefits enabled by a REST-like interface leveraging HTTP include easier understanding and debugging, flexible ALTO Server implementation strategies, and more importantly, simple caching and redistribution of ALTO information to increase scalability.

6.1. Existing Infrastructure

HTTP is a natural choice for integration with existing applications and infrastructure. In particular, the ALTO Protocol design leverages:

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

6.2. ALTO Information Reuse and Redistribution

ALTO information may be useful to a large number of applications and users. For example, an identical Network Map may be used by all ALTO Clients querying a particular ALTO Server. At the same time, distributing ALTO information must be efficient and not become a bottleneck.

Beyond integration with existing HTTP caching infrastructure, ALTO information may also be cached or redistributed using application-dependent mechanisms, such as P2P DHTs or P2P file-sharing. This document does not define particular mechanisms for such redistribution, but it does define the primitives (e.g., digital signatures) that may be needed to support such a mechanism. See [25] for further discussions.

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 12.4. Protocol support for redistribution is specified in Section 8.

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.

7.1. Notation

This document uses an adaptation of the C-style struct notation to define the required and optional members of JSON objects. Unless explicitly noted, each member of a struct is REQUIRED.

The types 'JSONString', 'JSONNumber', 'JSONBool' indicate the JSON string, number, and boolean types respectively.

This document only includes object members used by this specification. It is possible that protocol extensions include additional members to JSON objects defined in this document; such additional members will be silently ignored by ALTO Servers and Clients only implementing the base protocol defined in this document.

7.2. Message Format

Request and Response follow the standard format for HTTP Request and Response messages [2] [3].

The following subsections provide an overview of how ALTO Requests and Responses are encoded in HTTP, and discusses rationale for certain design decisions.

7.2.1. Protocol Versioning

The ALTO Protocol uses a simple versioning approach that permits evolution between versions even if ALTO information is being served as static, pre-generated files.

It is assumed that a single host responding to ALTO Requests implements a single protocol version. Virtual hosting may be used if multiple protocol versions need to be supported by a single physical server.

A common query (Server List, detailed in Section 7.8.1.1) to be present in all ALTO protocol versions allows an ALTO Client to discover additional ALTO Servers and the ALTO Protocol version number of each.

This approach keeps the ALTO Server implementation free from parsing

and directing each request based on version number. Although ALTO Requests are free from protocol version numbers, the protocol version number is echoed in each ALTO Response to keep responses self-contained to, for example, ease reading persisted or redistributed ALTO responses.

Using virtual hosting with TLS may require the Server Name Indication extension for TLS [5] [26].

This document specifies ALTO Protocol version 1.

7.2.2. Content Type

All ALTO Request and Response messages MUST set the Content-Type HTTP header to "application/alto".

7.2.3. Request Message

An ALTO Request is a standard HTTP Request generated by an ALTO Client, with certain components defined by the ALTO Protocol.

The basic syntax of an ALTO Request is:

```
<Method> /<Resource> HTTP/1.1
Host: <Host>
```

For example:

```
GET /info/capability HTTP/1.1
Host: alto.example.com:6671
```

7.2.3.1. Standard HTTP Headers

The Host header MUST follow the standard rules for the HTTP 1.1 Host Header.

The Content-Length header MUST follow the standard rules defined in HTTP 1.1.

The Content-Type HTTP Header MUST have value "application/alto" if the Body is non-empty.

7.2.3.2. Method and Resource

Next, both the HTTP Method and URI-Path (denoted as Resource) indicate the operation requested by the ALTO Client. In this example, the ALTO Client is requesting basic capability information from the ALTO Server.

7.2.3.3. Input Parameters

Certain operations defined by the ALTO Protocol (e.g., in the Map Filtering Service) allow the ALTO Client to supply additional input parameters. Such input parameters are encoded in a URI-Query-String where possible and appropriate. However, due to practical limitations (e.g. underlying HTTP implementations may have limitations on the total length of a URI and the Query-String is better-suited for simple unstructured parameters and lists), some operations in the ALTO Protocol use input parameters encoded in the HTTP Request Body.

7.2.4. Response Message

A Response message is a standard HTTP Response generated by an ALTO Server with certain components defined by the ALTO Protocol.

The basic syntax of an ALTO Response is:

```
HTTP/1.1 <StatusCode> <StatusMsg>
Content-Length: <ContentLength>
Content-Type: <ContentType>

<ALTOResponse>
```

where the HTTP Response Body is an ALTOResponse JSON Object (defined in Section 7.2.4.3). For example:

```
HTTP/1.1 200 OK
Content-Length: 1000
Content-Type: application/alto

{
    "meta" : {
        "version": 1,
        "status" : {
            "code" : "SUCCESS",
            "reason" : "Success"
        },
        ...
    },
    "type" : "capability",
    "data" : {
        ...
    }
}
```

7.2.4.1. Standard HTTP Headers

The Content-Length header MUST follow the standard rules defined in HTTP 1.1.

The Content-Type HTTP Header MUST have value "application/alto" if the Body is non-empty.

7.2.4.2. Status Code and Message

Two sets of status codes are used in the ALTO Protocol. First, an ALTO Status Code provides detailed information about the success or failure of a particular operation. Second, an HTTP Status Code indicates to HTTP processing elements (e.g., intermediaries and clients) how the response should be treated.

7.2.4.3. HTTP Body

The Response body MUST encode a single top-level JSON object of type ALTOResponse:

```
object {
    RspMetaData   meta;
    JSONString    type;
    [RspDataType] data;
} ALTOResponse;
```

The ALTOResponse object has distinct sections for:

- meta information encoded in an extensible way,
- the type of ALTO Information to follow, and
- the requested ALTO Information.

7.2.4.3.1. Meta Information

Meta information is encoded as a JSON object with type RspMetaData:

```
object {
    JSONString     code;
    JSONString     reason;          [OPTIONAL]
} RspStatus;

object {
    JSONNumber     version;
    RspStatus      status;
    RspRedistDesc  redistribution;  [OPTIONAL]
```

```
  } RspMetaData;
```

with members:

- version: the ALTO Protocol version, which MUST be an integer
- status: an ALTO Status Code from Section 7.4 and corresponding reason (free-form string) providing a human-readable explanation of the particular status code.
- redistribution: see Section 8.

7.2.4.3.2. ALTO Information

If the Response is successful (see Section 7.4), then the "type" and "data" members of the ALTOResponse object are REQUIRED. "type" encodes a Response-specific string which indicates to the ALTO Client the type of data encoded in the message. The "data" member encodes the actual Response-specific data; the structure of this member is detailed later in this section for each particular ALTO Response.

7.2.4.4. Signature

An ALTO Server MAY additionally supply a signature asserting that it generated a particular response. See Section 8.2.2.

7.3. 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 with which to communicate, establishes a TCP connection, and constructs and sends ALTO Request messages which MUST conform to Section 7.8. In response to Request messages, an ALTO Server constructs and sends ALTO Response messages which also MUST conform to Section 7.8.

7.4. ALTO Status Codes

This document defines ALTO Status Codes to support the operations defined in this document. Additional status codes may be defined in companion or extension documents.

An ALTO Server MUST return the SUCCESS status code if and only if the Request message is successfully processed and the requested ALTO information is returned by the ALTO Server.

The HTTP Status Codes corresponding to each ALTO Status Code are defined to provide correct behavior with HTTP intermediaries and

clients. When an ALTO Server returns a particular ALTO Status Code, it MUST indicate one of the corresponding HTTP Status Codes in Table 1.

If multiple errors are present in a single ALTO Request (e.g., a request uses a JSONString when a JSONInteger is expected and a required field is missing), then the ALTO Server MUST return exactly one of the detected errors. However, the reported error is implementation defined, since specifying a particular order for message processing encroaches needlessly on implementation technique.

ALTO Status Code	HTTP Status Code(s)	Description
SUCCESS	2xx	Success
E_JSON_SYNTAX	400	JSON parsing error in request
E_JSON_FIELD_MISSING	400	Required field missing
E_JSON_VALUE_TYPE	400	JSON Value of unexpected type
E_INTERNAL_ERROR	500	Server-side error
E_INVALID_OPERATION	501	Invalid operation requested
E_INVALID_COST_TYPE	501	Invalid cost type

Table 1: Defined ALTO Status Codes

Status codes described in Table 1 are a work in progress. This document will be modified to update the available status codes as implementation experience is gained. Feedback is welcomed.

In addition, feedback from implementers of ALTO Clients is welcomed to identify if there is a need to communicate multiple status codes in a single response.

7.5. Client Behavior

7.5.1. Successful Response

This specification does not indicate any required actions taken by ALTO Clients upon receiving a successful response from an ALTO Server. Although ALTO Clients are suggested to interpret the received ALTO Information and adapt application behavior, ALTO Clients are not required to do so.

7.5.2. Error Conditions

If an ALTO Client does not receive a successful response from the ALTO Server, it can either choose another server or fall back to a default behavior (e.g., perform peer selection without the use of ALTO information). An ALTO Client may also retry the request at a later time.

7.6. HTTP Usage

7.6.1. Authentication and Encryption

An ALTO Server MAY support SSL/TLS to implement server and/or client authentication, as well as encryption. See [6] for considerations regarding verifcation of server identity.

An ALTO Server MAY support HTTP Digest authentication.

7.6.2. Cookies

Cookies MUST NOT be used.

7.6.3. Caching Parameters

If the Response generated by the ALTO Server is cachable, the ALTO Server MAY include 'Cache-Control' and 'Expires' HTTP headers.

If a Response generated by the ALTO Server is not cachable, the ALTO Server MUST specify the "Cache-Control: no-cache" HTTP Header.

7.7. ALTO Types

This section details the encoding for particular data values used in the ALTO Protocol.

7.7.1. PID Name

A PID Name is encoded as a US-ASCII string. The string MUST be no more than 32 characters, and MUST NOT contain characters other than alphanumeric characters or the '.' separator. The '.' separator is reserved for future use and MUST NOT unless specifically indicated by a companion or extension document.

The type 'PIDName' is used in this document to indicate a string of this format.

7.7.2. Endpoints

7.7.2.1. Address Type

Address Types are encoded as US-ASCII strings consisting of only alphanumeric characters. This document defines the address type "ipv4" to refer to IPv4 addresses, and "ipv6" to refer to IPv6 addresses. Extension documents may define additional Address Types.

The type 'AddressType' is used in this document to indicate a string of this format.

7.7.2.2. Endpoint Address

Endpoint Addresses are encoded as US-ASCII strings. The exact characters and format depend on the type of endpoint address.

The type 'EndpointAddr' is used in this document to indicate a string of this format.

7.7.2.2.1. IPv4

IPv4 Endpoint Addresses are encoded as specified by the 'IPv4address' rule in Section 3.2.2 of [7].

7.7.2.2.2. IPv6

IPv6 Endpoint Addresses are encoded as specified in Section 2.2 of [8].

7.7.2.2.3. Typed Endpoint Addresses

When an Endpoint Address is used, an ALTO implemenation must be able to determine its type. For this purpose, the ALTO Protocol allows endpoint addresses to also explicitly indicate their type.

Typed Endpoint Addresses are encoded as US-ASCII strings of the format 'AddressType:EndpointAddr' (with the ':' character as a separator). The type 'TypedEndpointAddr' is used to indicate a string of this format.

7.7.2.3. Endpoint Prefixes

For efficiency, it is useful to denote a set of Endpoint Addresses using a special notation (if one exists). This specification makes use of the prefix notations for both IPv4 and IPv6 for this purpose.

Endpoint Prefixes are encoded as US-ASCII strings. The exact

characters and format depend on the type of endpoint address.

The type 'EndpointPrefix' is used in this document to indicate a string of this format.

7.7.2.3.1. IPv4

IPv4 Endpoint Prefixes are encoded as specified in Section 3.1 of [9].

7.7.2.3.2. IPv6

IPv6 Endpoint Prefixes are encoded as specified in Section 2.3 of [8].

7.7.2.4. Endpoint Address Group

The ALTO Protocol includes messages that specify potentially large sets of endpoint addresses. Endpoint Address Groups provide an efficient way to encode such sets, even when the set contains endpoint addresses of different types.

An Endpoint Address Group is defined as:

```
object {
    EndpointPrefix [AddressType]<0..*>;
    ...
} EndpointAddrGroup;
```

In particular, an Endpoint Address Group is a JSON object, with the name of each member being the string corresponding to the address type, and the member's corresponding value being a list of prefixes of addresses of that type.

The following is an example with both IPv4 and IPv6 endpoint addresses:

```
{
    "ipv4": [
        "192.0.2.0/24",
        "198.51.100.0/25"
    ],
    "ipv6": [
        "2001:db8:0:1::/64",
        "2001:db8:0:2::/64"
    ]
}
```

7.7.3. Cost Mode

A Cost Mode is encoded as a US-ASCII string. The string MUST either have the value 'numerical' or 'ordinal'.

The type 'CostMode' is used in this document to indicate a string of this format.

7.7.4. Cost Type

A Cost Type is encoded as a US-ASCII string. The string MUST be no more than 32 characters, and MUST NOT contain characters other than alphanumeric characters or the ':' separator.

Identifiers prefixed with 'priv:' are reserved for Private Use [10]. Identifiers prefixed with 'exp:' are reserved for Experimental use. All other identifiers appearing in an ALTO Request or Response MUST be registered in the ALTO Cost Types registry Section 11.

The type 'CostType' is used in this document to indicate a string of this format.

7.8. ALTO Messages

This section documents the individual operations supported in the ALTO Protocol. See Section 7.2.3 and Section 7.2.4 for specifications of HTTP Request/Response components common to all operations in the ALTO Protocol.

Table 2 provides an summary of the HTTP Method and URI-Paths used for ALTO Requests:

Service	Operation	HTTP Method and URI-Path
Server Info	List Servers	GET /info/servers
Server Info	Capability	GET /info/capability
Map	Network Map	GET /map/core/pid/net
Map	Cost Map	GET /map/core/pid/cost
Map Filtering	Network Map	POST /map/filter/pid/net
Map Filtering	Cost Map	POST /map/filter/pid/cost
Endpoint Prop.	Lookup	GET /endpoint/prop/<name>
		POST /endpoint/prop/lookup
Endpoint Cost	Lookup	POST /endpoint/cost/lookup

Table 2: Overview of ALTO Requests

7.8.1. Server Information Service

The Server Information Service provides information about available ALTO Servers and their capabilities (e.g., supported services).

An ALTO Server MUST support the Server Information Service and MUST implement all operations defined in this section.

7.8.1.1. Server List

The Server List request allows an ALTO Client to discover other ALTO Servers provided by the ALTO Service Provider. Upon discovering an additional ALTO Server, the ALTO Client may then query the server capabilities (see Section 7.8.1.2) to test if it supports desired functionality.

The Server List request is intended to help an ALTO Client find an ALTO Server supporting the desired ALTO Protocol version and capabilities. It is not intended to serve as a substitute for the ALTO Server Discovery which helps an ALTO Client locate an initial ALTO Server.

This operation MUST be supported by the ALTO Server.

7.8.1.1.1. Request Syntax

```
GET /info/servers HTTP/1.1
Host: <Host>
```

7.8.1.1.2. Response Syntax

```
HTTP/1.1 200 <StatusMsg>
Content-Length: <BodyLength>
Content-Type: application/alto

<ALTOResponse>
```

where the ALTOResponse object has "type" member equal to the string "server-list" and "data" member of type RspServerList:

```
object {
    JSONString   uri;
    JSONNumber   version;
} ServerItem;

object {
    ServerItem   servers<0..*>;
} RspServerList;
```

RspServerList has members:

- servers: Array of available ALTO Servers, detailing the URI of the ALTO Server and the ALTO Protocol version that it implements. The array must at least contain an entry corresponding to the ALTO Server at the URI from which it is retrieving the server list.

7.8.1.1.3. Example

```
GET /info/servers HTTP/1.1
Host: alto.example.com:6671
```

```
HTTP/1.1 200 OK
Content-Length: [TODO]
Content-Type: application/alto

{
    "meta" : {
        "version" : 1,
        "status" : {
            "code" : "SUCCESS"
        }
    },
    "type" : "server-list",
    "data" : {
        "servers" : [
            {
                "uri": "http://alto.example.com:6671",
                "version" : 1
            }
        ]
    }
}
```

7.8.1.2. Server Capability

The Server Capability request allows an ALTO Client to determine the functionality supported by the queried ALTO Server.

This operation MUST be supported by the ALTO Server.

7.8.1.2.1. Request Syntax

```
GET /info/capability HTTP/1.1
Host: <Host>
```

7.8.1.2.2. Response Syntax

```
HTTP/1.1 200 <StatusMsg>
Content-Length: <BodyLength>
Content-Type: application/alto

<ALTOResponse>
```

where the ALTOResponse object has "type" member equal to the string "capability" and "data" member of type RspCapability:

```
enum {
    map,
    map-filtering,
    endpoint-property,
    endpoint-cost
} ServiceType;          [Note: encoded as JSONString's]

object {
    ServiceType  services<0..*>;
    CostMode     cost-modes<0..*>;      [OPTIONAL]
    CostType     cost-types<0..*>;      [OPTIONAL]
    JSONBool     cost-constraints;      [OPTIONAL]
    JSONString   service-id;            [OPTIONAL]
    JSONString   certificates<0..*>;    [OPTIONAL]
} RspCapability;
```

RspCapability has members:

- services: Lists the services supported by the ALTO Server. The service names defined in this document are are "map", "map-filtering", "endpoint-property", and "endpoint-cost".

- cost-modes: Array of supported ALTO Cost Modes.

- cost-types: Array of supported ALTO Cost Types.

- cost-constraints: Indicates if the ALTO Server supports cost constraints. The value 'false' is implied if this member is not present.

- service-id: UUID [11] indicating an one or more ALTO Servers serving equivalent ALTO Information.

- certificates: List of PEM-encoded X.509 certificates used by the ALTO Server in the signing of responses.

If an ALTO Server denotes a response as redistributable, the 'service-id' and 'certificates' fields are REQUIRED instead of OPTIONAL. See Section 8 for detailed specification.

7.8.1.2.3. Example

```
GET /info/capability HTTP/1.1
Host: alto.example.com:6671
```

```
HTTP/1.1 200 OK
Content-Length: [TODO]
Content-Type: application/alto

{
    "meta" : {
        "version" : 1,
        "status" : {
            "code" : "SUCCESS"
        }
    },
    "type" : "capability",
    "data" : {
        "services" : [ "map", "map-filtering" ],
        "cost-modes": [
            "numerical",
            "ordinal"
        ],
        "cost-types": [
            "routingcost",
            "hopcount"
        ],
        "cost-constraints": false
    }
}
```

7.8.2. Map Service

The Map Service provides batch information to ALTO Clients in the form of two maps: a Network Map and Cost Map.

An ALTO Server MUST support the Map Service and MUST implement all operations defined in this section.

7.8.2.1. Network Map

The full Network Map lists for each PID, the network locations (endpoints) within the PID.

7.8.2.1.1. Request Syntax

```
GET /map/core/pid/net HTTP/1.1
Host: <Host>
```

7.8.2.1.2. Response Syntax

```
HTTP/1.1 200 <StatusMsg>
Content-Length: <BodyLength>
Content-Type: application/alto

<ALTOResponse>
```

where the ALTOResponse object has "type" member equal to the string "network-map" and "data" member of type RspNetworkMap:

```
object {
    EndpointAddrGroup [pidname]<0..*>;
    ...
} NetworkMapData;

object {
    JSONString     map-vtag;
    NetworkMapData map;
} RspNetworkMap;
```

RspNetworkMap has members:

- map-vtag: The Version Tag of the Network Map (Section 5.3)

- map: The network map data itself.

NetworkMapData is a JSON object with each member representing a single PID and its associated set of endpoint addresses. A member's name is a PIDName string denoting the PID's name.

7.8.2.1.3. Example

```
GET /map/core/pid/net HTTP/1.1
Host: alto.example.com:6671
```

```
HTTP/1.1 200 OK
Content-Length: [TODO]
Content-Type: application/alto

{
    "meta" : {
        "version" : 1,
        "status" : {
            "code" : "SUCCESS"
        }
    },
    "type" : "network-map",
    "data" : {
        "map-vtag" : "1266506139",
        "map" : {
            "PID1" : {
                "ipv4" : [
                  "192.0.2.0/24",
                  "198.51.100.0/25"
                ]
            },
            "PID2" : {
                "ipv4" : [
                  "198.51.100.128/25"
                ]
            },
            "PID3" : {
                "ipv4" : [
                  "0.0.0.0/0"
                ],
                "ipv6" : [
                  "::/0"
                ]
            ]
        }
    }
}
```

7.8.2.2. Cost Map

The Map Service Cost Map query is a batch operation in which the ALTO Server returns the Path Cost for each pair of source/destination PID defined by the ALTO Server.

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

7.8.2.2.1. Request Syntax

```
GET /map/core/pid/cost HTTP/1.1
Host: <Host>
```

7.8.2.2.2. Response Syntax

```
HTTP/1.1 200 <StatusMsg>
Content-Length: <BodyLength>
Content-Type: application/alto

<ALTOResponse>
```

where the ALTOResponse object has "type" member equal to the string "cost-map" and "data" member of type RspCostMap:

```
object DstCosts {
    JSONNumber [dstname];
    ...
};

object {
    DstCosts [srcname]<0..*>;
    ...
} CostMapData;

object {
    JSONString  map-vtag;
    CostType    cost-type;
    CostMode    cost-mode;
    CostMapData map;
} RspCostMap;
```

RspCostMap has members:

- map-vtag: The Version Tag of the Network Map used to generate the Cost Map (Section 5.3).
- cost-type: Cost Type used in the map (Section 5.1.1)
- cost-mode: Cost Mode used in the map (Section 5.1.2)
- map: The cost map data itself.

CostMapData is a JSON object with each member representing a single Source PID; the name for a member is the PIDName string identifying the corresponding Source PID. For each Source PID, a DstCosts object denotes the associated cost to a set of destination PIDs

(Section 5.2); the name for each member in the object is the PIDName string identifying the corresponding Destination PID. DstCosts has a single member for each destination PID in the map.

7.8.2.2.3. Example

```
GET /map/core/pid/cost HTTP/1.1
Host: alto.example.com:6671


HTTP/1.1 200 OK
Content-Length: [TODO]
Content-Type: application/alto

{
    "meta" : {
        "version" : 1,
        "status" : {
            "code" : "SUCCESS"
        }
    },
    "type" : "cost-map",
    "data" : {
        "map-vtag"  : "1266506139",
        "cost-type" : "routingcost",
        "cost-mode" : "numerical",
        "map" : {
            "PID1": { "PID1": 1,  "PID2": 5,  "PID3": 10 },
            "PID2": { "PID1": 5,  "PID2": 1,  "PID3": 15 },
            "PID3": { "PID1": 20, "PID2": 15, "PID3": 1  }
        }
    }
}
```

7.8.3. 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.

An ALTO Server MAY support the Map Filtering Service. If an ALTO Server supports the Map Filtering Service, all operations defined in this section MUST be implemented.

7.8.3.1. Network Map

ALTO Clients can query for a subset of the full network map (see Section 7.8.2.1).

7.8.3.1.1. Request Syntax

```
POST /map/filter/pid/net HTTP/1.1
Host: <Host>
Content-Length: <BodyLength>

<ReqNetworkMap>
```

where:

```
object {
    PIDName pids<0..*>;
} ReqNetworkMap;
```

The Body of the request encodes an array of PIDs to be included in the resulting Network Map. If the list of PIDs is empty, the ALTO Server MUST interpret the list as if it contained a list of all currently-defined PIDs.

7.8.3.1.2. Response Syntax

The Response syntax is identical to that of the Map Service's Network Map Response (Section 7.8.2.1.2).

The ALTO Server MUST only include PIDs in the Response that were specified (implicitly or explicitly) in the Request. If the Request contains a PID name that is not currently defined by the ALTO Server, the ALTO Server MUST behave as if the PID did not appear in the request.

7.8.3.1.3. Example

```
POST /map/filter/pid/net HTTP/1.1
Host: alto.example.com:6671
Content-Length: <BodyLength>

{
    pids: [ "PID1", "PID2" ]
}


HTTP/1.1 200 OK
Content-Length: [TODO]
Content-Type: application/alto

{
    "meta" : {
        "version" : 1,
        "status" : {
            "code" : "SUCCESS"
        }
    },
    "type" : "network-map",
    "data" : {
        "map-vtag" : "1266506139",
        "map" : {
            "PID1" : {
                "ipv4" : [
                    "192.0.2.0/24",
                    "198.51.100.0/24"
                ]
            },
            "PID2" : {
                "ipv4": [
                    "198.51.100.128/24"
                ]
            }
        }
    }
}
```

7.8.3.2. Cost Map

ALTO Clients can query for the Cost Map (see Section 7.8.2.2) based on additional parameters.

7.8.3.2.1. Request Syntax

```
POST /map/filter/pid/cost?<URI-Query-String> HTTP/1.1
Host: <Host>

<ReqCostMap>
```

where:

```
object {
    PIDName srcs<0..*>;
    PIDName dsts<0..*>;
} ReqCostMap;
```

The Query String may contain the following parameters:

- type: The requested Cost Type (Section 5.1.1). If not specified, the default value is "routingcost". This parameter MUST NOT be specified multiple times.

- mode: The requested Cost mode (Section 5.1.2). If not specified, the default value is "numerical". This parameter MUST NOT be specified multiple times.

- constraint: Defines a constraint on which elements of the Cost Map are returned. This parameter MUST NOT be used if the Server Capability Response (Section 7.8.1.2) indicates that constraint support is not available. A constraint contains two entities separated by whitespace (before URL encoding): (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 Server Capability Response. If multiple 'constraint' parameters are specified, the ALTO Server assumes they are related to each other with a logical AND. If no 'constraint' parameters are specified, then the ALTO Server returns the full Cost Map.

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 currently-defined PIDs. The ALTO Server returns costs between each pair of source/destination PID. If the Request body is empty, both lists are interpreted to be empty.

7.8.3.2.2. Response Syntax

The Response syntax is identical to that of the Map Service's Cost Map Response (Section 7.8.2.2.2).

The Response MUST NOT contain any source/destination pair that was not indicated (implicitly or explicitly) in the Request. If the Request contains a PID name that is not currently defined by the ALTO Server, the ALTO Server MUST behave as if the PID did not appear in the request.

7.8.3.2.3. Example

```
POST /map/filter/pid/cost?type=hopcount HTTP/1.1
Host: alto.example.com:6671

{
    "srcs" : [ "PID1" ],
    "dsts" : [ "PID1", "PID2", "PID3" ]
}


HTTP/1.1 200 OK
Content-Length: [TODO]
Content-Type: application/alto

{
    "meta" : {
        "version" : 1,
        "status" : {
            "code" : "SUCCESS"
        }
    },
    "type" : "cost-map",
    "data" : {
        "map-vtag"  : "1266506139",
        "cost-type" : "hopcount",
        "cost-mode" : "numerical",
        "map" : {
            "PID1": { "PID1": 0,  "PID2": 1,  "PID3": 2 }
        }
    }
}
```

7.8.4. Endpoint Property Service

The Endpoint Property Lookup query allows an ALTO Client to lookup 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.

An ALTO Server MAY support the Endpoint Property Service. If an ALTO Server supports the Endpoint Property Service, all operations defined

in this section MUST be implemented.

7.8.4.1. Endpoint Property Lookup

7.8.4.1.1. Request Syntax

```
POST /endpoint/prop/lookup?<URI-Query-String> HTTP/1.1
Host: <Host>
Content-Length: <BodyLength>

<ReqEndpointProp>
```

where:

```
object {
    TypedEndpointAddr endpoints<0..*>;
} ReqEndpointProp;
```

The Query String may contain the following parameters:

- prop: The requested property type. This parameter MUST be specified at least once, and MAY be specified multiple times (e.g., to query for multiple different properties at once).

The body encodes a list of typed endpoint addresses.

An alternate syntax is supported for the case when properties are requested for a single endpoint:

```
GET /endpoint/prop/<TypedEndpointAddr>?<URI-Query-String> HTTP/1.1
Host: <Host>
```

where the Query String is the same as in the first form.

7.8.4.1.2. Response Syntax

```
HTTP/1.1 200 <StatusMsg>
Content-Length: <BodyLength>
Content-Type: application/alto

<ALTOResponse>
```

where the ALTOResponse object has "type" member equal to the string "endpoint-property" and "data" member of type RspEndpointProperty:

```
object {
    JSONString [propertyname];
    ...
} EndpointProps;

object {
    EndpointProps  [TypedEndpointAddr]<0..*>;
    ...
} RspEndpointProperty;
```

RspEndpointProperty has one member for each endpoint indicated in the Request (with the name being the endpoint encoded as a TypedEndpointAddr). The requested properties for each endpoint are encoded in a corresponding EndpointProps object, which encodes one name/value pair for each requested property. Note that property values are JSON Strings. If the ALTO Server does not define a requested property for a particular endpoint, then it MUST omit it from the Response for only that endpoint.

7.8.4.1.3. Example

```
POST /endpoint/prop/lookup?prop=pid HTTP/1.1
Host: alto.example.com:6671
Content-Length: [TODO]

{
    "endpoints" : [ "ipv4:192.0.2.34", "ipv4:203.0.113.129" ]
}


HTTP/1.1 200 OK
Content-Length: [TODO]
Content-Type: application/alto

{
    "meta" : {
        "version" : 1,
        "status" : {
            "code" : "SUCCESS"
        }
    },
    "type" : "endpoint-property",
    "data": {
        "ipv4:192.0.2.34"    : { "pid": "PID1" },
        "ipv4:203.0.113.129" : { "pid": "PID3" }
    }
}
```

7.8.5. Endpoint Cost Service

The Endpoint Cost Service allows ALTO Clients to directly supply 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 prefixes (and addresses, as a special case) to be ranked (ordered) by an ALTO Server.

An ALTO Server MAY support the Endpoint Cost Service. If an ALTO Server supports the Endpoint Cost Service, all operations defined in this section MUST be implemented.

7.8.5.1. Endpoint Cost Lookup

7.8.5.1.1. Request Syntax

```
POST /endpoint/cost/lookup?<URI-Query-String> HTTP/1.1
Host: <Host>
Content-Length: <BodyLength>

<ReqEndpointCostMap>
```

where:

```
object {
    TypedEndpointAddr srcs<0..*>;
    TypedEndpointAddr dsts<0..*>;
} ReqEndpointCostMap;
```

The request body includes a list of source and destination endpoints that should be assigned a cost by the ALTO Server. The allowed Query String parameters are defined identically to Section 7.8.3.2.

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.

7.8.5.1.2. Response Syntax

```
HTTP/1.1 200 <StatusMsg>
Content-Length: <BodyLength>
Content-Type: application/alto

<ALTOResponse>
```

where ALTOResponse is encoded identically to Section 7.8.2.2.2 with the following exceptions:

- ALTO Response's "type" member must be equal to "endpoint-cost-map",
- The "map-vtag" member of RspCostMap MUST be omitted, and
- Identifiers refer to TypedEndpointAddres instead of PIDs.

7.8.5.1.3. Example

```
POST /endpoint/cost/lookup?mode=ordinal HTTP/1.1
Host: alto.example.com:6671
Content-Length: [TODO]

{
  "src": [ "ipv4:192.0.2.2" ],
  "dst": [
      "ipv4:192.0.2.89",
      "ipv4:198.51.100.34",
      "ipv4:203.0.113.45"
  ]
}


HTTP/1.1 200 OK
Content-Length: [TODO]
Content-Type: application/alto

{
    "meta" : {
        "version" : 1,
        "status" : {
            "code" : "SUCCESS"
        }
    },
    "type" : "endpoint-cost-map",
    "data" : {
        "cost-type" : "routingcost",
        "cost-mode" : "ordinal",
        "map" : {
            "ipv4:192.0.2.2": {
                "ipv4:192.0.2.89"    : 1,
                "ipv4:198.51.100.34" : 2,
                "ipv4:203.0.113.45"  : 3
            }
        }
    }
}
```

8. Redistributable Responses

This section defines how an ALTO Server enables certain responses to be redistributed by ALTO Clients. Concepts are first introduced, followed by the protocol specification.

8.1. Concepts

8.1.1. Service ID

The Service ID is a UUID that identifies a set of ALTO Servers that would provide identical ALTO Information for any ALTO Request for any ALTO Client. Each ALTO Server within such a set is configured with an identical Service ID.

If a pair of ALTO Servers would provide the same ALTO Information (same information sources, configuration, internal computations, update timescales, etc) in response to a particular ALTO Client request, then the pair of ALTO Servers SHOULD have the same Service ID. If this condition is not true, the pair of ALTO Servers MUST have a different Service ID.

8.1.1.1. Rationale

For scalability and fault tolerance, multiple ALTO Servers may be deployed to serve equivalent ALTO Information. In such a scenario, ALTO Responses from any such redundant server should be seen as equivalent for the purposes of redistribution. For example, if two ALTO Servers A and B are deployed by the service provider to distribute equivalent ALTO Information, then clients contacting Server A should be able to redistribute ALTO Responses to clients contacting Server B.

To accomplish this behavior, ALTO Clients must be able to determine that Server A and Server B serve identical ALTO Information. One technique would be to rely on the ALTO Server's DNS name. However, such an approach would mandate that all ALTO Servers resolved by a particular DNS name would need to provide equivalent ALTO information, which may be unneccessarily restrictive. Another technique would be to rely on the server's IP adddress. However, this suffers similar problems as the DNS name in deployment scenarios using IP Anycast.

To avoid such restrictions, the ALTO Protocol allows an ALTO Service Provider to explicitly denote ALTO Servers that provide equivalent ALTO Information by giving them identical Service IDs. Service IDs decouple the identification of equivalent ALTO Servers from the discovery process.

8.1.1.2. Server Capability Response

If an ALTO Server generates redistributable responses, the Server Capability response's 'service-id' field MUST be set to the ALTO Server's Service ID.

8.1.1.3. Configuration

To help prevent ALTO Servers from mistakenly claiming to distribute equivalent ALTO Information, ALTO Server implementations SHOULD by default generate a new UUID at installation time or startup if one has not explicitly been configured.

8.1.2. Expiration Time

ALTO Responses marked as redistributable should indicate a time after which the information is considered stale and should be refreshed from the ALTO Server (or possibly another ALTO Client).

If an expiration time is present, the ALTO Server SHOULD ensure that it is reasonably consistent with the expiration time that would be computed by HTTP header fields. This specification makes no recommendation on which expiration time takes precedence, but implementers should be cognizant that HTTP intermediaries will obey only the HTTP header fields.

8.1.3. Signature

ALTO Responses marked as redistributable include a signature used to assert that the ALTO Server Provider generated the ALTO Information.

8.1.3.1. Rationale

Verification of the signature requires the ALTO Client to retrieve the ALTO Server's public key. To reduce requirements on the underlying transport (i.e., requiring SSL/TLS), an ALTO Client retrieves the public key as part of an X.509 certificate from the ALTO Server's Server Capability Response.

8.1.3.2. Certificates

8.1.3.2.1. Local Certificate

The ALTO Server's public key is encoded within an X.509 certificate. The corresponding private key MUST be used to sign redistributable responses. This certificate is termed the Local Certificate for an ALTO Server.

8.1.3.2.2. Certificate Chain

To ease key provisioning, the ALTO Protocol is designed such that each ALTO Server with an identical Service ID may have a unique private key (and hence certificate).

The ALTO Service Provider may configure a certificate chain at each such ALTO Server. The Local Certificate for a single ALTO Server is the bottom-most certificate in the chain. The Certificate Chains of each ALTO Server with an identical Service ID MUST share a common Root Certificate.

Note that there are two simple deployment scenarios:

- One-Level Certificate Chain (Local Certificate Only): In this deployment scenario, each ALTO Server with an identical Service ID may provisioned with an identical Local Certificate.

- Two-Level Certificate Chain: In this deployment scenario, a Root Certificate is maintained for a set of ALTO Servers with the same Service ID. A unique Local Certificate signed by this CA is provisioned to each ALTO Server.

There are advantages to using a Certificate Chain instead of deploying the same Local Certificate to each ALTO Server. Specifically, it avoids storage of the CA's private key at ALTO Servers. It is possible to revoke and re-issue a key to a single ALTO Server.

8.1.3.2.3. Server Capability Response

If an ALTO Server generates redistributable responses, the Server Capability response's 'certificates' field MUST be populated with the ALTO Server's full certificate chain. The first element MUST be the ALTO Server's Local Certificate, followed by the remaining Certificate Chain in ascending order to the Root Certificate.

8.1.3.3. Signature Verification

ALTO Clients SHOULD verify the signature on any ALTO information received via redistribution before adjusting application behavior based on it.

An ALTO Client SHOULD cache its ALTO Server's Service ID and corresponding Certificate Chain included in the Server Capability response. Recall that the last certificate in this chain is the Root Certificate. The retrieval of the Service ID and certificates SHOULD be secured using HTTPS with proper validation of the server endpoint of the SSL/TLS connection [6].

An ALTO Response received via redistribution from Service ID S is declared valid if an ALTO Client can construct a transitive certificate chain from the certificate (public key) used to sign the ALTO Response to the Root Certificate corresponding to Service ID S

obtained by the ALTO Client in a Server Capability response.

To properly construct the chain and complete this validation, an ALTO Client may need to request additional certificates from other ALTO Clients. A simple mechanism is to request the certificate chain from the ALTO Client that received the ALTO Response. Note that these additional received certificates may be cached locally by an ALTO CLient.

ALTO Clients SHOULD verify ALTO Responses received via redistribution.

8.1.3.4. Redistribution by ALTO Clients

ALTO Clients SHOULD pass the ALTO Server Certificate, Signature, and Signature Algorithm along with the body of the ALTO Response. The mechanism for redistributing such information is not specified by the ALTO Protocol, but one possibility is to add additional messages or fields to the application's native protocol.

8.2. Protocol

An ALTO Server MAY indicate that a response is suitable for redistribution by including the "redistribution" member in the RspMetaData JSON object of an ALTO Response message. This additional member, called the Response Redistribution Descriptor, has type RspRedistDesc:

```
object {
    JSONString service-id;
    JSONString request-uri;
    JSONValue  request-body;
    JSONString expires;
} RspRedistDesc;
```

The fields encoded in the Response Redistribution Descriptor allows an ALTO Client receiving redistributed ALTO Information to understand the context of the query (the ALTO Service generating the response and any input parameters) and to interpret the results.

Information about ALTO Client performing the Request and any HTTP Headers passed in the request are not included in the Response Redistribution Descriptor. If any such information or headers influence the response generated by the ALTO Server, the response SHOULD NOT be indicated as redistributable.

8.2.1. Response Redistribution Descriptor Fields

This section defines the fields of the Response Redistribution Descriptor.

8.2.1.1. Service ID

The 'service-id' member is REQUIRED and MUST have a value equal to the ALTO Server's Service ID.

8.2.1.2. Request URI

The 'request-uri' member is REQUIRED and MUST specify the HTTP Request-URI that was passed in the HTTP Request.

8.2.1.3. Request Body

If the HTTP Request body was non-empty, the 'request-body' member MUST specify full JSON value passed in the HTTP Request (note that whitespace may differ, as long as the JSON Value is identical). If the HTTP Request was empty, then the 'request-body' MUST NOT be included.

8.2.1.4. Expiration Time

The 'expires' element is RECOMMENDED and, if present, MUST specify a time in UTC formatted according to [12].

8.2.2. Signature

The Hash Algorithm, Signature Algorithm, and Signature are included as either HTTP Headers or Trailers. Headers may be useful if Responses are pre-generated, while Trailers may be useful if Responses are dynamically generated (e.g., to avoid buffering large responses in memory while the hash value is computed).

The following HTTP Headers (the ALTO Server MAY specify them as HTTP Trailers instead) MUST be used to encode the Signature parameters for redistributable ALTO Responses:

```
ALTO-HashAlgorithm: <HashAlgorithm>
ALTO-SignatureAlgorithm: <SignatureAlgorithm>
ALTO-SignatureDigest: <Signature>
```

where <HashAlgorithm> and <SignatureAlgorithm> are an integer values from the IANA TLS HashAlgorithm and SignatureAlgorithm registries, and <Signature> is the corresponding Base64-encoded signature.

9. Use Cases

The sections below depict typical use cases.

9.1. 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.

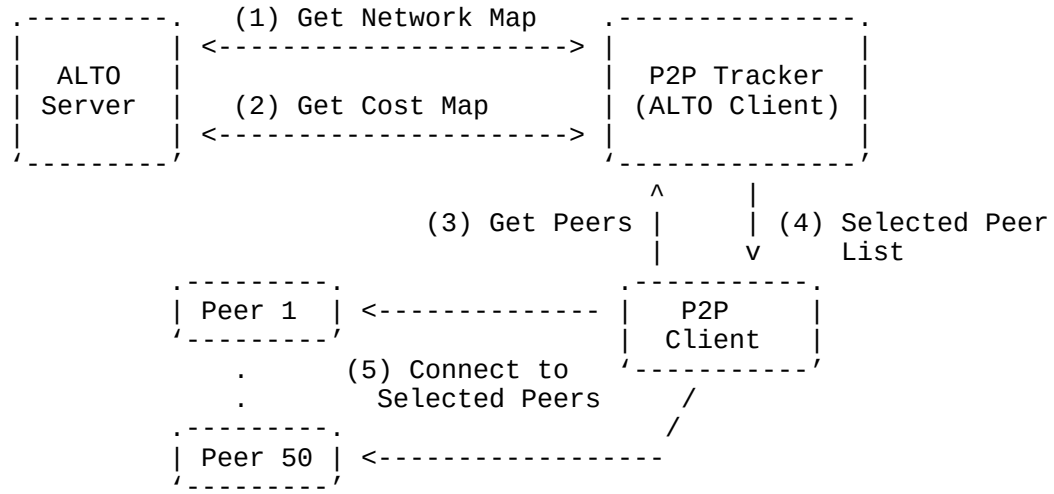

Figure 4: ALTO Client Embedded in P2P Tracker

Figure 4 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 Network Map 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.

9.2. 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.

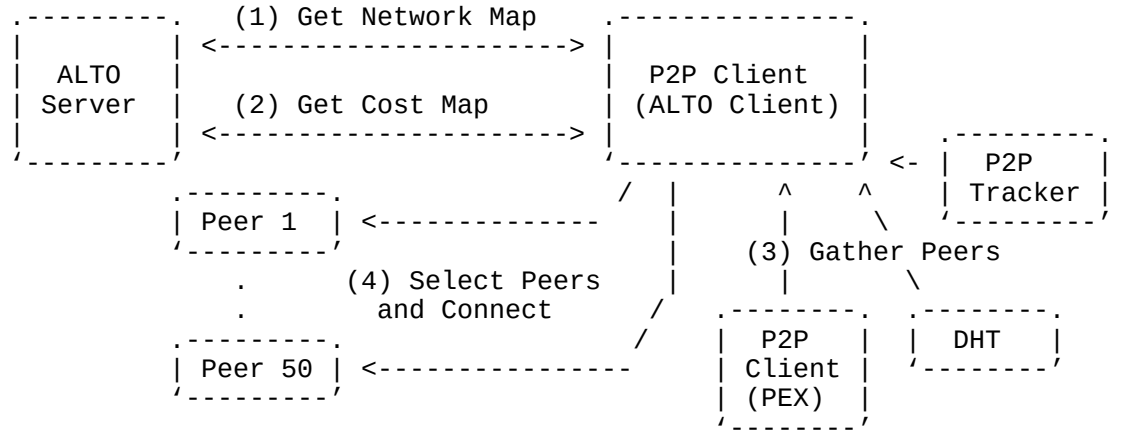

Figure 5: ALTO Client Embedded in P2P Client

Figure 5 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.

9.3. 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.

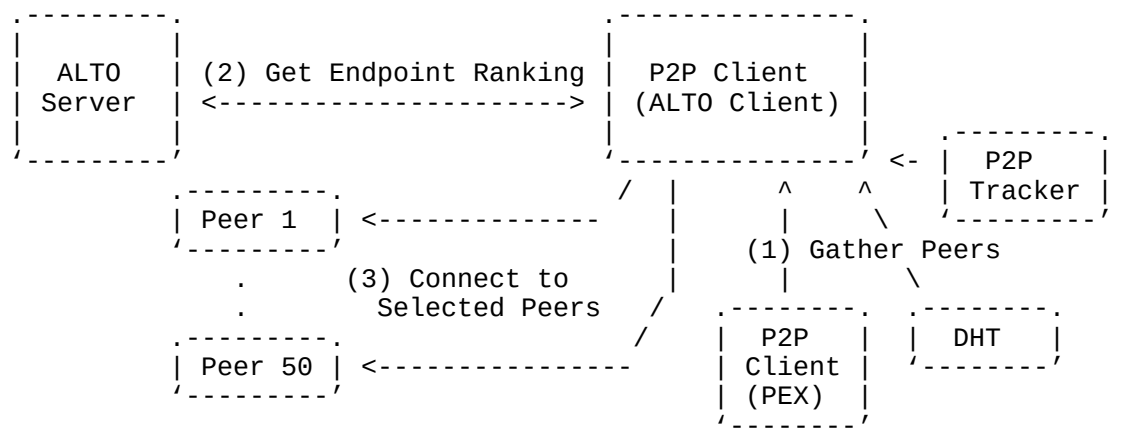

Figure 6: ALTO Client Embedded in P2P Client: Ranking

Figure 6 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.

10. Discussions

10.1. Discovery

The discovery mechanism by which an ALTO Client locates an appropriate ALTO Server is out of scope for this document. This document assumes that an ALTO Client can discover an appropriate ALTO Server. Once it has done so, the ALTO Client may use the Server List query Section 7.8.1.1 to locate an ALTO Server with capabilities necessary for its application.

10.2. Hosts with Multiple Endpoint Addresses

In practical deployments, especially during the transition from IPv4 to IPv6, a particular host may be reachable using multiple addresses. Furthermore, the particular network path followed when sending packets to the host may differ based on the address that is used. Network providers may perfer one path over another (e.g., one path my have a NAT64 middlebox). An additional consideration may be how to handle private address spaces (e.g., behind carrier-grade NATs).

To support such behavior, this document allows multiple types of endpoint addresses. In supporting multiple address types, the ALTO Protocol also allows ALTO Service Provider the flexibility to indicate preferences for paths from an endpoint address of one type to an endpoint address of a different type. Note that in general, the path through the network may differ dependent on the types of addresses that are used (one such example is DS-Lite).

Note that there are limitations as to what information ALTO can provide in this regard. In particular, a particular ALTO Service provider may not be able to determine if connectivity with a particular endhost will succeed over IPv4 or IPv6, as this may depend upon information unknown to the ISP such as particular application implementations.

10.3. Network Address Translation Considerations

At this day and age of NAT v4<->v4, v4<->v6 [27], and possibly v6<->v6[28], 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 network location 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 [29]) operate.

The ALTO client SHOULD use the Session Traversal Utilities for NAT (STUN) [13] to determine a public IP address to use as a source Endpoint address. 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.

10.4. 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, the Network Map query in the Map Filtering Service 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.

10.5. 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. This specification focuses on network location and grouping, but the protocol may be extended to handle other Endpoint properties.

10.6. REST-ful Protocol Structure

There is an ongoing discussion as to whether the ALTO Protocol should be restructured to be REST-ful. The discussion has been captured at http://www.ietf.org/mail-archive/web/alto/current/msg00792.html and the ensuing thread.

Three possible paths forward for the ALTO Protocol are:

1. Keep the ALTO Protocol as it is;

2. Restructure this document to allow a REST-ful protocol as an extension, while keeping the protocol unchanged;

3. Restructure the protocol.

This should be resolved by the ALTO Working Group before the next revision of this draft.

11. IANA Considerations

11.1. application/alto Media Type

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

Type name: application

Subtype name: alto

Required parameters: n/a

Optional parameters: n/a

Encoding considerations: Encoding considerations are identical to those specified for the 'application/json' media type. See [4].

Security considerations: Security considerations relating to the generation and consumption of ALTO protocol messages are discussed in Section 12.

Interoperability considerations: This document specifies format of conforming messages and the interpretation thereof.

Published specification: This document.

Applications that use this media type: ALTO Servers and ALTO Clients either standalone or embedded within other applications.

Additional information:

Magic number(s): n/a

File extension(s): This document uses the mime type to refer to protocol messages and thus does not require a file extension.

Macintosh file type code(s): n/a

Person & email address to contact for further information: See "Authors' Addresses" section.

Intended usage: COMMON

Restrictions on usage: n/a

Author: See "Authors' Addresses" section.

Change controller: See "Authors' Addresses" section.

11.2. ALTO Cost Type Registry

This document requests the creation of an ALTO Cost Type registry to be maintained by IANA.

This registry serves two purposes. First, it ensures uniqueness of identifiers referring to ALTO Cost Types. Second, it provides references to particular semantics of allocated Cost Types to be applied by both ALTO Servers and applications utilizing ALTO Clients.

New ALTO Cost Types are assigned after Expert Review [10]. The Expert Reviewer will generally consult the ALTO Working Group or its successor. Expert Review is used to ensure that proper documentation regarding ALTO Cost Type semantics and security considerations has been provided. The provided documentation should be detailed enough to provide guidance to both ALTO Service Providers and applications utilizing ALTO Clients as to how values of the registered ALTO Cost Type should be interpreted. Updates and deletions of ALTO Cost Types follow the same procedure.

Registered ALTO Cost Type identifiers MUST conform to the syntatical requirements specified in Section 7.7.4. Identifiers are to be recorded and displayed as ASCII strings.

Identifiers prefixed with 'priv:' are reserved for Private Use. Identifiers prefixed with 'exp:' are reserved for Experimental use.

Requests to add a new value to the registry MUST include the following information:

- Identifier: The name of the desired ALTO Cost Type.
- Intended Semantics: ALTO Costs carry with them semantics to guide their usage by ALTO Clients. For example, if a value refers to a measurement, the measurement units must be documented. For proper implementation of the ordinal Cost Mode (e.g., by a third-party service), it should be documented whether higher or lower values of the cost are more preferred.

- Security Considerations: ALTO Costs expose information to ALTO Clients. As such, proper usage of a particular Cost Type may require certain information to be exposed by an ALTO Service Provider. Since network information is frequently regarded as proprietary or confidential, ALTO Service Providers should be made aware of the security ramifications related to usage of a Cost Type.

This specification requests registration of the identifier 'routingcost'. Semantics for the this Cost Type are documented in Section 5.1.1.1, and security considerations are documented in Section 12.1.

12. Security Considerations

12.1. Privacy Considerations for 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. On the one hand, providing overly fine-grained information may make it easier for attackers to infer network topology. In particular, attackers may try to infer details regarding ISPs' operational policies or inter-ISP business relationships by intentionally posting a multitude of selective queries to an ALTO server and analyzing the responses. Such sophisticated attacks may reveal more information than an ISP hosting an ALTO server intends to disclose. On the other hand, revealing overly coarse-grained information may not provide benefits to network efficiency or performance improvements to ALTO Clients.

12.2. ALTO Clients

Applications using the information must be cognizant of the possibility that the information is malformed or incorrect. Even if an ALTO Server has been properly authenticated by the ALTO Client, the information provided may be malicious because the ALTO Server and its credentials have been compromised (e.g., through malware). Other considerations (e.g., relating to application performance) can be found in Section 6 of [22].

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.

In addition, ALTO clients should be cautious not to unintentionally or indirectly disclose the resource identifier (of which they try to improve the retrieval through ALTO-guidance), e.g., the name/ identifier of a certain video stream in P2P live streaming, to the ALTO server. Note that the ALTO Protocol specified in this document does not explicitly reveal any resource identifier to the ALTO Server. However, for instance, depending on the popularity or other specifics (such as language) of the resource, an ALTO server could potentially deduce information about the desired resource from information such as the Network Locations the client sends as part of its request to the server.

12.3. Authentication, Integrity Protection, and Encryption

SSL/TLS can provide encryption of transmitted messages as well as authentication of the ALTO 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).

An ALTO Server may optionally use authentication (and potentially encryption) to protect ALTO information it provides. This can be achieved by digitally signing a hash of the ALTO information itself and attaching the signature to the ALTO information. There may be special use cases where encryption of ALTO information is desirable. In many cases, however, information sent out by an ALTO Server may be regarded as non-confidential information.

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.

12.4. 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 generated by an appropriate 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. Contacting an ALTO server for information validation would also undermine the intended effect of redistribution and is therefore not desirable.

Note that the redistribution scheme must additionally handle details

such as ensuring ALTO Clients retrieve ALTO information from the correct ALTO Server. See [25] 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 with its private key. The corresponding public key is included in the Server Capability response Section 7.8.1.2, along with the certificate chain to a Root Certificate generated by the ALTO Service Provider. To prevent man-in-the-middle attacks, an ALTO Client SHOULD perform the Server Capability Query over SSL/TLS and verify the server identity according to [6].

The signature verification algorithm is detailed in Section 8.1.3.3.

12.5. 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. In particular, queries which can be generated with low effort but result in expensive workloads at the ALTO Server could be exploited for Denial-of-Service attacks. For instance, a simple ALTO query with n Source Network Locations and m Destination Network Locations can be generated fairly easily but results in the computation of n*m Path Costs between pairs by the ALTO Server (see Section 5.2). One way to limit Denial-of-Service attacks is to employ access control to the ALTO server. Another possible mechanism for an ALTO Server to protect itself against a multitude of computationally expensive bogus requests is to demand that each ALTO Client to solve a computational puzzle first before allocating resources for answering a request (see, e.g., [30]). The current specification does not use such computational puzzles, and discussion regarding tradeoffs of such an approach would be needed before including such a technique in the ALTO Protocol.

ISPs should also leverage the fact that the the Map Service allows ALTO Servers to pre-generate maps that can be useful to many ALTO Clients.

12.6. ALTO Server Access Control

In order to limit access to an ALTO server (e.g., for an ISP to only allow its users to access its ALTO server, or to prevent Denial-of-Service attacks by arbitrary hosts from the Internet), an ALTO server may employ access control policies. Depending on the use-case and scenario, an ALTO server may restrict access to its services more

strictly or rather openly (see [31] for a more detailed discussion on this issue).

13. References

13.1. 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] Crockford, D., "The application/json Media Type for JavaScript Object Notation (JSON)", RFC 4627, July 2006.

[5] Blake-Wilson, S., Nystrom, M., Hopwood, D., Mikkelsen, J., and T. Wright, "Transport Layer Security (TLS) Extensions", RFC 4366, April 2006.

[6] Saint-Andre, P. and J. Hodges, "Representation and Verification of Domain-Based Application Service Identity within Internet Public Key Infrastructure Using X.509 (PKIX) Certificates in the Context of Transport Layer Security (TLS)", draft-saintandre-tls-server-id-check-14 (work in progress), January 2011.

[7] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, January 2005.

[8] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, February 2006.

[9] Fuller, V. and T. Li, "Classless Inter-domain Routing (CIDR): The Internet Address Assignment and Aggregation Plan", BCP 122, RFC 4632, August 2006.

[10] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.

[11] Leach, P., Mealling, M., and R. Salz, "A Universally Unique IDentifier (UUID) URN Namespace", RFC 4122, July 2005.

[12] Klyne, G., Ed. and C. Newman, "Date and Time on the Internet: Timestamps", RFC 3339, July 2002.

[13] 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.

13.2. Informative References

[14] 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.

[15] 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.

[16] H. Xie, YR. Yang, A. Krishnamurthy, Y. Liu, and A. Silberschatz., "P4P: Provider Portal for (P2P) Applications", In SIGCOMM 2008.

[17] 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.

[18] Shalunov, S., Penno, R., and R. Woundy, "ALTO Information Export Service", draft-shalunov-alto-infoexport-00 (work in progress), October 2008.

[19] Das, S. and V. Narayanan, "A Client to Service Query Response Protocol for ALTO", draft-saumitra-alto-queryresponse-00 (work in progress), March 2009.

[20] 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.

[21] Akonjang, O., Feldmann, A., Previdi, S., Davie, B., and D. Saucez, "The PROXIDOR Service", draft-akonjang-alto-proxidor-00 (work in progress), March 2009.

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

[23] 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.

[24] Zyp, K. and G. Court, "A JSON Media Type for Describing the Structure and Meaning of JSON Documents", draft-zyp-json-schema-03 (work in progress), November 2010.

[25] Yingjie, G., Alimi, R., and R. Even, "ALTO Information Redistribution", draft-gu-alto-redistribution-03 (work in progress), July 2010.

[26] 3rd, D., "Transport Layer Security (TLS) Extensions: Extension Definitions", draft-ietf-tls-rfc4366-bis-12 (work in progress), September 2010.

[27] Baker, F., Li, X., and C. Bao, "Framework for IPv4/IPv6 Translation", draft-baker-behave-v4v6-framework-02 (work in progress), February 2009.

[28] Wasserman, M. and F. Baker, "IPv6-to-IPv6 Network Address Translation (NAT66)", draft-mrw-behave-nat66-02 (work in progress), March 2009.

[29] "Bittorrent Protocol Specification v1.0", http://wiki.theory.org/BitTorrentSpecification, 2009.

[30] Jennings, C., "Computational Puzzles for SPAM Reduction in SIP", draft-jennings-sip-hashcash-06 (work in progress), July 2007.

[31] Stiemerling, M. and S. Kiesel, "ALTO Deployment Considerations", draft-stiemerling-alto-deployments-06 (work in progress), January 2011.

Appendix A. Acknowledgments

Thank you to Jan Seedorf for contributions to the Security Considerations section. We would like to thank Yingjie Gu and Roni Even for helpful input and design concerning ALTO Information redistribution.

We would like to thank the following people whose input and involvement was indispensable in achieving this merged proposal:

Obi Akonjang (DT Labs/TU Berlin),

Saumitra M. Das (Qualcomm Inc.),

Syon Ding (China Telecom),

Doug Pasko (Verizon),

Laird Popkin (Pando Networks),

Satish Raghunath (Juniper Networks),

Albert Tian (Ericsson/Redback),

Yu-Shun Wang (Microsoft),

David Zhang (PPLive),

Yunfei Zhang (China Mobile).

We would also 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 Poese (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).

Appendix B. Authors

[[Comment.1: RFC Editor: Please move information in this section to the Authors' Addresses section at publication time.]]

Stefano Previdi
Cisco

Email: sprevidi@cisco.com

Stanislav Shalunov
BitTorrent

Email: shalunov@bittorrent.com

Richard Woundy
Comcast

Richard_Woundy@cable.comcast.com

Authors' Addresses

Richard Alimi (editor)
Google
1600 Amphitheatre Parkway
Mountain View CA
USA

Email: ralimi@google.com

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

Email: rpenno@juniper.net

Y. Richard Yang (editor)
Yale University
51 Prospect St
New Haven CT
USA

Email: yry@cs.yale.edu