ALTO
Internet-Draft
Intended status: Informational
Expires: September 3, 2015

M. Stiemerling
NEC Europe Ltd.
S. Kiesel
University of Stuttgart
S. Previdi
Cisco
M. Scharf
Alcatel-Lucent Bell Labs

March 2, 2015

ALTO Deployment Considerations
draft-ietf-alto-deployments-11

Abstract

Many Internet applications are used to access resources such as pieces of information or server processes that are available in several equivalent replicas on different hosts. This includes, but is not limited to, peer-to-peer file sharing applications. The goal of Application-Layer Traffic Optimization (ALTO) is to provide guidance to applications that have to select one or several hosts from a set of candidates, which are able to provide a desired resource. This memo discusses deployment related issues of ALTO. It addresses different use cases of ALTO such as peer-to-peer file sharing and CDNs and presents corresponding examples. The document also includes recommendations for network administrators and application designers planning to deploy ALTO.

Status of This Memo

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

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at http://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on September 3, 2015.

Table of Contents

1. Introduction

Many Internet applications are used to access resources such as pieces of information or server processes that are available in several equivalent replicas on different hosts. This includes, but is not limited to, peer-to-peer (P2P) file sharing applications and Content Delivery Networks (CDNs). The goal of Application-Layer Traffic Optimization (ALTO) is to provide guidance to applications that have to select one or several hosts from a set of candidates, which are able to provide a desired resource. The basic ideas and problem space of ALTO is described in [RFC5693] and the set of requirements is discussed in [RFC6708]. The ALTO protocol is

specified in [RFC7285]. An ALTO server discovery procedure is defined in [RFC7286].

This document discusses use cases and operational issues that can be expected when ALTO gets deployed. This includes, but is not limited to, location of the ALTO server, imposed load to the ALTO server, or from whom the queries are performed. The document also provides guidance which ALTO services to use, and it summarizes known challenges. It thereby complements the management considerations in the protocol specification [RFC7285], which are independent of any specific use of ALTO.

2. General Considerations

2.1. ALTO Entities

2.1.1. Baseline Scenario

The ALTO protocol [RFC7285] is a client/server protocol, operating between a number of ALTO clients and an ALTO server, as sketched in Figure 1.

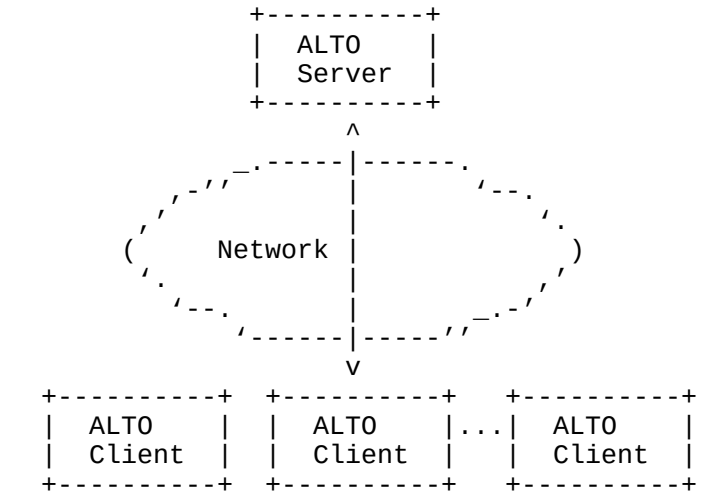

Figure 1: Baseline deployment scenario of the ALTO protocol

This document uses the terminology introduced in [RFC5693]. In particular, the following terms are defined by [RFC5693]:

- ALTO Service: Several resource providers may be able to provide the same resource. The ALTO service gives guidance to a resource consumer and/or resource directory about which resource

provider(s) to select in order to optimize the client's performance or quality of experience, while improving resource consumption in the underlying network infrastructure.

- o ALTO Server: A logical entity that provides interfaces to the queries to the ALTO service.

- o ALTO Client: The logical entity that sends ALTO queries. Depending on the architecture of the application, one may embed it in the resource consumer and/or in the resource directory.

According to that definition, both an ALTO server and an ALTO client are logical entities. An ALTO service may be offered by more than one ALTO servers. In ALTO deployments, the functionality of an ALTO server can therefore be realized by several server instances, e.g., by using load balancing between different physical servers. The term ALTO server should not be confused with use of a single physical server.

2.1.2. Placement of ALTO Entities

The ALTO server and ALTO clients can be situated at various entities in a network deployment. The first differentiation is whether the ALTO client is located on the actual host that runs the application, as shown in Figure 2, or if the ALTO client is located on a resource directory, as shown in Figure 3.

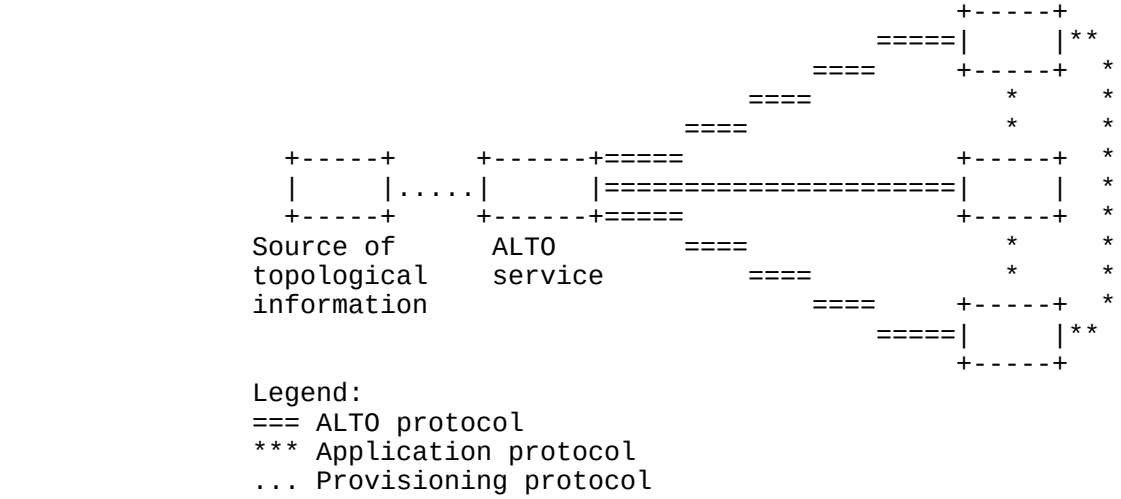

Figure 2: Overview of protocol interaction between ALTO elements without a resource directory

Figure 2 shows the operational model for an ALTO client running at endpoints. An example would be a peer-to-peer file sharing application that does not use a tracker, such as edonkey. In addition, ALTO clients at peers could also be used in a similar way even if there is a tracker, as further discussed in Section 4.1.2.

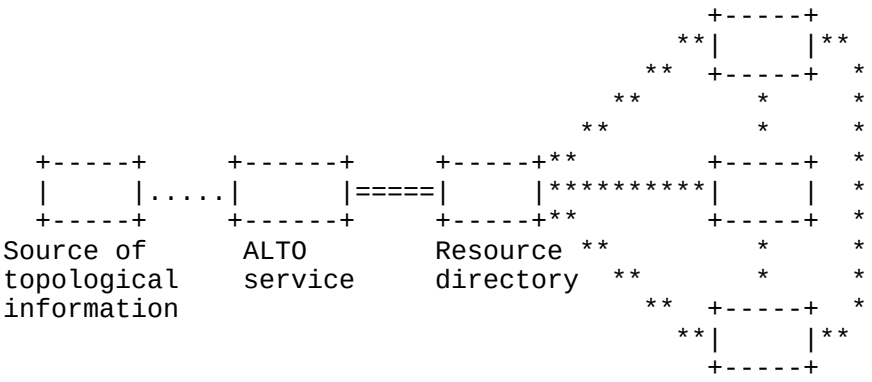

Legend:
=== ALTO protocol
*** Application protocol
... Provisioning protocol

Figure 3: Overview of protocol interaction between ALTO elements with a resource directory

In Figure 3, a use case with a resource directory is illustrated, e.g., a tracker in peer-to-peer file-sharing. Both deployment scenarios may differ in the number of ALTO clients that access an ALTO service: If ALTO clients are implemented in a resource directory, ALTO servers may be accessed by a limited and less dynamic set of clients, whereas in the general case any host could be an ALTO client. This use case is further detailed in Section 4.

Using ALTO in CDNs may be similar to a resource directory [I-D.jenkins-alto-cdn-use-cases]. The ALTO server can also be queried by CDN entities to get guidance about where the a particular client accessing data in the CDN is exactly located in the Internet Service Provider's network, as discussed in Section 5.

2.2. Classification of Deployment Scenarios

2.2.1. Roles in ALTO Deployments

ALTO is a general-purpose protocol and it is intended to be used by a wide range of applications. This implies that there are different possibilities where the ALTO entities are actually located, i.e., if the ALTO clients and the ALTO server are in the same Internet Service Provider (ISP) domain, or if the clients and the ALTO server are managed/owned/located in different domains.

An ALTO service includes four types of entities:

1. Source of topological information
2. ALTO server
3. ALTO client
4. Resource consumer (using the ALTO guidance)

Each of these entities corresponds to a certain role, which results in requirements and constraints on the interaction between the entities.

A key design objective of the ALTO service is that each these four roles can be separated, i.e., they can be realized by different organizations or disjoint system components. ALTO is inherently designed for use in multi-domain environments. Most importantly, ALTO is designed to enable deployments in which the ALTO server and the ALTO client are not located within the same administrative domain.

As explained in [RFC5693], from this follows that at least three different kinds of entities can operate an ALTO server:

1. Network operators. Network Service Providers (NSPs) such as Internet Service Providers (ISPs) may have detailed knowledge of their network topology and policies. In this case, the source of the topology information and the provider of the ALTO server may be part of the same organization.
2. Third parties. Topology information could also be collected by entities separate from network operators but that may either have collected network information or have arrangements with network operators to learn the network information. Examples of such entities could be Content Delivery Network (CDN) operators or companies specialized on offering ALTO services on behalf of ISPs.

3. User communities. User communities could run distributed measurements for estimating the topology of the Internet. In this case the topology information may not originate from ISP data.

Regarding the interaction between ALTO server and client, ALTO deployments can be differentiated e.g. according to the following aspects:

1. Applicable trust model: The deployment of ALTO can differ depending on whether ALTO client and ALTO server are operated within the same organization and/or network, or not. This affects a lot of constraints, because the trust model is very different. For instance, as discussed later in this memo, the level-of-detail of maps can depend on who the involved parties actually are.

2. Size of user group: The main use case of ALTO is to provide guidance to any Internet application. However, an operator of an ALTO server could also decide to only offer guidance to a set of well-known ALTO clients, e. g., after authentication and authorization. In the peer-to-peer application use case, this could imply that only selected trackers are allowed to access the ALTO server. The security implications of using ALTO in closed groups differ from the public Internet.

3. Covered destinations: In general, an ALTO server has to be able to provide guidance for all potential destinations. Yet, in practice a given ALTO client may only be interested in a subset of destinations, e.g., only in the network cost between a limited set of resource providers. For instance, CDN optimization may not need the full ALTO cost maps, because traffic between individual residential users is not in scope. This may imply that an ALTO server only has to provide the costs that matter for a given user, e. g., by customized maps.

The following sections enumerate different classes of use cases for ALTO, and they discuss deployment implications of each of them. An ALTO server can in principle be operated by any organization, and there is no requirement that an ALTO server is deployed and operated by ISPs. Yet, since the ALTO solution is designed for ISPs, most examples in this document assume that the operator of an ALTO server is a network operator (e.g., an ISP or the network department in a large enterprise) that offers ALTO guidance in particular to users if this network.

It must be emphasized that any application using ALTO must also work if no ALTO servers can be found or if no responses to ALTO queries

are received, e.g., due to connectivity problems or overload situations (see also [RFC6708]).

2.2.2. Information Exposure

An ALTO server stores information about preferences (e.g., for IP address ranges) and ALTO clients can retrieve these preferences. There are basically two different approaches on where the preferences are actually processed:

1. The ALTO server has a list of preferences and clients can retrieve this list via the ALTO protocol. This preference list can partially be updated by the server. The actual processing of the data is done on the client and thus there is no data of the client's operation revealed to the ALTO server.

2. The ALTO server has a list of preferences or preferences calculated during runtime and the ALTO client is sending information of its operation (e.g., a list of IP addresses) to the server. The server is using this operational information to determine its preferences and returns these preferences (e.g., a sorted list of the IP addresses) back to the ALTO client.

Approach 1 has the advantage (seen from the client) that all operational information stays within the client and is not revealed to the provider of the server. On the other hand, approach 1 requires that the provider of the ALTO server, i.e., the network operator, reveals information about its network structure (e.g., IP ranges or topology information in general) to the ALTO client. The ALTO protocol supports this scheme by the Network and Cost Map Service.

Approach 2 has the advantage (seen from the operator) that all operational information stays with the ALTO server and is not revealed to the ALTO client. On the other hand, approach 2 requires that the clients send their operational information to the server. This approach is realized by the ALTO Endpoint Cost Service (ECS).

Both approaches have their pros and cons, as further detailed in Section 3.3.

2.2.3. More Advanced Deployments

From an ALTO client's perspective, there are different ways to use ALTO:

1. Single service instance with single metric guidance: An ALTO client only obtains guidance regarding a single metric from a

single ALTO service, e.g., an ALTO server that is offered by the network service provider of the corresponding access network. Corresponding ALTO server instances can be discovered e.g. by ALTO server discovery [RFC7286] [I-D.kiesel-alto-xdom-disc]. Being a REST-ful protocol, an ALTO service can use known methods to balance the load between different server instances or between clusters of servers, i.e., an ALTO server can be realized by many instances with a load balancing scheme. The ALTO protocol also supports the use of different URIs for different ALTO features.

2. Single service instance with multiple metric guidance: An ALTO client could also query an ALTO service for different kinds of information, e.g., cost maps with different metrics. The ALTO protocol is extensible and permits such operation. However, ALTO does not define how a client shall deal with different forms of guidance, and it is up to the client to determine what provided information may indeed be useful.

3. Multiple service offers: An ALTO client can also decide to access multiple ALTO servers providing guidance, possibly from different operators or organizations. Each of these services may only offer partial guidance, e.g., for a certain network partition. In that case, it may be difficult for an ALTO client to compare the guidance from different services. Different organization may use different methods to determine maps, and they may also have different (possibly even contradicting or competing) guidance objectives. How to discover multiple ALTO servers and how to deal with conflicting guidance is an open issue.

There are also different options regarding the guidance offered by an ALTO service:

1. Authoritative servers: An ALTO server instance can provide guidance for all destinations for all kinds of ALTO clients.

2. Cascaded servers: An ALTO server may itself include an ALTO client and query other ALTO servers, e.g., for certain destinations. This results is a cascaded deployment of ALTO servers, as further explained below.

3. Inter-server synchronization: Different ALTO servers my communicate by other means. This approach is not further discussed in this document.

An assumption of the ALTO design is that ISP operate ALTO servers independently, irrespectively of other ISPs. This may true for most envisioned deployments of ALTO but there may be certain deployments that may have different settings. Figure 4 shows such setting with a

university network that is connected to two upstream providers. NREN is a National Research and Education Network and ISP is a commercial upstream provider to this university network. The university, as well as ISP, are operating their own ALTO server. The ALTO clients, located on the peers will contact the ALTO server located at the university.

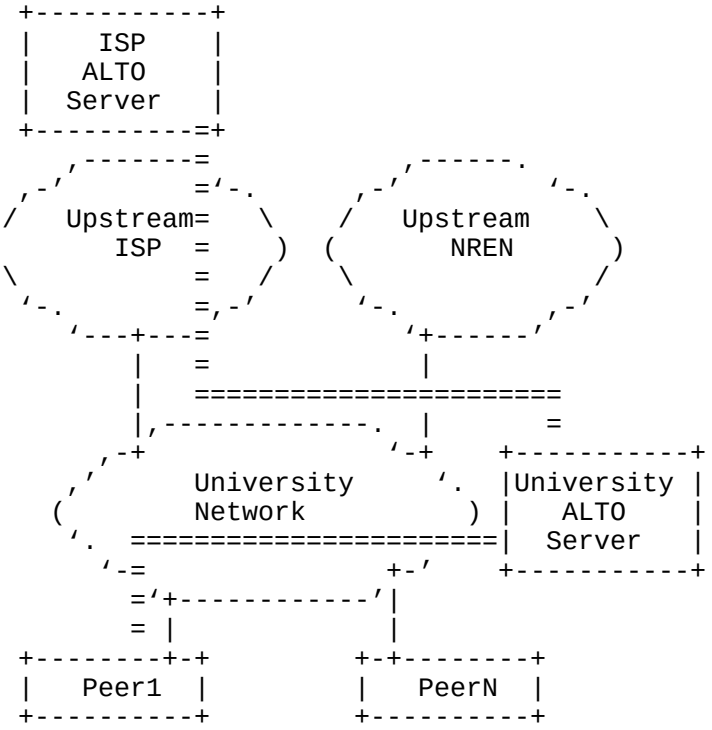

Figure 4: Example of a cascaded ALTO server

In this setting all "destinations" useful for the peers within NREN are free-of-charge for the peers located in the university network (i.e., they are preferred in the rating of the ALTO server). However, all traffic that is not towards NREN will be handled by the ISP upstream provider. Therefore, the ALTO server at the university may also include the guidance given by the ISP ALTO server in its replies to the ALTO clients. This is an example for cascaded ALTO servers.

3. Deployment Considerations by ISPs

3.1. Objectives for the Guidance to Applications

3.1.1. General Objectives for Traffic Optimization

The Internet consists of many networks. The networks are operated by Network Service Providers (NSP) or Internet Service Providers (ISP), which also include e.g. universities, enterprises, or other organizations. The Internet provides network connectivity e.g. by access networks, such as cable networks, xDSL networks, 3G/4G mobile networks, etc. Network operators need to manage, to control and to audit the traffic. Therefore, it is important to understand how to deploy an ALTO service and its expected impact.

The general objective of ALTO is to give guidance to applications on what endpoints (e.g., IP addresses or IP prefixes) are to be preferred according to the operator of the ALTO server. The ALTO protocol gives means to let the ALTO server operator express its preference, whatever this preference is.

ALTO enables ISPs to support application-level traffic engineering by influencing application resource selections. This traffic engineering for overlay formed by the application can have different objectives:

1. Inter-network traffic localization: ALTO can help to reduce inter-domain traffic. The networks of ISPs are connected through peering points. From a business view, the inter-network settlement is needed for exchanging traffic between these networks. These peering agreements can be costly. To reduce these costs, a simple objective is to decrease the traffic exchange across the peering points and thus keep the traffic in the own network or Autonomous System (AS) as far as possible.

2. Intra-network traffic localization: In case of large ISPs, the network may be grouped into several networks, domains, or Autonomous Systems (ASs). The core network includes one or several backbone networks, which are connected to multiple aggregation, metro, and access networks. If traffic can be limited to certain areas such as access networks, this decreases the usage of backbone and thus helps to save resources and costs.

3. Network off-loading: Compared to fixed networks, mobile networks have some special characteristics, including smaller link bandwidth, high cost, limited radio frequency resource, and limited terminal battery. In mobile networks, wireless links should be used efficiently. For example, in the case of a P2P

service, it is likely that hosts in fixed networks should avoid retrieving data from hosts in mobile networks, and hosts in mobile networks should prefer retrieval of data from hosts in fixed networks.

4. Application tuning: ALTO is also a tool to optimize the performance of applications that depend on the network and perform resource selection decisions among network endpoints. And example is the network-aware selection of Content Delivery Network (CDN) caches.

In the following, these objectives are explained in more detail with examples.

3.1.2. Inter-Network Traffic Localization

ALTO guidance can be used to keep traffic local in a network. An ALTO server can let applications prefer other hosts within the same network operator’s network instead of randomly connecting to other hosts that are located in another operator’s network. Here, a network operator would always express its preference for hosts in its own network, while hosts located outside its own network are to be avoided (i.e., they are undesired to be considered by the applications). Figure 5 shows such a scenario where hosts prefer hosts in the same network (e.g., Host 1 and Host 2 in ISP1 and Host 3 and Host 4 in ISP2).

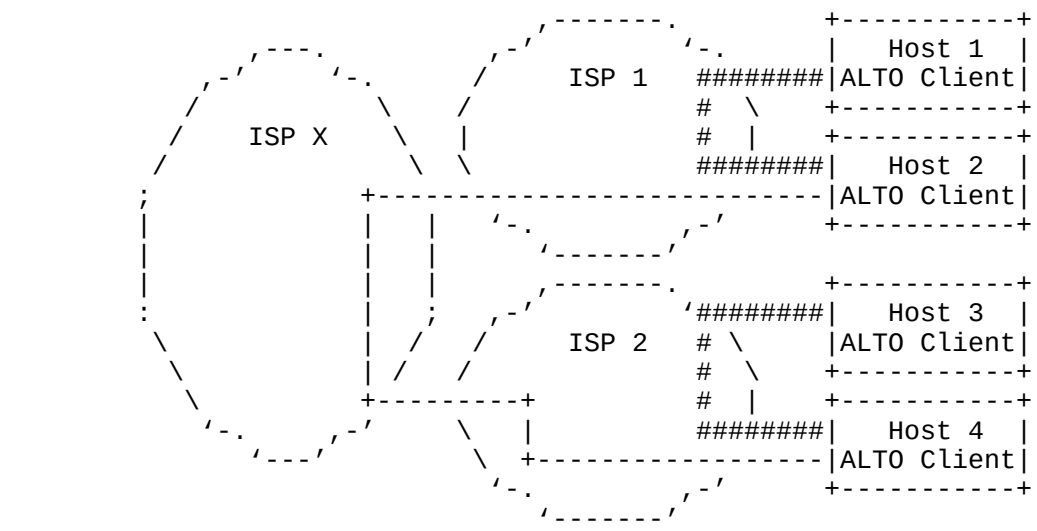

Figure 5: Inter-network traffic localization

Examples for corresponding ALTO maps can be found in Section 3.5. Depending on the application characteristics, it may not be possible or even not be desirable to completely localize all traffic.

3.1.3. Intra-Network Traffic Localization

The above sections described the results of the ALTO guidance on an inter-network level. However, ALTO can also be used for intra-network localization. In this case, ALTO provides guidance which internal hosts are to be preferred inside a single network or, e.g., one AS. Figure 6 shows such a scenario where Host 1 and Host 2 are located in Net 2 of ISP1 and connect via a low capacity link to the core (Net 1) of the same ISP1. If Host 1 and Host 2 exchange their data with remote hosts, they would probably congest the bottleneck link.

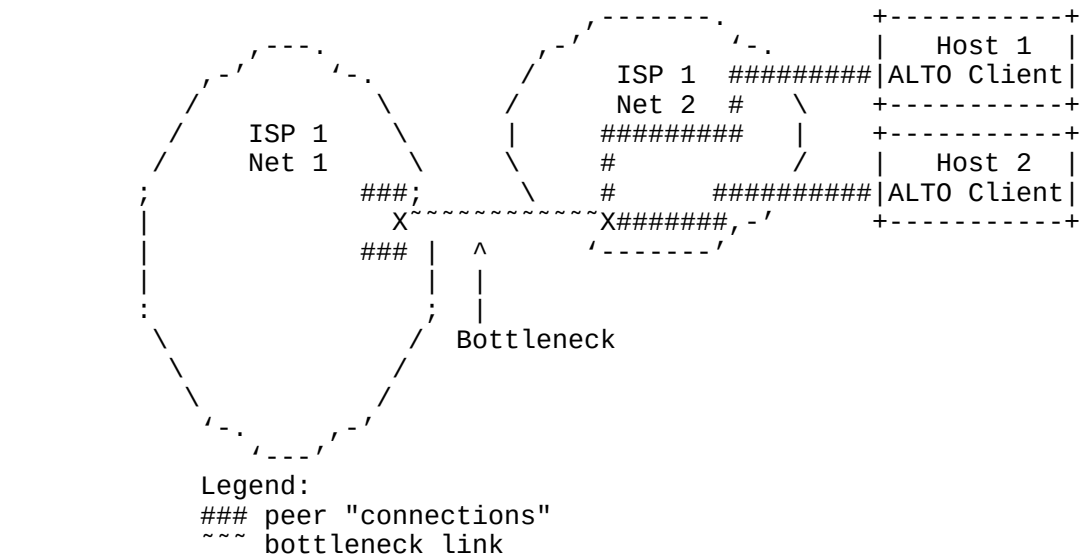

Figure 6: Without intra-network ALTO traffic localization

The operator can guide the hosts in such a situation to try first local hosts in the same network islands, avoiding or at least lowering the effect on the bottleneck link, as shown in Figure 7.

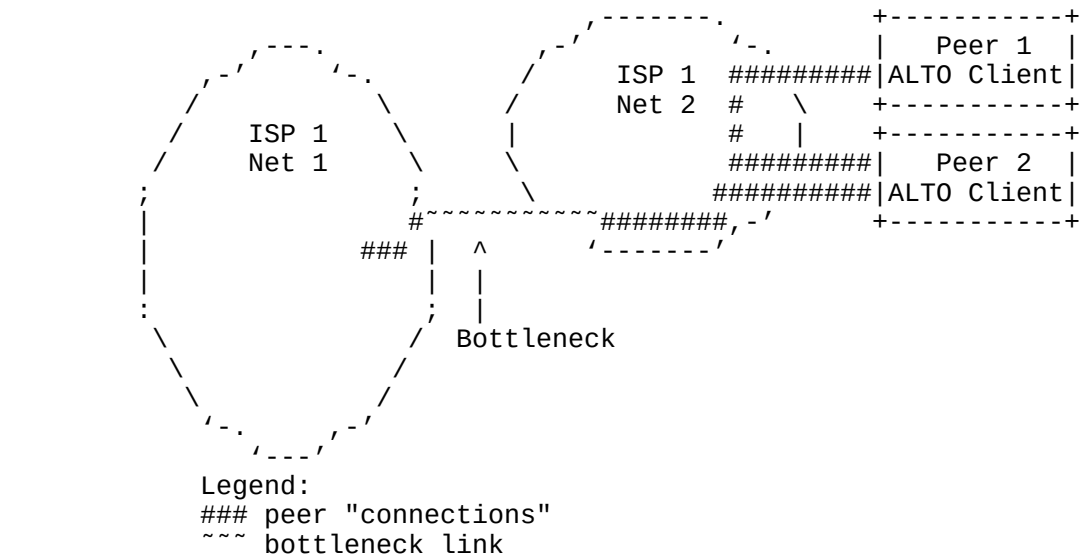

Figure 7: With intra-network ALTO traffic localization

The objective here is to avoid bottlenecks by optimized endpoint selection at application level. ALTO is not a method to deal with the congestion at the bottleneck.

3.1.4. Network Off-Loading

Another scenario is off-loading traffic from networks. This use of ALTO can be beneficial in particular in mobile networks. The network operator may have the desire to guide hosts in its own network to use hosts in remote networks. One reason can be that the wireless network is not made for the load cause by, e.g., peer-to-peer applications, and the operator has the need that peers fetch their data from remote peers in other parts of the Internet.

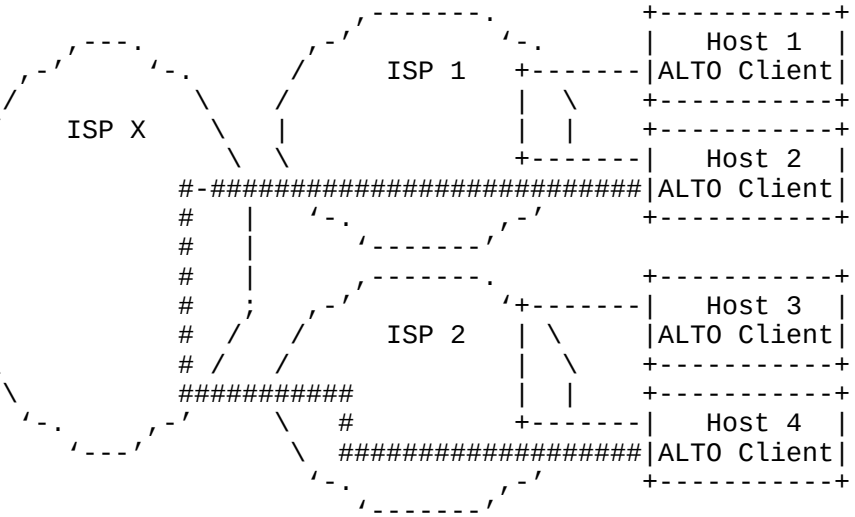

Legend:
=== preferred "connections"
--- non-preferred "connections"

Figure 8: ALTO traffic network de-localization

Figure 8 shows the result of such a guidance process where Host 2 prefers a connection with Host 4 instead of Host 1, as shown in Figure 5.

A realization of this scenario may have certain limitations and may not be possible in all cases. For instance, it may require that the ALTO server can distinguish mobile and non-mobile hosts, e.g., based on their IP address. This may depend on mobility solutions and may not be possible or accurate. In general, ALTO is not intended as a

fine-grained traffic engineering solution for individual hosts. Instead, it typically works on aggregates (e.g., if it is known that certain IP prefixes are often assigned to mobile users).

3.1.5. Application Tuning

ALTO can also provide guidance to optimize the application-level topology of networked applications, e.g., by exposing network performance information. Applications can often run own measurements to determine network performance, e.g., by active delay measurements or bandwidth probing, but such measurements result in overhead and complexity. Accessing an ALTO server can be a simpler alternative. In addition, an ALTO server may also expose network information that applications cannot easily measure or reverse-engineer.

3.2. Provisioning of ALTO Topology Data

3.2.1. Data Sources

An ALTO server can collect topological information from a variety of sources in the network and provides a cohesive, abstracted view of the network topology to applications using an ALTO client. Sources that may include routing protocols, network policies, state and performance information, geo-location, etc. Based on the input, the ALTO server builds an ALTO-specific network topology that represents the network as it should be understood and utilized by applications (resource consumers) at endpoints using ALTO services (e.g., Network/ Cost Map Service or ECS).

The ALTO protocol does not assume a specific network topology. In principle, ALTO can be used with various types of addresses (Endpoint Addresses). [RFC7285] defines the use of IPv4/IPv6 addresses or prefixes in ALTO, but further address types could be added by extensions. In this document, only the use of IPv4/IPv6 addresses is considered.

The exposure of network topology information is controlled and managed by the ALTO server. ALTO abstract network topologies can be automatically generated from the physical or logical topology of the network. The generation would typically be based on policies and rules set by the network operator. The maps and the guidance can significantly differ depending on the use case, the network architecture, and the trust relationship between ALTO server and ALTO client, etc. Besides the security requirements that consist of not delivering any confidential or critical information about the infrastructure, there are efficiency requirements in terms of what aspects of the network are visible and required by the given use case and/or application.

The ALTO server operator has to ensure that the ALTO topology does not contain any details that would endanger the network integrity and security. For instance, ALTO is not intended to leak raw Interior Gateway Protocol (IGP) or Border gateway Protocol (BGP) databases to ALTO clients.

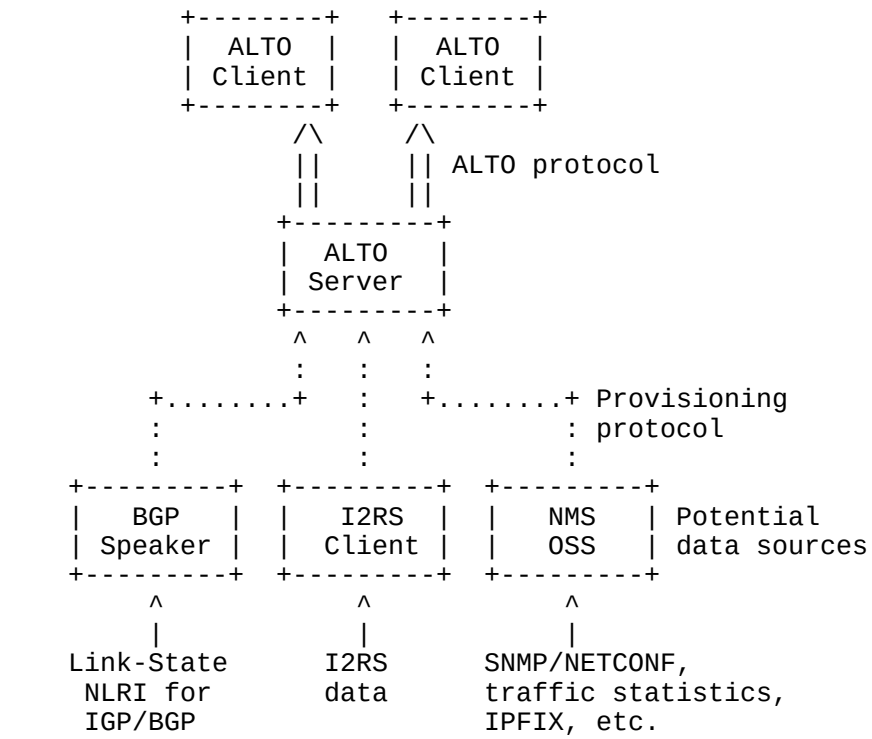

Figure 9: Potential data sources for ALTO

As illustrated in Figure 9, the topology data used by an ALTO server can originate from different data sources:

- The document [I-D.ietf-idr-ls-distribution] describes a mechanism by which links state and traffic engineering information can be collected from networks and shared with external components using the BGP routing protocol. This is achieved using a new BGP Network Layer Reachability Information (NLRI) encoding format. The mechanism is applicable to physical and virtual IGP links and can also include Traffic Engineering (TE) data. For instance, prefix data can be carried and originated in BGP, while TE data is originated and carried in an IGP. The mechanism described is subject to policy control. An ALTO Server can also use other

mechanisms to get network data, for example, peering with multiple IGP and BGP speakers.

- The Interface to the Routing System (I2RS) is a solution for state transfer in and out of the Internet's routing system [I-D.ietf-i2rs-architecture]. An ALTO server could use an I2RS client to observe routing-related information.

- An ALTO server can also leverage a Network Management System (NMS) or an Operations Support System (OSS) as data sources. NMS or OSS solutions are used to control, operate, and manage a network, e.g., using the Simple Network Management Protocol (SNMP) or NETCONF. As explained for instance in [I-D.farrkingel-pce-abno-architecture], the NMS and OSS can be consumers of network events reported and can act on these reports as well as displaying them to users and raising alarms. The NMS and OSS can also access the Traffic Engineering Database (TED) and Label Switched Path Database (LSP-DB) to show the users the current state of the network. In addition, NMS and OSS systems may have access to IGP/BGP routing information, network inventory data (e.g., links, nodes, or link properties not visible to routing protocols, such as Shared Risk Link Groups), statistics collection system that provides traffic information, such as traffic demands or link utilization obtained from IP Flow Information Export (IPFIX), as well as other Operations, Administration, and Maintenance (OAM) information (e.g., syslog). NMS or OSS systems also may have functions to correlate and orchestrate information originating from other data sources. For instance, it could be required to correlate IP prefixes with routers (Provider, Provider Edge, Customer Edge, etc.), IGP areas, VLAN IDs, or policies.

3.2.2. Privacy Requirements

Providing ALTO guidance can result in a win-win situation both for network providers and users of the ALTO information. Applications possibly get a better performance, while the network provider has means to optimize the traffic engineering and thus its costs. Yet, there can be security concerns with exposing topology data. Corresponding limitations are discussed in Section 7.2.

ISPs may have important privacy requirements when deploying ALTO. In particular, an ISP may not be willing to expose sensitive operational details of its network. The topology abstraction of ALTO enables an ISP to expose the network topology at a desired granularity only, determined by security policies.

With the Endpoint Cost Service (ECS), the ALTO client does not to have to implement any specific algorithm or mechanism in order to retrieve, maintain and process network topology information (of any kind). The complexity of the network topology (computation, maintenance and distribution) is kept in the ALTO server and ECS is delivered on demand. This allows the ALTO server to enhance and modify the way the topology information sources are used and combined. This simplifies the enforcement of privacy policies of the ISP.

The ALTO Network Map and Cost Map service expose an abstracted view on the ISP network topology. Therefore, in this case care is needed when constructing those maps in order to take into account privacy policies, as further discussed in Section 3.2.3. The ALTO protocol also supports further features such as endpoint properties, which could also be used to expose topology guidance. The privacy considerations for ALTO maps also apply to such ALTO extensions.

3.2.3. Partitioning and Grouping of IP Address Ranges

ALTO introduces provider-defined network location identifiers called Provider-defined Identifiers (PIDs) to aggregate network endpoints in the Map Services. Endpoints within one PID may be treated as single entity, assuming proximity based on network topology or other similarity. A key use case of PIDs is to specify network preferences (costs) between PIDs instead of individual endpoints. It is up to the operator of the ALTO server how to group endpoints and how to assign PIDs. For example, a PID may denote a subnet, a set of subnets, a metropolitan area, a POP, an autonomous system, or a set of autonomous systems.

This document only considers deployment scenarios in which PIDs expand to a set of IP address ranges (CIDR). A PID is characterized by a string identifier and its associated set of endpoint addresses [RFC7285]. If an ALTO server offers the Map Service, corresponding identifiers have to be configured.

An automated ALTO implementation may use dynamic algorithms to aggregate network topology. However, it is often desirable to have a mechanism through which the network operator can control the level and details of network aggregation based on a set of requirements and constraints. This will typically be governed by policies that enforce a certain level of abstraction and prevent leakage of sensitive operational data.

For instance, an ALTO server may leverage BGP information that is available in a networks service provider network layer and compute the group of prefix. An example are BGP communities, which are used

in MPLS/IP networks as a common mechanism to aggregate and group prefixes. A BGP community is an attribute used to tag a prefix to group prefixes based on mostly any criteria (as an example, most ISP networks originate BGP prefixes with communities identifying the Point of Presence (PoP) where the prefix has been originated). These BGP communities could be used to map IP address ranges to PIDs. By an additional policy, the ALTO server operator may decide an arbitrary cost defined between groups. Alternatively, there are algorithms that allow a dynamic computation of cost between groups. The ALTO protocol itself is independent of such algorithms and policies.

3.2.4. Rating Criteria and/or Cost Calculation

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 the operator of the ALTO server according to its own policy. For a given ALTO network map, an ALTO cost map defines directional path costs pairwise amongst the set of source and destination network locations defined by the PIDs.

The ALTO protocol permits the use of different cost types. An ALTO cost type is defined by the combination of a cost metric and a cost mode. The cost metric identifies what the costs represent. The cost mode identifies how the costs should be interpreted, e.g., whether returned costs should be interpreted as numerical values or ordinal rankings. The ALTO protocol also allows the definition of additional constraints defining which elements of a cost map shall be returned.

The ALTO protocol specification [RFC7285] defines the "routingcost" cost metric as basic set of rating criteria, which has to be supported by all implementations. This cost metric conveys a generic measure for the cost of routing traffic from a source to a destination. A lower value indicates a higher preference for traffic to be sent from a source to a destination. It is up to the ALTO server how that metric is calculated.

There is also an extension procedure for adding new ALTO cost types. The following list gives an overview on further rating criteria that have been proposed or which are in use by ALTO-related prototype implementations. This list is not intended as normative text; a definition of further metrics can be found for instance in [I-D.wu-alto-te-metrics]. Instead, the only purpose of the following list is to document and discuss rating criteria that have been proposed so far. It can also depend on the use case of ALTO whether such rating criteria are useful, and whether the corresponding information would indeed be made available by ISPs.

Distance-related rating criteria:

- Relative topological distance: The term relative means that a larger numerical value means greater distance, but it is up to the ALTO service how to compute the values, and the ALTO client will not be informed about the nature of the information. One way of generating this kind of information may be counting AS hops, but when querying this parameter, the ALTO client must not assume that the numbers actually are AS hops. In addition to the AS path, a relative cost value could also be calculated taking into account other routing protocol parameters, such as BGP local preference or multi-exit discriminator (MED) attributes.

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

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

- Absolute physical distance, based on knowledge of the approximate geo-location (e.g., continent, country) of an IP address.

Performance-related rating criteria:

- The minimum achievable throughput between the resource consumer and the candidate resource provider, which is considered useful by the application (only in ALTO queries).

- An arbitrary upper bound for the throughput from/to the candidate resource provider (only in ALTO responses). This may be, but is not necessarily the provisioned access bandwidth of the candidate resource provider.

- The maximum round-trip time (RTT) between resource consumer and the candidate resource provider, which is acceptable for the application for useful communication with the candidate resource provider (only in ALTO queries).

- An arbitrary lower bound for the RTT between resource consumer and the candidate resource provider (only in ALTO responses). This may be, for example, based on measurements of the propagation delay in a completely unloaded network.

Charging-related rating criteria:

- Traffic volume caps, in case the Internet access of the resource consumer is not charged by "flat rate". For each candidate

resource provider, the ALTO service could indicate the amount of data that may be transferred from/to this resource provider until a given point in time, and how much of this amount has already been consumed. Furthermore, it would have to be indicated how excess traffic would be handled (e.g., blocked, throttled, or charged separately at an indicated price). The interaction of several applications running on a host, out of which some use this criterion while others don't, as well as the evaluation of this criterion in resource directories, which issue ALTO queries on behalf of other endpoints, are for further study.

- Other metrics representing an abstract cost, e.g., determined by policies that distinguish "cheap" from "expensive" IP subnet ranges, e.g., without detailing the cost function.

These rating criteria are subject to the remarks below:

The ALTO client must be aware that with high probability the actual performance values differs from whatever an ALTO server exposes. In particular, an ALTO client must not consider a throughput parameter as a permission to send data at the indicated rate without using congestion control mechanisms.

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

- the ALTO service is not an admission control system
- the ALTO service may not know the instantaneous congestion status of the network
- the ALTO service may not know all link bandwidths, i.e., where the bottleneck really is, and there may be shared bottlenecks
- the ALTO service may not have all information about the actual routing
- the ALTO service may not know whether the candidate endpoints itself is overloaded
- the ALTO service may not know whether the candidate endpoints throttles the bandwidth it devotes for the considered application
- the ALTO service may not know whether the candidate endpoints will throttle the data it sends to us (e.g., because of some fairness algorithm, such as tit-for-tat).

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

Rating criteria that should not be defined for and used by the ALTO service include:

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

- Performance metrics that raise privacy concerns. For instance, it has been questioned whether an ALTO service could publicly expose the provisioned access bandwidth, e.g. of cable / DSL customers, because this could enables identification of "premium" customers.

3.3. Known Limitations of ALTO

3.3.1. Limitations of Map-based Approaches

The specification of the Map Service in the ALTO protocol [RFC7285] is based on the concept of network maps. A network map partitions the network into Provider-defined Identifier (PID) that group one or multiple endpoints (e.g., subnetworks) to a single aggregate. The "costs" between the various PIDs is stored in a cost map. Map-based approaches lower the signaling load on the server as maps have to be retrieved only if they change.

One main assumption for map-based approaches is that the information provided in these maps is static for a longer period of time. This assumption is fine as long as the network operator does not change any parameter, e.g., routing within the network and to the upstream peers, IP address assignment stays stable (and thus the mapping to the partitions). However, there are several cases where this assumption is not valid:

1. ISPs reallocate IP subnets from time to time;

2. ISPs reallocate IP subnets on short notice;

3. IP prefix blocks may be assigned to a router that serves a variety of access networks;

4. Network costs between IP prefixes may change depending on the ISP’s routing and traffic engineering.

These effects can be explained as follows:

Case 1: ISPs may reallocate IP subnets within their infrastructure from time to time, partly to ensure the efficient usage of IPv4 addresses (a scarce resource), and partly to enable efficient route tables within their network routers. The frequency of these "renumbering events" depend on the growth in number of subscribers and the availability of address space within the ISP. As a result, a subscriber’s household device could retain an IP address for as short as a few minutes, or for months at a time or even longer.

It has been suggested that ISPs providing ALTO services could sub-divide their subscribers’ devices into different IP subnets (or certain IP address ranges) based on the purchased service tier, as well as based on the location in the network topology. The problem is that this sub-allocation of IP subnets tends to decrease the efficiency of IP address allocation, in particular for IPv4. A growing ISP that needs to maintain high efficiency of IP address utilization may be reluctant to jeopardize their future acquisition of IP address space.

However, this is not an issue for map-based approaches if changes are applied in the order of days.

Case 2: ISPs can use techniques that allow the reallocation of IP prefixes on very short notice, i.e., within minutes. An IP prefix that has no IP address assignment to a host anymore can be reallocated to areas where there is currently a high demand for IP addresses.

Case 3: In residential access networks (e.g., DSL, cable), IP prefixes are assigned to broadband gateways, which are the first IP-hop in the access-network between the Customer Premises Equipment (CPE) and the Internet. The access-network between CPE and broadband gateway (called aggregation network) can have varying characteristics (and thus associated costs), but still using the same IP prefix. For instance one IP addresses IP11 out of a IP prefix IP1 can be assigned to a VDSL (e.g., 2 MBit/s uplink) access line while the subsequent IP address IP12 is assigned to a slow ADSL line (e.g., 128 kbit/s uplink). These IP addresses are assigned on a first come first served basis, i.e., a single IP address out of the same IP prefix can change its associated costs quite fast. This may not be an issue

with respect to the used upstream provider (thus the cross ISP traffic) but depending on the capacity of the aggregation-network this may raise to an issue.

Case 4: The routing and traffic engineering inside an ISP network, as well as the peering with other autonomous systems, can change dynamically and affect the information exposed by an ALTO server. As a result, cost map and possibly also network maps can change.

3.3.2. Limitations of Non-Map-based Approaches

The specification of the ALTO protocol [RFC7285] also includes the Endpoint Cost Service (ECS) mechanism. ALTO clients can ask guidance for specific IP addresses to the ALTO server, thereby avoiding the need of processing maps. This can mitigate some of the problems mentioned in the previous section.

However, asking for IP addresses, asking with long lists of IP addresses, and asking quite frequently may overload the ALTO server. The server has to rank each received IP address, which causes load at the server. This may be amplified by the fact that not only a single ALTO client is asking for guidance, but a larger number of them. The results of the ECS are also more difficult to cache than ALTO maps. Therefore, the ALTO client may have to await the server response before starting a communication, which results in an additional delay.

Caching of IP addresses at the ALTO client or the usage of the H12 approach [I-D.kiesel-alto-h12] in conjunction with caching may lower the query load on the ALTO server.

When ALTO server receives an ECS request, it may not have the most appropriate topology information in order to accurately determine the ranking. [RFC7285] generally assumes that a server can always offer some guidance. In such a case the ALTO server could adopt one of the following strategies:

- Reply with available information (best effort).
- Query another ALTO server presumed to have better topology information and return that response (cascaded servers).
- Redirect the request to another ALTO server presumed to have better topology information (redirection).

The protocol mechanisms and decision processes that would be used to determine if redirection is necessary and which mode to use is out of

the scope of this document, since protocol extensions could be required.

3.3.3. General Limitations

ALTO is designed as a protocol between clients integrated in applications and servers that provide network information and guidance (e.g., basic network location structure and preferences of network paths). The objective is to modify network resource consumption patterns at application level while maintaining or improving application performance. This design focus results in a number of characteristics of ALTO:

- Endpoint focus: In typical ALTO use cases, neither the consumer of the topology information (i.e., the ALTO client) nor the considered resources (e.g., files at endpoints) are part of the network. The ALTO server presents an abstract network topology containing only information relevant to an application overlay for better-than-random resource selections among its endpoints. The ALTO protocol specification [RFC7285] is not designed to expose network internals such as routing tables or configuration data that are not relevant for application-level resource selection decisions in network endpoints.

- Abstraction: The ALTO services such as the Network/Cost Map Service or the ECS provide an abstract view of the network only. The operator of the ALTO server has full control over the granularity (e.g., by defining policies how to aggregate subnets into PIDs) and the level-of-detail of the abstract network representation (e.g., by deciding what cost types to support).

- Multiple administrative domains: The ALTO protocol is designed for use cases where the ALTO server and client can be located in different organizations or trust domains. ALTO assumes a loose coupling between server and client. In addition, ALTO does not assume that an ALTO client has any a priori knowledge about the ALTO server and its supported features. An ALTO server can be discovered automatically.

- Read-only: ALTO is a query/response protocol to retrieve guidance information. Neither network/cost map queries nor queries to the endpoint cost service are designed to affect state in the network.

If ALTO shall be deployed for use cases violating these assumptions, the protocol design may result in limitations.

For instance, in an Application-Based Network Operation (ABNO) environment the application could issue explicit service request to

the network [I-D.farrkingel-pce-abno-architecture]. In this case, the application would require detailed knowledge about the internal network topology and the actual state. A network configuration would also require a corresponding security solution for authentication and authorization. ALTO is not designed for operations to control, operate, and manage a network.

Such deployments could be addressed by network management solutions, e.g., based on SNMP [RFC3411] or NETCONF [RFC6241] and YANG [RFC6020] that are typically designed to manipulate configuration state. Reference [I-D.farrkingel-pce-abno-architecture] contains a more detailed discussion of interfaces between components such as Element Management System (EMS), Network Management System (NMS), Operations Support System (OSS), Traffic Engineering Database (TED), Label Switched Path Database (LSP-DB), Path Computation Element (PCE), and other Operations, Administration, and Maintenance (OAM) components.

3.4. Monitoring ALTO

3.4.1. Impact and Observation on Network Operation

ALTO presents a new opportunity for managing network traffic by providing additional information to clients. In particular, the deployment of an ALTO Server may shift network traffic patterns, and the potential impact to network operation can be large. An ISP providing ALTO may want to assess the benefits of ALTO as part of the management and operations (cf. [RFC7285]). For instance, the ISP might be interested in understanding whether the provided ALTO maps are effective, and in order to decide whether an adjustment of the ALTO configuration would be useful. Such insight can be obtained from a monitoring infrastructure. An ISP offering ALTO could consider the impact on (or integration with) traffic engineering and the deployment of a monitoring service to observe the effects of ALTO operations. The measurement of impacts can be challenging because ALTO-enabled applications may not provide related information back to the ALTO Service Provider.

To construct an effective monitoring infrastructure, the ALTO Service Provider should decide how to monitor the performance of ALTO and identify and deploy data sources to collect data to compute the performance metrics. In certain trusted deployment environments, it may be possible to collect information directly from ALTO clients. It may also be possible to vary or selectively disable ALTO guidance for a portion of ALTO clients either by time, geographical region, or some other criteria to compare the network traffic characteristics with and without ALTO. Monitoring an ALTO service could also be realized by third parties. In this case, insight into ALTO data may

require a trust relationship between the monitoring system operator and the network service provider offering an ALTO service.

The required monitoring depends on the network infrastructure and the use of ALTO, and an exhaustive description is outside the scope of this document.

3.4.2. Measurement of the Impact

ALTO realizes an interface between the network and applications. This implies that an effective monitoring infrastructure may have to deal with both network and application performance metrics. This document does not comprehensively list all performance metrics that could be relevant, nor does it formally specify metrics.

The impact of ALTO can be classified regarding a number of different criteria:

- Total amount and distribution of traffic: ALTO enables ISPs to influence and localize traffic of applications that use the ALTO service. An ISP may therefore be interested in analyzing the impact on the traffic, i.e., whether network traffic patterns are shifted. For instance, if ALTO shall be used to reduce the inter-domain P2P traffic, it makes sense to evaluate the total amount of inter-domain traffic of an ISP. Then, one possibility is to study how the introduction of ALTO reduces the total inter-domain traffic (inbound and/our outbound). If the ISPs intention is to localize the traffic inside his network, the network-internal traffic distribution will be of interest. Effectiveness of localization can be quantified in different ways, e.g., by the load on core routers and backbone links, or by considering more advanced effects, such as the average number of hops that traffic traverses inside a domain.

- Application performance: The objective of ALTO is improve application performance. ALTO can be used by very different types applications, with different communication characteristics and requirements. For instance, if ALTO guidance achieves traffic localization, one would expect that applications achieve a higher throughput and/or smaller delays to retrieve data. If application-specific performance characteristics (e.g., video or audio quality) can be monitored, such metrics related to user experience could also help to analyze the benefit of an ALTO deployment. If available, selected statistics from the TCP/IP stack in hosts could be leveraged, too.

Of potential interest can also be the share of applications or customers that actually use an offered ALTO service, i.e., the adoption of the service.

Monitoring statistics can be aggregated, averaged, and normalized in different ways. This document does not mandate specific ways how to calculate metrics.

3.4.3. System and Service Performance

A number of interesting parameters can be measured at the ALTO server. [RFC7285] suggests certain ALTO-specific metrics to be monitored:

- Requests and responses for each service listed in a Information Directory (total counts and size in bytes).
- CPU and memory utilization
- ALTO map updates
- Number of PIDs
- ALTO map sizes (in-memory size, encoded size, number of entries)

This data characterizes the workload, the system performance as well as the map data. Obviously, such data will depend on the implementation and the actual deployment of the ALTO service. Logging is also recommended in [RFC7285].

3.4.4. Monitoring Infrastructures

Understanding the impact of ALTO may require interaction between different systems, operating at different layers. Some information discussed in the preceding sections is only visible to an ISP, while application-level performance can hardly be measured inside the network. It is possible that not all information of potential interest can directly be measured, either because no corresponding monitoring infrastructure or measurement method exists, or because it is not easily accessible.

One way to quantify the benefit of deploying ALTO is to measure before and after enabling the ALTO service. In addition to passive monitoring, some data could also be obtained by active measurements, but due to the resulting overhead, the latter should be used with care. Yet, in all monitoring activities an ALTO service provider has to take into account that ALTO clients are not bound to ALTO server

guidance as ALTO is only one source of information, and any measurement result may thus be biased.

Potential sources for monitoring the use of ALTO include:

- Network Operations, Administration, and Maintenance (OAM) systems: Many ISPs deploy OAM systems to monitor the network traffic, which may have insight into traffic volumes, network topology, and bandwidth information inside the management area. Data can be obtained by SNMP, NETCONF, IP Flow Information Export (IPFIX), syslog, etc.

- Applications/clients: Relevant data could be obtained by instrumentation of applications.

- ALTO server: If available, log files or other statistics data could be analyzed.

- Other application entities: In several use cases, there are other application entities that could provide data as well. For instance, there may be centralized log servers that collect data.

In many ALTO use cases some data sources are located within an ISP network while some other data is gathered at application level. Correlation of data could require a collaboration agreement between the ISP and an application owner, including agreements of data interchange formats, methods of delivery, etc. In practice, such a collaboration may not be possible in all use cases of ALTO, because the monitoring data can be sensitive, and because the interacting entities may have different priorities. Details of how to build an over-arching monitoring system for evaluating the benefits of ALTO are outside the scope of this memo.

3.5. Map Examples for Different Types of ISPs

3.5.1. Small ISP with Single Internet Uplink

The ALTO protocol does not mandate how to determine costs between endpoints and/or determine map data. In complex usage scenarios this can be a non-trivial problem. In order to show the basic principle, this and the following sections explain for different deployment scenarios how ALTO maps could be structured.

For a small ISP, the inter-domain traffic optimizing problem is how to decrease the traffic exchanged with other ISPs, because of high settlement costs. By using the ALTO service to optimize traffic, a small ISP can define two "optimization areas": one is its own network; the other one consists of all other network destinations.

The cost map can be defined as follows: the cost of link between clients of inner ISP's networks is lower than between clients of outer ISP's networks and clients of inner ISP's network. As a result, a host with ALTO client inside the network of this ISP will prefer retrieving data from hosts connected to the same ISP.

An example is given in Figure 10. It is assumed that ISP A is a small ISP only having one access network. As operator of the ALTO service, ISP A can define its network to be one optimization area, named as PID1, and define other networks to be the other optimization area, named as PID2. C1 is denoted as the cost inside the network of ISP A. C2 is denoted as the cost from PID2 to PID1, and C3 from PID1 to PID2. For the sake of simplicity, in the following C2=C3 is assumed. In order to keep traffic local inside ISP A, it makes sense to define: C1<C2

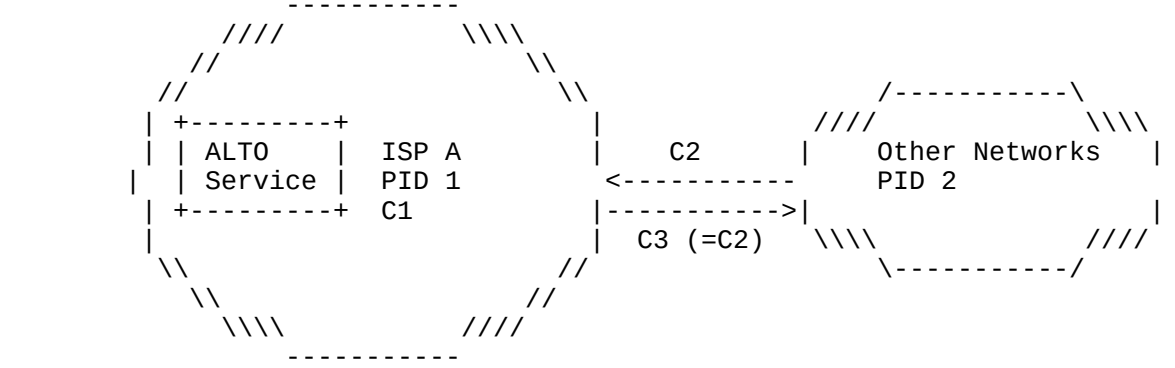

Figure 10: Example ALTO deployment in small ISPs

A simplified extract of the corresponding ALTO network and cost maps is listed in Figure 11 and Figure 12, assuming that the network of ISP A has the IPv4 address ranges 192.0.2.0/24 and 198.51.100.0/25. In this example, the cost values C1 and C2 can be set to any number C1<C2.

```
HTTP/1.1 200 OK
...
Content-Type: application/alto-networkmap+json

{
 ...
  "network-map" : {
    "PID1" : {
      "ipv4" : [
        "192.0.2.0/24",
        "198.51.100.0/25"
      ]
    },
    "PID2" : {
      "ipv4" : [
        "0.0.0.0/0"
      ],
      "ipv6" : [
        "::/0"
      ]
    }
  }
}
```

Figure 11: Example ALTO network map

```
HTTP/1.1 200 OK
...
Content-Type: application/alto-costmap+json

{
    ...
    "cost-type" : {"cost-mode"  : "numerical",
                   "cost-metric": "routingcost"
    }
  },
  "cost-map" : {
    "PID1": { "PID1": C1,  "PID2": C2 },
    "PID2": { "PID1": C2,  "PID2": 0 },
  }
}
```

Figure 12: Example ALTO cost map

3.5.2. ISP with Several Fixed Access Networks

This example discusses a P2P application traffic optimization use case for a larger ISP with a fixed network comprising several access networks and a core network. The traffic optimizing objectives include (1) using the backbone network efficiently, (2) adjusting the traffic balance in different access networks according to traffic conditions and management policies, and (3) achieving a reduction of settlement costs with other ISPs.

Such a large ISP deploying an ALTO service may want to optimize its traffic according to the network topology of its access networks. For example, each access network could be defined to be one optimization area, i.e., traffic should be kept local withing that area if possible. This can be achieved by mapping each area to a PID. Then the costs between those access networks can be defined according to a corresponding traffic optimizing requirement by this ISP. One example setup is further described below and also shown in Figure 13.

In this example, ISP A has one backbone network and three access networks, named as AN A, AN B, and AN C. A P2P application is used in this example. For a reasonable application-level traffic optimization, the first requirement could be a decrease of the P2P traffic on the backbone network inside the Autonomous System of ISP A and the second requirement could be a decrease of the P2P traffic to other ISPs, i.e., other Autonomous Systems. The second requirement can be assumed to have priority over the first one. Also, we assume that the settlement rate with ISP B is lower than with other ISPs. ISP A can deploy an ALTO service to meet these traffic distribution requirements. In the following, we will give an example of an ALTO setting and configuration according to these requirements.

In the network of ISP A, the operator of the ALTO server can define each access network to be one optimization area, and assign one PID to each access network, such as PID 1, PID 2, and PID 3. Because of different peerings with different outer ISPs, one can define ISP B to be one additional optimization area and assign PID 4 to it. All other networks can be added to a PID to be one further optimization area (PID 5).

In the setup, costs (C1, C2, C3, C4, C5, C6, C7, C8) can be assigned as shown in Figure 13. Cost C1 is denoted as the link cost in inner AN A (PID 1), and C2 and C3 are defined accordingly. C4 is denoted as the link cost from PID 1 to PID 2, and C5 is the corresponding cost from PID 3, which is assumed to have a similar value. C6 is the cost between PID 1 and PID 3. For simplicity, this scenario assumes symmetrical costs between the AN this example. C7 is denoted as the

link cost from the ISP B to ISP A. C8 is the link cost from other networks to ISP A.

According to previous discussion of the first requirement and the second requirement, the relationship of these costs will be defined as: (C1, C2, C3) < (C4, C5, C6) < (C7) < (C8)

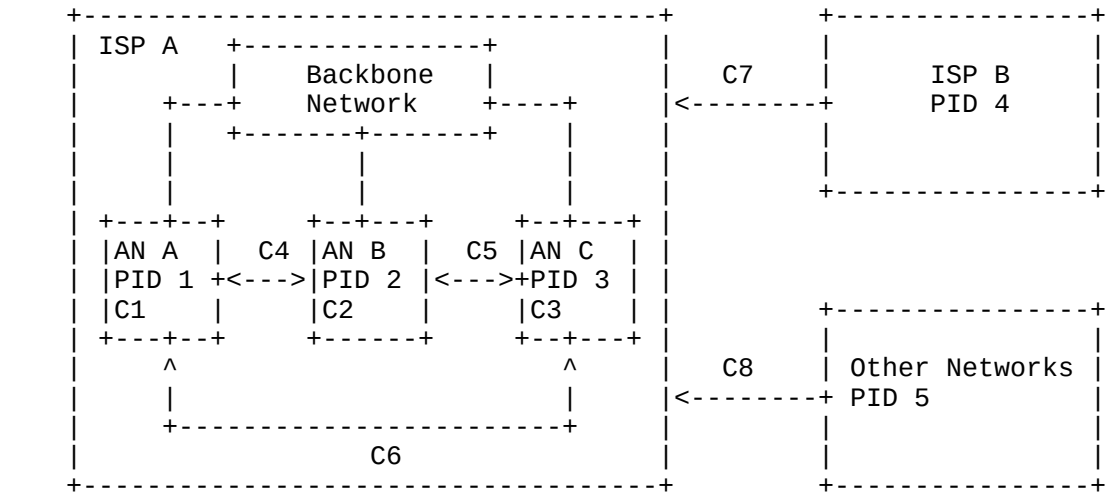

Figure 13: ALTO deployment in large ISPs with layered fixed network structures

3.5.3. ISP with Fixed and Mobile Network

An ISP with both mobile network and fixed network my focus on optimizing the mobile traffic by keeping traffic in the fixed network as far as possible, because wireless bandwidth is a scarce resource and traffic is costly in mobile network. In such a case, the main requirement of traffic optimization could be decreasing the usage of radio resources in the mobile network. An ALTO service can be deployed to meet these needs.

Figure 14 shows an example: ISP A operates one mobile network, which is connected to a backbone network. The ISP also runs two fixed access networks AN A and AN B, which are also connected to the backbone network. In this network structure, the mobile network can be defined as one optimization area, and PID 1 can be assigned to it. Access networks AN A and B can also be defined as optimization areas, and PID 2 and PID 3 can be assigned, respectively. The cost values are then defined as shown in Figure 14.

To decrease the usage of wireless link, the relationship of these costs can be defined as follows:

From view of mobile network: C4 < C1. This means that clients in mobile network requiring data resource from other clients will prefer clients in AN A to clients in the mobile network. This policy can decrease the usage of wireless link and power consumption in terminals.

From view of AN A: C2 < C6, C5 = maximum cost. This means that clients in other optimization area will avoid retrieving data from the mobile network.

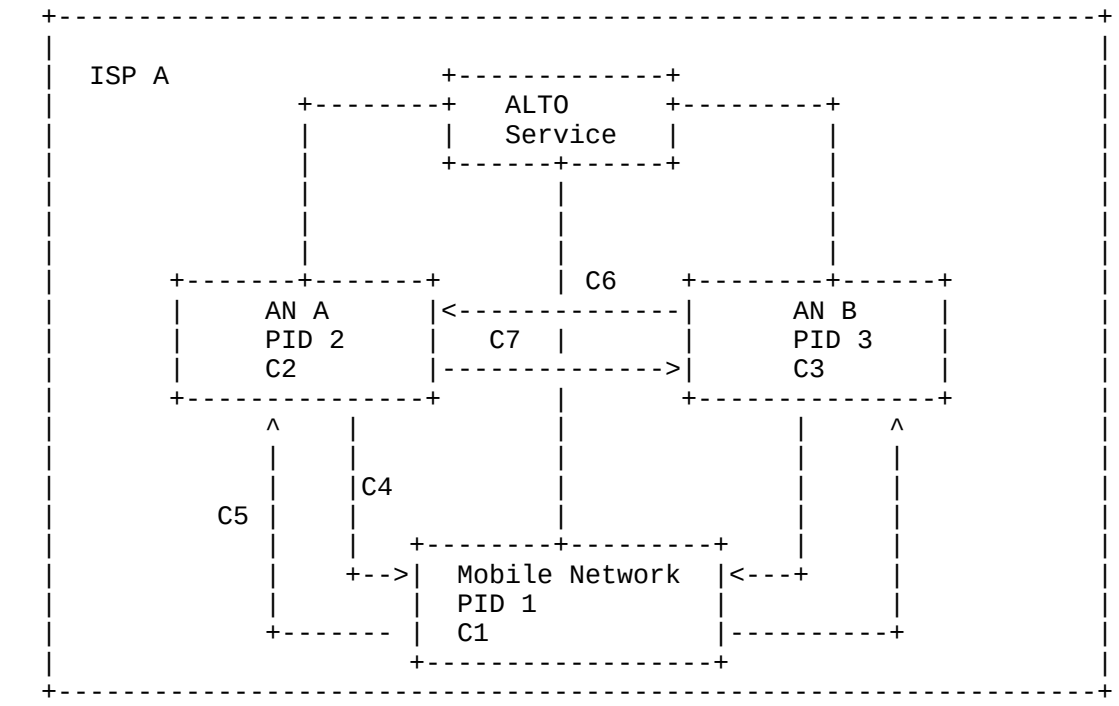

Figure 14: ALTO deployment in ISPs with mobile network

These examples show that for ALTO in particular the relations between different costs matter; the operator of the server has several degrees of freedom how to set the absolute values.

3.6. Deployment Experiences

The examples in the previous section are simple and do not consider specific requirements inside access networks, such as different link types. Deploying an ALTO service in real network may require dealing with further network conditions and requirements. One real example is described in greater detail in reference [I-D.lee-alto-chinatelecom-trial].

Also, experiments have been conducted with ALTO-like deployments in Internet Service Provider (ISP) networks. For instance, NTT performed tests with their HINT server implementation and dummy nodes to gain insight on how an ALTO-like service influence peer-to-peer systems [I-D.kamei-p2p-experiments-japan]. The results of an early experiment conducted in the Comcast network are documented in [RFC5632].

4. Using ALTO for P2P Traffic Optimization

4.1. Overview

4.1.1. Usage Scenario

Originally, peer-to-peer (P2P) applications have been the main driver for the development of ALTO. In this use case it is assumed that one party (usually the operator of a "managed" IP network domain) will disclose information about the network through ALTO. The application overlay will query this information and optimize its behavior in order to improve performance or Quality of Experience in the application while reducing the utilization of the underlying network infrastructure. The resulting win-win situation is assumed to be the incentive for both parties to provide or consume the ALTO information, respectively.

P2P systems can be build without or with use of a centralized resource directory ("tracker"). The scope of this section is the interaction of P2P applications with the ALTO service. In this scenario, the resource consumer ("peer") asks the resource directory for a list of candidate resource providers, which can provide the desired resource. There are different options how ALTO can be deployed in such use cases with a centralized resource directory.

For efficiency reasons (i.e., message size), usually only a subset of all resource providers known to the resource directory will be returned to the resource consumer. Some or all of these resource providers, plus further resource providers learned by other means such as direct communication between peers, will be contacted by the resource consumer for accessing the resource. The purpose of ALTO is

giving guidance on this peer selection, which is supposed to yield better-than-random results. The tracker response as well as the ALTO guidance are most beneficial in the initial phase after the resource consumer has decided to access a resource, as long as only few resource providers are known. Later, when the resource consumer has already exchanged some data with other peers and measured the transmission speed, the relative importance of ALTO may dwindle.

4.1.2. Applicability of ALTO

A tracker-based P2P application can leverage ALTO in different ways. In the following, the different alternatives and their pros and cons are discussed.

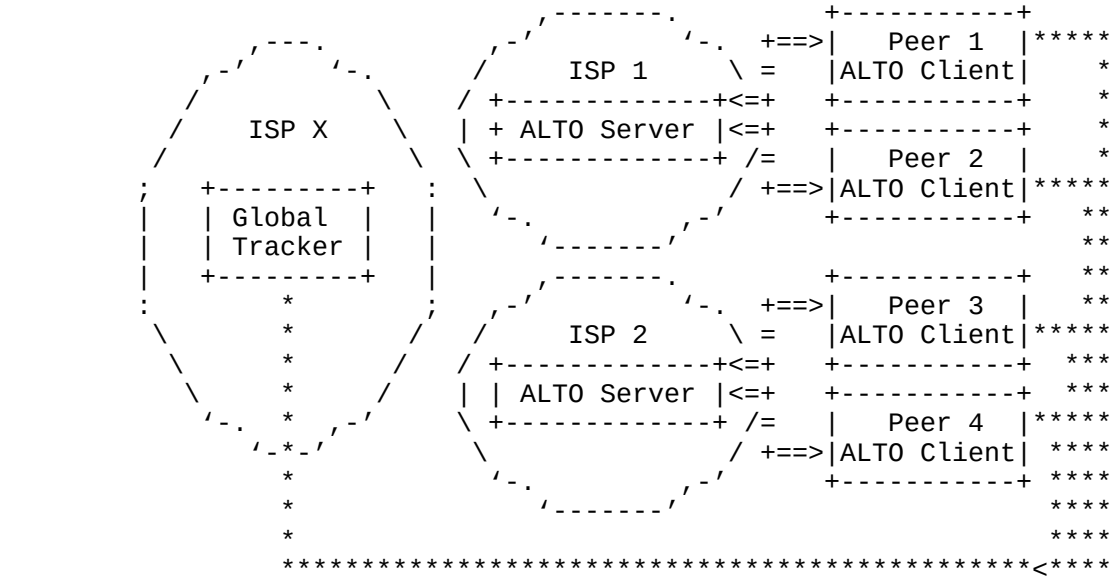

Figure 15: Global tracker and local ALTO servers

Figure 15 depicts a tracker-based P2P system with several peers. The peers (i.e., resource consumers) embed an ALTO client to improve the resource selection. The tracker (i.e., resource directory) itself may be hosted and operated by another entity. A tracker outside the networks of the ISPs of the peers may be a typical use case. For instance, a tracker like Pirate Bay can serve Bittorrent peers world-wide. The figure only shows one tracker instance, but deployments with several trackers could be possible, too.

In the scenario depicted in Figure 15 lets the peers directly communicate with their ISP's ALTO server (i.e., ALTO client embedded in the peers), giving thus the peers the most control on which information they query for, as they can integrate information received from one tracker or several trackers and through direct peer-to-peer knowledge exchange. For instance, the latter approach is called peer exchange (PEX) in bittorent. In this deployment scenarios, the peers have to discover a suitable ALTO server e.g. offered by their ISP, as described in [RFC7286].

There are also tracker-less P2P system architectures that do not rely on centralized resource directories, e.g., unstructured P2P networks. Regarding the use of ALTO, their deployment would be similar to Figure 15, since the ALTO client would be embedded in the peers as well. This option is not further considered in this memo.

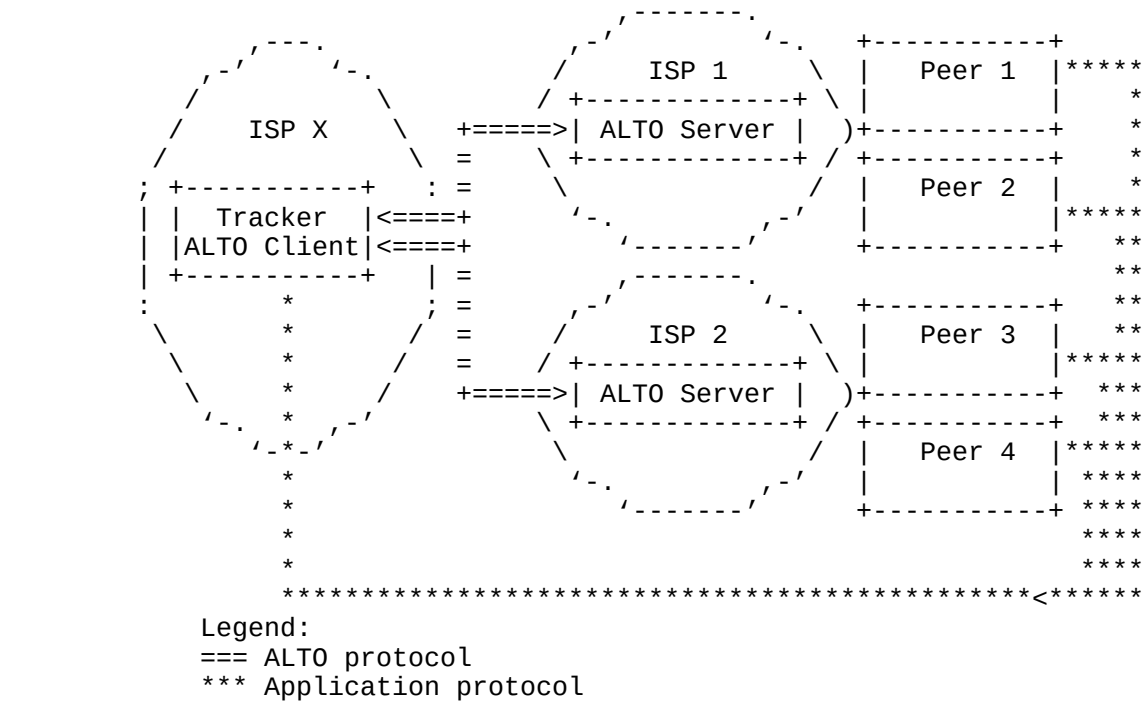

Figure 16: Global tracker accessing ALTO server at various ISPs

An alternative deployment scenario for a tracker-based system is depicted in Figure 16. Here, the tracker embeds the ALTO client. As already explained, the tracker itself may be hosted and operated by an entity different than the ISP hosting and operating the ALTO server. The key difference to the previously discussed use case is

that the ALTO client is different to the resource consumer. Initially, the tracker has to look-up the ALTO server in charge for each peer where it receives a ALTO query for. Therefore, the ALTO server has to discover the handling ALTO server for a pear [RFC7286] [I-D.kiesel-alto-xdom-disc]. The peers do not have any way to query the ALTO server themselves. This setting allows giving the peers a better selection of candidate peers for their operation at an initial time, but does not consider peers learned through direct peer-to-peer knowledge exchange.

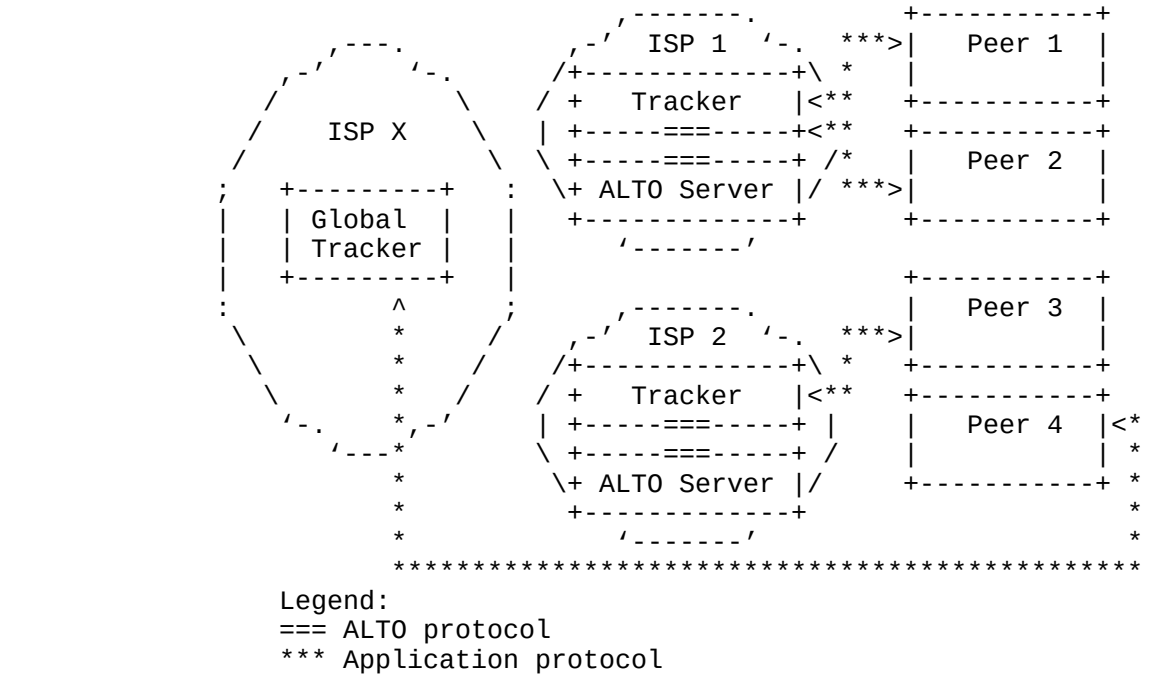

Figure 17: Local trackers and local ALTO servers (P4P approach)

There are some attempts to let ISP's to deploy their own trackers, as shown in Figure 17. In this case, the client has no chance to get guidance from the ALTO server, other than talking to the ISP's tracker. However, the peers would have still chance the contact other trackers, deployed by entities other than the peer's ISP.

4.2. Deployment Recommendations

4.2.1. ALTO Services

The ALTO protocol specification [RFC7285] details how an ALTO client can query an ALTO server for guiding information and receive the corresponding replies. In case of peer-to-peer networks, two different ALTO services can be used: The Cost Map Service is often preferred as solution by peer-to-peer software implementors and users, since it avoids disclosing peer IP addresses to a centralized entity. Different to that, network operators may have a preference for the Endpoint Cost Service (ECS), since it does not require exposure of the network topology.

For actual use of ALTO in P2P applications, both software vendors and network operators have to agree which ALTO services to use. The ALTO protocol is flexible and supports both services. Note that for other use cases of ALTO, in particular in more controlled environments, both the Cost Map Service as well as Endpoint Cost Service might be feasible and it is more an engineering trade-off whether to use a map-based or query-based ALTO service.

4.2.2. Guidance Considerations

As explained in Section 4.1.2, for a tracker-based P2P application there are two fundamentally different possibilities where to place the ALTO client:

1. ALTO client in the resource consumer ("peer")

2. ALTO client in the resource directory ("tracker")

Both approaches have advantages and drawbacks that have to be considered. If the ALTO client is in the resource consumer (Figure 15), a potentially very large number of clients has to be deployed. Instead, when using an ALTO client in the resource directory (Figure 16 and Figure 17), ostensibly peers do not have to directly query the ALTO server. In this case, an ALTO server could even not permit access to peers.

However, it seems to be beneficial for all participants to let the peers directly query the ALTO server. Considering the plethora of different applications that could use ALTO, e.g. multiple tracker or non-tracker based P2P systems or other applications searching for relays, this renders the ALTO service more useful. The peers are also the single point having all operational knowledge to decide whether to use the ALTO guidance and how to use the ALTO guidance. For a given peer one can also expect that an ALTO server of the corresponding ISP provides useful guidance and can be discovered.

Yet, ALTO clients in the resource consumer also have drawbacks compared to use in the resource directory. In the following, both scenarios are compared more in detail in order to explain the impact on ALTO guidance and the need for third-party ALTO queries.

In the first scenario (see Figure 18), the peer (resource consumer) queries the tracker (resource directory) for the desired resource (F1). The resource directory returns a list of potential resource providers without considering ALTO (F2). It is then the duty of the resource consumer to invoke ALTO (F3/F4), in order to solicit guidance regarding this list.

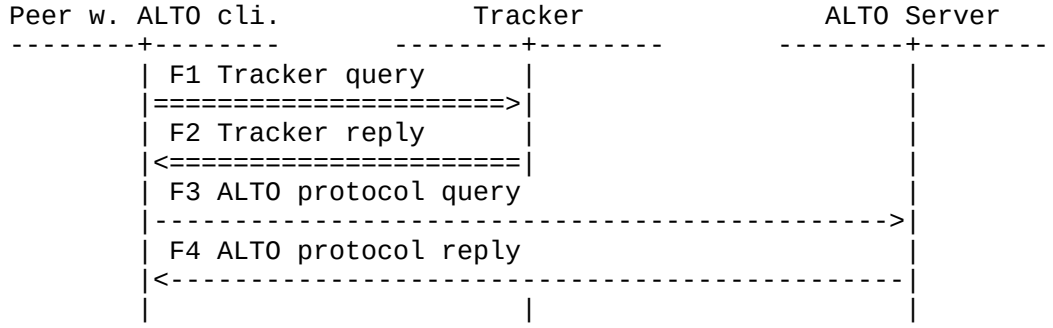

==== Application protocol (i.e., tracker-based P2P app protocol)
---- ALTO protocol

Figure 18: Basic message sequence chart for resource consumer-initiated ALTO query

In the second scenario (see Figure 19), the resource directory has an embedded ALTO client, which we will refer to as Resource Directory ALTO Client (RDAC) in this document. After receiving a query for a given resource (F1) the resource directory invokes the RDAC to evaluate all resource providers it knows (F2/F3). Then it returns a, possibly shortened, list containing the "best" resource providers to the resource consumer (F4).

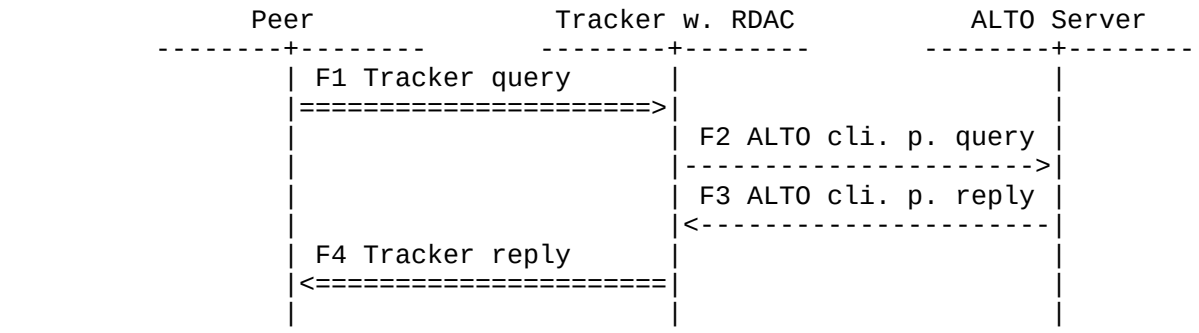

==== Application protocol (i.e., tracker-based P2P app protocol)
---- ALTO protocol

Figure 19: Basic message sequence chart for third-party ALTO query

Note: The message sequences depicted in Figure 18 and Figure 19 may occur both in the target-aware and the target-independent query mode (cf. [RFC6708]). In the target-independent query mode no message exchange with the ALTO server might be needed after the tracker query, because the candidate resource providers could be evaluated using a locally cached "map", which has been retrieved from the ALTO server some time ago.

The first approach has the following problem: While the resource directory might know thousands of peers taking part in a swarm, the list returned to the resource consumer is usually shortened for efficiency reasons. Therefore, the "best" (in the sense of ALTO) potential resource providers might not be contained in that list anymore, even before ALTO can consider them.

Much better traffic optimization could be achieved if the tracker would evaluate all known peers using ALTO. This list would then include a significantly higher fraction of "good" peers. If the tracker returned "good" peers only, there might be a risk that the swarm might disconnect and split into several disjunct partitions. However, finding the right mix of ALTO-biased and random peer selection is out of the scope of this document.

Therefore, from an overall optimization perspective, the second scenario with the ALTO client embedded in the resource directory is advantageous, because it is ensured that the addresses of the "best" resource providers are actually delivered to the resource consumer. An architectural implication of this insight is that the ALTO server discovery procedures must support third-party discovery. That is, as the tracker issues ALTO queries on behalf of the peer which contacted the tracker, the tracker must be able to discover an ALTO server that

can give guidance suitable for that respective peer (see [I-D.kiesel-alto-xdom-disc]).

5. Using ALTO for CDNs

5.1. Overview

5.1.1. Usage Scenario

This section briefly introduces the usage of ALTO for Content Delivery Networks (CDNs), as explained e.g. in [I-D.jenkins-alto-cdn-use-cases]. CDNs are used in the delivery of some Internet services (e.g. delivery of websites, software updates and video delivery) from a location closer to the location of the user. A CDN typically consists of a network of servers often attached to Internet Service Provider (ISP) networks. The point of attachment is often as close to content consumers and peering points as economically or operationally feasible in order to decrease traffic load on the ISP backbone and to provide better user experience measured by reduced latency and higher throughput.

CDNs use several techniques to redirect a client to a server (surrogate). A request routing function within a CDN is responsible for receiving content requests from user agents, obtaining and maintaining necessary information about a set of candidate surrogates, and for selecting and redirecting the user agent to the appropriate surrogate. One common way is relying on the DNS system, but there are many other ways, see [RFC3568].

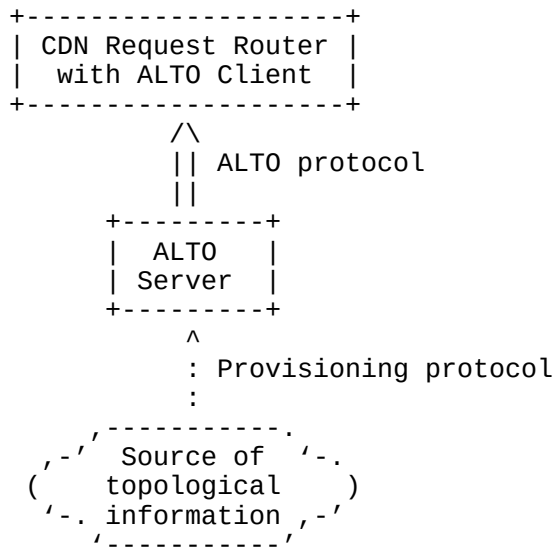

Figure 20: Use of ALTO information for CDN request routing

In order to derive the optimal benefit from a CDN it is preferable to deliver content from the servers (caches) that are "closest" to the end user requesting the content. "closest" may be as simple as geographical or IP topology distance, but it may also consider other combinations of metrics and CDN or Internet Service Provider (ISP) policies. As illustrated in Figure 20, ALTO could provide this information.

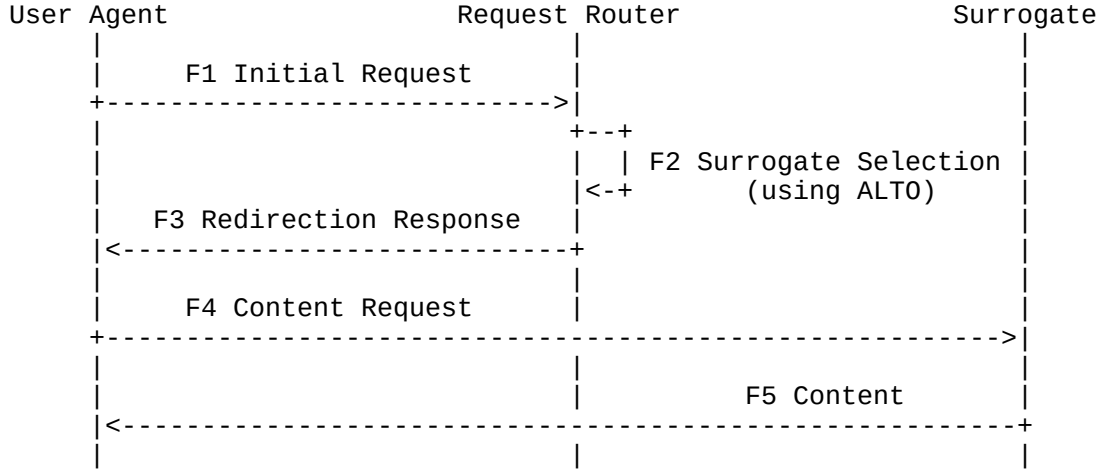

Figure 21: Example of CDN surrogate selection

Figure 21 illustrates the interaction between a user agent, a request router, and a surrogate for the delivery of content in a single CDN. As explained in [I-D.jenkins-alto-cdn-use-cases], the user agent makes an initial request to the CDN (F1). This may be an application-level request (e.g., HTTP) or a DNS request. In the second step (F2), the request router selects an appropriate surrogate (or set of surrogates) based on the user agent's (or its proxy's) IP address, the request router's knowledge of the network topology (which can be obtained by ALTO) and reachability cost between CDN caches and end users, and any additional CDN policies. Then (F3), the request router responds to the initial request with an appropriate response containing a redirection to the selected cache, for example by returning an appropriate DNS A/AAAA record, a HTTP 302 redirect, etc. The user agent uses this information to connect directly to the surrogate and request the desired content (F4), which is then delivered (F5).

5.1.2. Applicability of ALTO

The most simple use case for ALTO in a CDN context is to improve the selection of a CDN surrogate or origin. In this case, the CDN makes use of an ALTO server to choose a better CDN surrogate or origin than would otherwise be the case. Although it is possible to obtain raw network map and cost information in other ways, for example passively listening to the ISP's routing protocols or use of active probing, the use of an ALTO service to expose that information may provide additional control to the ISP over how their network map/cost is exposed. Additionally it may enable the ISP to maintain a functional separation between their routing plane and network map computation functions. This may be attractive for a number of reasons, for example:

- The ALTO service could provide a filtered view of the network and/or cost map that relates to CDN locations and their proximity to end users, for example to allow the ISP to control the level of topology detail they are willing to share with the CDN.

- The ALTO service could apply additional policies to the network map and cost information to provide a CDN-specific view of the network map/cost, for example to allow the ISP to encourage the CDN to use network links that would not ordinarily be preferred by a Shortest Path First routing calculation.

- The routing plane may be operated and controlled by a different operational entity (even within a single ISP) to the CDN. Therefore, the CDN may not be able to passively listen to routing protocols, nor may it have access to other network topology data (e.g., inventory databases).

When CDN servers are deployed outside of an ISP's network or in a small number of central locations within an ISP's network, a simplified view of the ISP's topology or an approximation of proximity is typically sufficient to enable the CDN to serve end users from the optimal server/location. As CDN servers are deployed deeper within ISP networks it becomes necessary for the CDN to have more detailed knowledge of the underlying network topology and costs between network locations in order to enable the CDN to serve end users from the most optimal servers for the ISP.

The request router in a CDN will typically also take into account criteria and constraints that are not related to network topology, such as the current load of CDN surrogates, content owner policies, end user subscriptions, etc. This document only discusses use of ALTO for network information.

A general issue for CDNs is that the CDN logic has to match the client's IP address with the closest CDN surrogate, both for DNS or HTTP redirect based approaches (see, for instance, [I-D.penno-alto-cdn]). This matching is not trivial, for instance, in DNS based approaches, where the IP address of the DNS original requester is unknown (see [I-D.vandergaast-edns-client-ip] for a discussion of this and a solution approach).

In addition to use by a single CDN, ALTO can also be used in scenarios that interconnect several CDNs. This use case is detailed in [I-D.seedorf-cdni-request-routing-alto].

5.2. Deployment Recommendations

5.2.1. ALTO Services

In its simplest form an ALTO server would provide an ISP with the capability to offer a service to a CDN that provides network map and cost information. The CDN can use that data to enhance its surrogate and/or origin selection. If an ISP offers an ALTO network and cost map service to expose a cost mapping/ranking between end user IP subnets (within that ISP's network) and CDN surrogate IP subnets/ locations, periodic updates of the maps may be needed. As introduced in Section 3.3), it is common for broadband subscribers to obtain their IP addresses dynamically and in many deployments the IP subnets allocated to a particular network region can change relatively frequently, even if the network topology itself is reasonably static.

An alternative would be to use the ALTO Endpoint Cost Service (ECS): When an end user request a given content, the CDN request router issues an ECS request with the endpoint address (IPv4/IPv6) of the end user (content requester) and the set of endpoint addresses of the

surrogate (content targets). The ALTO server receives the request and ranks the list of content targets addresses based on their distance from the content requester. Once the request router obtained from the ALTO Server the ranked list of locations (for the specific user), it can incorporate this information into its selection mechanisms in order to point the user to the most appropriate surrogate.

Since CDNs operate in a controlled environment, the ALTO network/cost map service and ECS have a similar level of security and confidentiality of network-internal information. However, the network/cost map service and ECS differ in the way the ALTO service is delivered and address a different set of requirements in terms of topology information and network operations.

If a CDN already has means to model connectivity policies, the map-based approaches could possibly be integrated into that. If the ECS service is preferred, a request router that uses ECS could cache the results of ECS queries for later usage in order to address the scalability limitations of ECS and to reduce the number of transactions between CDN and ALTO server. The ALTO server may indicate in the reply message how long the content of the message is to be considered reliable and insert a lifetime value that will be used by the CDN in order to cache (and then flush or refresh) the entry.

5.2.2. Guidance Considerations

In the following it is discussed how a CDN could make use of ALTO services.

In one deployment scenario, ALTO could expose ISP end user reachability to a CDN. The request router needs to have information which end user IP subnets are reachable via which networks or network locations. The network map services offered by ALTO could be used to expose this topology information while avoiding routing plane peering between the ISP and the CDN. For example, if CDN surrogates are deployed within the access or aggregation network, the ISP is likely to want to utilize the surrogates deployed in the same access/ aggregation region in preference to surrogates deployed elsewhere, in order to alleviate the cost and/or improve the user experience.

In addition, CDN surrogates could also use ALTO guidance, e.g., if there is more than one upstream source of content or several origins. In this case, ALTO could help a surrogate with the decision which upstream source to use. This specific variant of using ALTO is not further detailed in this document.

If content can be provided by several CDNs, there may be a need to interconnect these CDNs. In this case, ALTO can be uses as interface [I-D.seedorf-cdni-request-routing-alto], in particular for footprint and capabilities advertisement interface.

Other and more advanced scenarios of deploying ALTO are also listed in [I-D.jenkins-alto-cdn-use-cases] and [I-D.penno-alto-cdn].

The granularity of ALTO information required depends on the specific deployment of the CDN. For example, an over-the-top CDN whose surrogates are deployed only within the Internet "backbone" may only require knowledge of which end user IP subnets are reachable via which ISPs' networks, whereas a CDN deployed within a particular ISP's network requires a finer granularity of knowledge.

ALTO server ranks addresses based on topology information it acquires from the network. By default, according to [RFC7285], distance in ALTO represents an abstract "routingcost" that can be computed for instance from routing protocol information. But an ALTO server may also take into consideration other criteria or other information sources for policy, state, and performance information (e.g., geo-location), as explained in Section 3.2.1.

The different methods and algorithms through which the ALTO server computes topology information and rankings is out of the scope of this document. If rankings are based on routing protocol information, it is obvious that network events may impact the ranking computation. Due to internal redundancy and resilience mechanisms inside current networks, most of the network events happening in the infrastructure will be handled internally in the network, and they should have limited impact on a CDN. However, catastrophic events such as main trunks failures or backbone partitioning will have to take into account by the ALTO server to redirect traffic away from the impacted area.

An ALTO server implementation may want to keep state about ALTO clients so to inform and signal to these clients when a major network event happened, e.g., by a notification mechanism. In a CDN/ALTO interworking architecture with few CDN components interacting with the ALTO server there are less scalability issues in maintaining state about clients in the ALTO server, compared to ALTO guidance to any Internet user.

6. Other Use Cases

This section briefly surveys and references other use cases that have been tested or suggested for ALTO deployments.

6.1. Application Guidance in Virtual Private Networks (VPNs)

Virtual Private Network (VPN) technology is widely used in public and private networks to create groups of users that are separated from other users of the network and allows these users to communicate among them as if they were on a private network. Network Service Providers (NSPs) offer different types of VPNs. [RFC4026] distinguishes between Layer 2 VPN (L2VPN) and Layer 3 VPN (L3VPN) using different sub-types. In the following, the term "VPN" is used to refer to provider supplied virtual private networking.

From the perspective of an application at an endpoint, a VPN may not be very different to any other IP connectivity solution, but there are a number of specific applications that could benefit from ALTO topology exposure and guidance in VPNs. Similar like in the general Internet, one advantage is that applications do not have to perform excessive measurements on their own. For instance, potential use cases for ALTO application guidance in VPNs environments are:

- Enterprise application optimization: Enterprise customers often run distributed applications that exchange large amounts of data, e.g., for synchronization of replicated data bases. Both for placement of replicas as well as for the scheduling of transfers insight into network topology information could be useful.

- Private cloud computing solution: An enterprise customer could run own data centers at the four sites. The cloud management system could want to understand the network costs between different sites for intelligent routing and placement decisions of Virtual Machines (VMs) among the VPN sites.

- Cloud-bursting: One or more VPN endpoints could be located in a public cloud. If an enterprise customer needs additional resources, they could be provided by a public cloud, which is accessed through the VPN. Network topology awareness would help to decide in which data center of the public cloud those resources should be allocated.

These examples focus on enterprises, which are typical users of VPNs. VPN customers typically have no insight into the network topology that transports the VPN. Similar like in other ALTO use cases, better-than-random application-level decisions would be enabled by an ALTO server offered by the NSP, as illustrated in Figure 22.

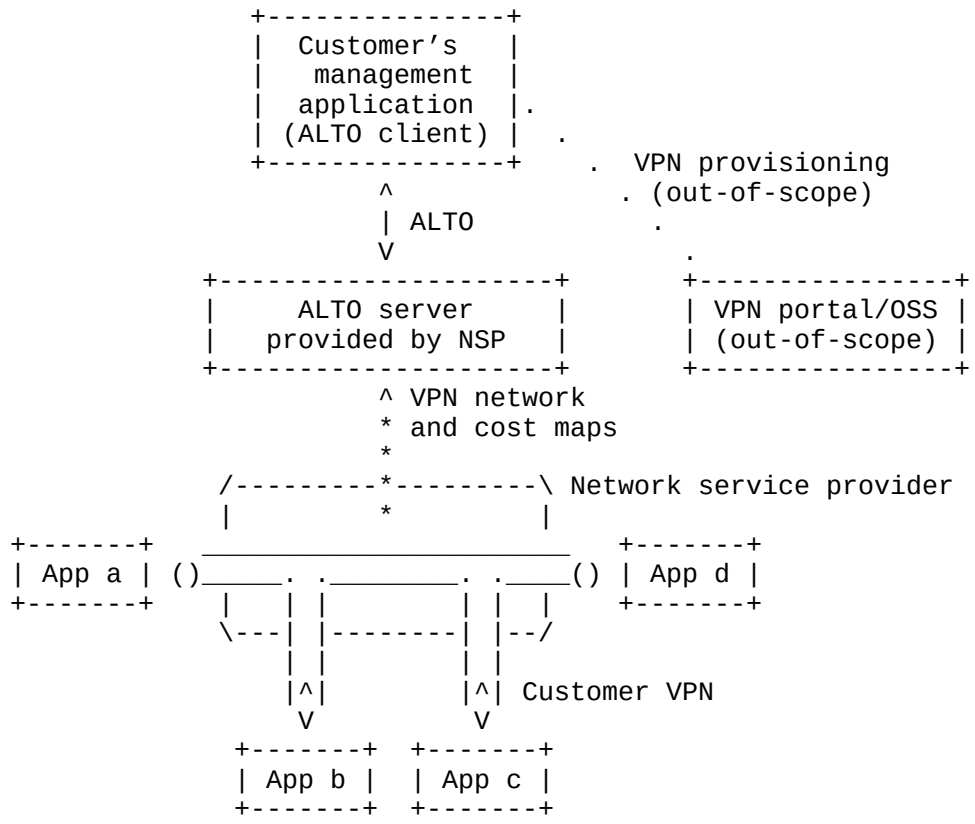

Figure 22: Using ALTO in VPNs

A common characteristic of these use cases is that applications will not necessarily run in the public Internet, and that the relationship between the provider and customer of the VPN is rather well-defined. Since VPNs run often in a managed environment, an ALTO server may have access to topology information (e.g., traffic engineering data) that would not be available for the public Internet, and it may expose it to the customer of the VPN only.

Also, a VPN will not necessarily be static. The customer could possibly modify the VPN and add new VPN sites by a Web portal, network management systems, or other Operation Support Systems (OSS) solutions. Prior to adding a new VPN site, an application will not be have connectivity to that site, i.e., an ALTO server could offer access to information that an application cannot measure on its own (e.g., expected delay to a new VPN site).

The VPN use cases, requirements, and solutions are further detailed in [I-D.scharf-alto-vpn-service].

6.2. In-Network Caching

Deployment of intra-domain P2P caches has been proposed for a cooperations between the network operator and the P2P service providers, e.g., to reduce the bandwidth consumption in access networks [I-D.deng-alto-p2pcache].

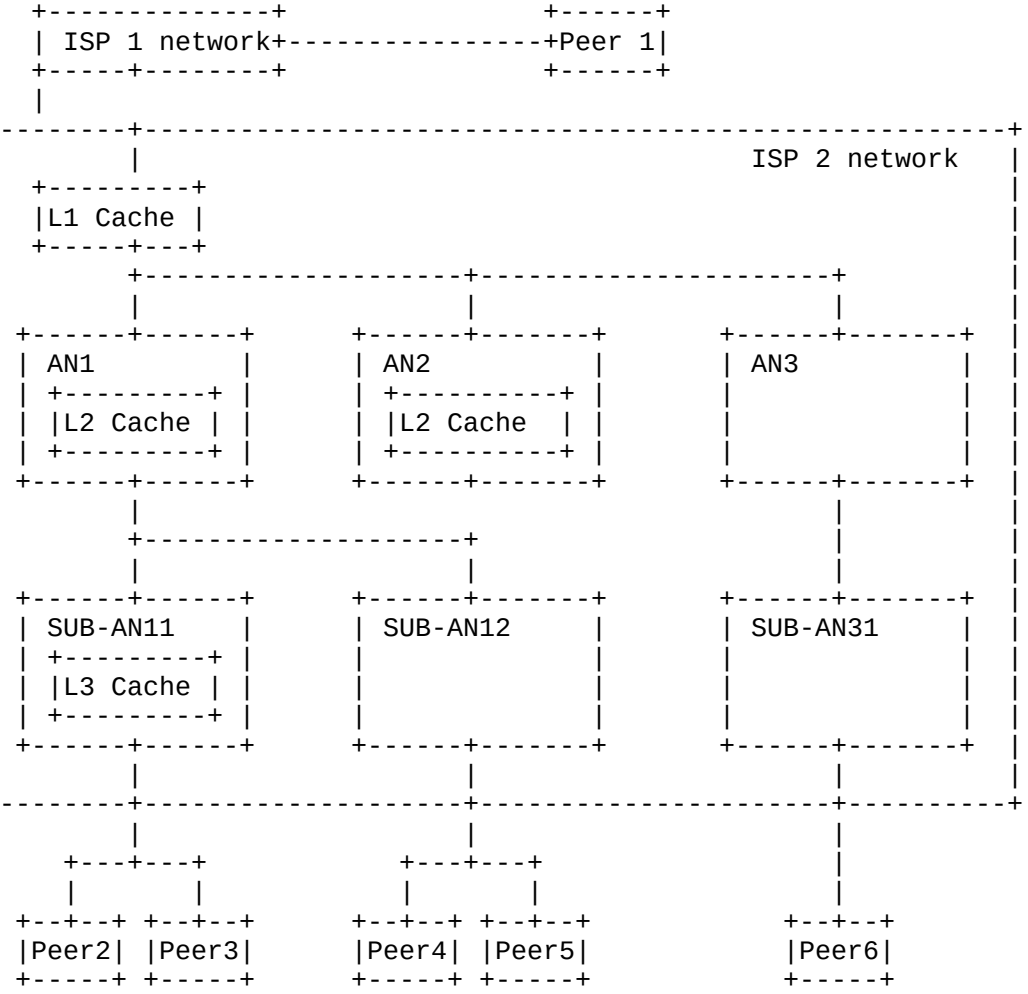

Figure 23: General architecture of intra-ISP caches

Figure 23 depicts the overall architecture of a potential P2P cache deployments inside an ISP 2 with various access network types. As shown in the figure, P2P caches may be deployed at various levels, including the interworking gateway linking with other ISPs, internal access network gateways linking with different types of accessing networks (e.g. WLAN, cellular and wired), and even within an accessing network at the entries of individual WLAN sub-networks. Moreover, depending on the network context and the operator's policy, each cache can be a Forwarding Cache or a Bidirectional Cache [I-D.deng-alto-p2pcache].

In such a cache architecture, the locations of caches could be used as dividers of different PIDs to guide intra-ISP network abstraction and mark costs among them according to the location and type of relevant caches.

Further details and deployment considerations can be found in [I-D.deng-alto-p2pcache].

6.3. Other Application-based Network Operations

An ALTO server can be part of an overall framework for Application-Based Network Operations (ABNO) [I-D.farrkingel-pce-abno-architecture] that brings together different technologies for gathering information about the resources available in a network, for consideration of topologies and how those topologies map to underlying network resources, for requesting path computation, and for provisioning or reserving network resources. Such an architecture may include additional components such as a Path Computation Element (PCE) for on-demand and application-specific reservation of network connectivity, reliability, and resources (such as bandwidth). Some use cases how to leverage ALTO for joint network and application-layer optimization are explained in [I-D.farrkingel-pce-abno-architecture].

7. Security Considerations

Security concerns were extensively discussed from the very beginning of the development of the ALTO protocol, and they have been considered in detail in the ALTO requirements document [RFC6708] as well as in the ALTO protocol specification document [RFC7285]. The two main security concerns are related to the unwanted disclosure of information through ALTO and the negative impact of specially crafted, wrong ("faked") guidance presented to an ALTO client. In addition to this, the usual concerns related to the operation of any networked application apply.

This section focuses on the peer-to-peer use case, which is - from a security perspective - probably the most difficult ALTO use case that has been considered. Special attention is given to the two main security concerns.

7.1. ALTO as a Protocol Crossing Trust Boundaries

The optimization of peer-to-peer applications was the first use case and the impetus for the development of the ALTO protocol, in particular file sharing applications such as BitTorrent [RFC5594].

As explained in Section 4.1.1, for the publisher of the ALTO information (i.e., the ALTO server operator) it is not always clear who is in charge of the P2P application overlay. Some P2P applications do not have any central control entity and the whole overlay consists only of the peers, which are under control of the individual users. Other P2P applications may have some control entities such as super peers or trackers, but these may be located in foreign countries and under the control of unknown organizations. As outlined in Section 4.2.2, in some scenarios it may be very beneficial to forward ALTO information to such trackers, super peers, etc. located in remote networks. This somewhat intransparent situation is aggravated by the vast number of different P2P applications which are evolving quickly and often without any coordination with the network operators.

In summary it can be said that in many instances of the P2P use case, the ALTO protocol bridges the border between the "managed" IP network infrastructure under strict administrative control and one or more "unmanaged" application overlays, i.e., overlays for which it is hard to tell who is in charge of them. This is different to more controlled environments (e.g., in the CDN use case), in which bilateral agreements between the producer and consumer of guidance are possible.

7.2. Information Leakage from the ALTO Server

An ALTO server will be provisioned with information about the ISP’s network and possibly also with information about neighboring ISPs. This information (e.g., network topology, business relations, etc.) is often considered to be confidential to the ISP and can include very sensitive information. ALTO does not require any particular level of details of information disclosure, and hence the provider should evaluate how much information is revealed and the associated risks.

Furthermore, if the ALTO information is very fine grained, it may also be considered sensitive with respect to user privacy. For

example, consider a hypothetical endpoint property "provisioned access link bandwidth" or "access technology (ADSL, VDSL, FTTH, etc.)" and an ALTO service that publishes this property for individual IP addresses. This information could not only be used for traffic optimization but, for example, also for targeted advertising to residential users with exceptionally good (or bad) connectivity, such as special banner ads. For an advertisement system it would be more complex to obtain such information otherwise, e.g., by bandwidth probing.

Different scenarios related to the unwanted disclosure of an ALTO server's information have been itemized and categorized in RFC 6708, Section 5.2.1., cases (1)-(3) [RFC6708].

In some use cases it is not possible to use access control (see Section 7.3) to limit the distribution of ALTO knowledge to a small set of trusted clients. In these scenarios it seems tempting not to use network maps and cost maps at all, and instead completely rely on endpoint cost service and endpoint ranking in the ALTO server. While this practice may indeed reduce the amount of information that is disclosed to an individual ALTO client, some issues should be considered: First, when using the map based apporach, it is trivial to analyze the maximum amount of information that could be disclosed to a client: the full maps. In contrast, when providing endpoint cost service only, the ALTO server operator could be prone to a false feeling of security, while clients use repeated queries and/or collaboration to gather more information than they are expected to get (see Section 5.2.1., case (3) in [RFC6708]). Second, the endpoint cost service reveals more information about the user or application behavior to the ALTO server, e.g., which other hosts are considered as peers for the exchange of a significant amount of data (see Section 5.2.1., cases (4)-(6) in [RFC6708]).

Consequently, users may be more reluctant to use the ALTO service at all if it is based on the endpoint property service instead of providing network and cost maps. Given that some popular P2P applications are sometimes used for purposes such as distribution of files without the explicit permission from the copyright owner, it may also be in the interest of the ALTO server operator that an ALTO server cannot infer the behavior of the application to be optimized. One possible conclusion could be to publish network and cost maps through ALTO that are so coarse-grained that they do not violate the network operator's or the user's interests.

In other use cases in more controlled environments (e.g., in the CDN use case) bilateral agreements, access control (see Section 7.3), and encryption could be used to reduce the risk of information leakage.

7.3. ALTO Server Access

Depending on the use case of ALTO, it may be desired to apply access restrictions to an ALTO server, i.e., by requiring client authentication. According to [RFC7285], ALTO requires that HTTP Digest Authentication is supported, in order to achieve client authentication and possibly to limit the number of parties with whom ALTO information is directly shared. TLS Client Authentication may also be supported.

In general, well-known security management techniques and best current practices [RFC4778] for operational ISP infrastructure also apply to an ALTO service, including functions to protect the system from unauthorized access, key management, reporting security-relevant events, and authorizing user access and privileges.

For peer-to-peer applications, a potential deployment scenario is that an ALTO server is solely accessible by peers from the ISP network (as shown in Figure 15). For instance, the source IP address can be used to grant only access from that ISP network to the server. This will "limit" the number of peers able to attack the server to the user’s of the ISP (however, including botnet computers).

If the ALTO server has to be accessible by parties not located in the ISP’s network (see Figure 16), e.g., by a third-party tracker or by a CDN system outside the ISP’s network, the access restrictions have to be looser. In the extreme case, i.e., no access restrictions, each and every host in the Internet can access the ALTO server. This might no be the intention of the ISP, as the server is not only subject to more possible attacks, but also the server load could increase, since possibly more ALTO clients have to be served.

There are also use cases where the access to the ALTO server has to be much more strictly controlled, i. e., where an authentication and authorization of the ALTO client to the server may be needed. For instance, in case of CDN optimization the provider of an ALTO service as well as potential users are possibly well-known. Only CDN entities may need ALTO access; access to the ALTO servers by residential users may neither be necessary nor be desired.

Access control can also help to prevent Denial-of-Service attacks by arbitrary hosts from the Internet. Denial-of-Service (DoS) can both affect an ALTO server and an ALTO client. A server can get overloaded if too many requests hit the server, or if the query load of the server surpasses the maximum computing capacity. An ALTO client can get overloaded if the responses from the sever are, either intentionally or due to an implementation mistake, too large to be handled by that particular client.

7.4. Faking ALTO Guidance

The ALTO services enables an ALTO service provider to influence the behavior of network applications. An attacker who is able to generate false replies, or e.g. an attacker who can intercept the ALTO server discovery procedure, can provide faked ALTO guidance.

Here is a list of examples how the ALTO guidance could be faked and what possible consequences may arise:

Sorting: An attacker could change to sorting order of the ALTO guidance (given that the order is of importance, otherwise the ranking mechanism is of interest), i.e., declaring peers located outside the ISP as peers to be preferred. This will not pose a big risk to the network or peers, as it would mimic the "regular" peer operation without traffic localization, apart from the communication/processing overhead for ALTO. However, it could mean that ALTO is reaching the opposite goal of shuffling more data across ISP boundaries, incurring more costs for the ISP.

Preference of a single peer: A single IP address (thus a peer) could be marked as to be preferred all over other peers. This peer can be located within the local ISP or also in other parts of the Internet (e.g., a web server). This could lead to the case that quite a number of peers to trying to contact this IP address, possibly causing a Denial-of-Service (DoS) attack.

It has not yet been investigated how a faked or wrong ALTO guidance by an ALTO server can impact the operation of the network and also the applications, e.g., peer-to-peer applications.

8. IANA Considerations

This document makes no specific request to IANA.

9. Conclusion

This document discusses how the ALTO protocol can be deployed in different use cases and provides corresponding guidance and recommendations to network administrators and application developers.

10. Acknowledgments

This memo is the result of contributions made by several people:

- Xianghue Sun, Lee Kai, and Richard Yang contributed text on ISP deployment requirements and monitoring.

- Stefano Previdi contributed parts of the Section 5 on "Using ALTO for CDNs".

- Rich Woundy contributed text to Section 3.3.

- Lingli Deng, Wei Chen, Qiuchao Yi, and Yan Zhang contributed Section 6.2.

Thomas-Rolf Banniza, Vinayak Hegde, and Qin Wu provided very useful comments and reviewed the document.

Martin Stiemerling is partially supported by the CHANGE project (http://www.change-project.eu), a research project supported by the European Commission under its 7th Framework Program (contract no. 257422). The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the CHANGE project or the European Commission.

11. References

11.1. Normative References

[RFC5693] Seedorf, J. and E. Burger, "Application-Layer Traffic Optimization (ALTO) Problem Statement", RFC 5693, October 2009.

[RFC6708] Kiesel, S., Previdi, S., Stiemerling, M., Woundy, R., and Y. Yang, "Application-Layer Traffic Optimization (ALTO) Requirements", RFC 6708, September 2012.

[RFC7285] Alimi, R., Penno, R., Yang, Y., Kiesel, S., Previdi, S., Roome, W., Shalunov, S., and R. Woundy, "Application-Layer Traffic Optimization (ALTO) Protocol", RFC 7285, September 2014.

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

11.2. Informative References

[I-D.deng-alto-p2pcache] Lingli, D., Chen, W., Yi, Q., and Y. Zhang, "Considerations for ALTO with network-deployed P2P caches", draft-deng-alto-p2pcache-03 (work in progress), February 2014.

[I-D.farrkingel-pce-abno-architecture]
King, D. and A. Farrel, "A PCE-based Architecture for Application-based Network Operations", draft-farrkingel-pce-abno-architecture-16 (work in progress), January 2015.

[I-D.ietf-i2rs-architecture]
Atlas, A., Halpern, J., Hares, S., Ward, D., and T. Nadeau, "An Architecture for the Interface to the Routing System", draft-ietf-i2rs-architecture-08 (work in progress), January 2015.

[I-D.ietf-idr-ls-distribution]
Gredler, H., Medved, J., Previdi, S., Farrel, A., and S. Ray, "North-Bound Distribution of Link-State and TE Information using BGP", draft-ietf-idr-ls-distribution-10 (work in progress), January 2015.

[I-D.jenkins-alto-cdn-use-cases]
Niven-Jenkins, B., Watson, G., Bitar, N., Medved, J., and S. Previdi, "Use Cases for ALTO within CDNs", draft-jenkins-alto-cdn-use-cases-03 (work in progress), June 2012.

[I-D.kamei-p2p-experiments-japan]
Kamei, S., Momose, T., Inoue, T., and T. Nishitani, "ALTO-Like Activities and Experiments in P2P Network Experiment Council", draft-kamei-p2p-experiments-japan-09 (work in progress), October 2012.

[I-D.kiesel-alto-h12]
Kiesel, S. and M. Stiemerling, "ALTO H12", draft-kiesel-alto-h12-02 (work in progress), March 2010.

[I-D.kiesel-alto-xdom-disc]
Kiesel, S. and M. Stiemerling, "Application Layer Traffic Optimization (ALTO) Cross-Domain Server Discovery", draft-kiesel-alto-xdom-disc-00 (work in progress), July 2014.

[I-D.lee-alto-chinatelecom-trial]
Li, K. and G. Jian, "ALTO and DECADE service trial within China Telecom", draft-lee-alto-chinatelecom-trial-04 (work in progress), March 2012.

[I-D.penno-alto-cdn]
Penno, R., Medved, J., Alimi, R., Yang, R., and S. Previdi, "ALTO and Content Delivery Networks", draft-penno-alto-cdn-03 (work in progress), March 2011.

[I-D.scharf-alto-vpn-service]
Scharf, M., Gurbani, V., Soprovich, G., and V. Hilt, "The Virtual Private Network (VPN) Service in ALTO: Use Cases, Requirements and Extensions", draft-scharf-alto-vpn-service-02 (work in progress), February 2014.

[I-D.seedorf-cdni-request-routing-alto]
Seedorf, J., Yang, Y., and J. Peterson, "CDNI Footprint and Capabilities Advertisement using ALTO", draft-seedorf-cdni-request-routing-alto-07 (work in progress), June 2014.

[I-D.vandergaast-edns-client-ip]
Contavalli, C., Gaast, W., Leach, S., and D. Rodden, "Client IP information in DNS requests", draft-vandergaast-edns-client-ip-01 (work in progress), May 2010.

[I-D.wu-alto-te-metrics]
Wu, W., Yang, Y., Lee, Y., Dhody, D., and S. Randriamasy, "ALTO Traffic Engineering Cost Metrics", draft-wu-alto-te-metrics-05 (work in progress), October 2014.

[RFC3411] Harrington, D., Presuhn, R., and B. Wijnen, "An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks", STD 62, RFC 3411, December 2002.

[RFC3568] Barbir, A., Cain, B., Nair, R., and O. Spatscheck, "Known Content Network (CN) Request-Routing Mechanisms", RFC 3568, July 2003.

[RFC4026] Andersson, L. and T. Madsen, "Provider Provisioned Virtual Private Network (VPN) Terminology", RFC 4026, March 2005.

[RFC4778] Kaeo, M., "Operational Security Current Practices in Internet Service Provider Environments", RFC 4778, January 2007.

[RFC5594] Peterson, J. and A. Cooper, "Report from the IETF Workshop on Peer-to-Peer (P2P) Infrastructure, May 28, 2008", RFC 5594, July 2009.

[RFC5632] Griffiths, C., Livingood, J., Popkin, L., Woundy, R., and Y. Yang, "Comcast’s ISP Experiences in a Proactive Network Provider Participation for P2P (P4P) Technical Trial", RFC 5632, September 2009.

[RFC6020] Bjorklund, M., "YANG - A Data Modeling Language for the Network Configuration Protocol (NETCONF)", RFC 6020, October 2010.

[RFC6241] Enns, R., Bjorklund, M., Schoenwaelder, J., and A. Bierman, "Network Configuration Protocol (NETCONF)", RFC 6241, June 2011.

Authors' Addresses

Martin Stiemerling
NEC Laboratories Europe
Kurfuerstenanlage 36
Heidelberg 69115
Germany

Phone: +49 6221 4342 113
Fax: +49 6221 4342 155
Email: martin.stiemerling@neclab.eu
URI: http://ietf.stiemerling.org

Sebastian Kiesel
University of Stuttgart Information Center
Networks and Communication Systems Department
Allmandring 30
Stuttgart 70550
Germany

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

Stefano Previdi
Cisco Systems, Inc.
Via Del Serafico 200
Rome 00191
Italy

Email: sprevidi@cisco.com

Michael Scharf
Alcatel-Lucent Bell Labs
Lorenzstrasse 10
Stuttgart 70435
Germany

Email: michael.scharf@alcatel-lucent.com