

Internet Engineering Task Force (IETF)
Request for Comments: 8686
Category: Standards Track
ISSN: 2070-1721

S. Kiesel
University of Stuttgart
M. Stiemerling
H-DA
February 2020

Application-Layer Traffic Optimization (ALTO) Cross-Domain Server Discovery

Abstract

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 capable of providing a desired resource. ALTO is realized by a client-server protocol. Before an ALTO client can ask for guidance, it needs to discover one or more ALTO servers that can provide suitable guidance.

In some deployment scenarios, in particular if the information about the network topology is partitioned and distributed over several ALTO servers, it may be necessary to discover an ALTO server outside of the ALTO client's own network domain, in order to get appropriate guidance. This document details applicable scenarios, itemizes requirements, and specifies a procedure for ALTO cross-domain server discovery.

Technically, the procedure specified in this document takes one IP address or prefix and a U-NAPTR Service Parameter (typically, "ALTO:https") as parameters. It performs DNS lookups (for NAPTR resource records in the "in-addr.arpa." or "ip6.arpa." trees) and returns one or more URIs of information resources related to that IP address or prefix.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc8686>.

Copyright Notice

Copyright (c) 2020 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction
 - 1.1. Terminology and Requirements Language

2. ALTO Cross-Domain Server Discovery Procedure: Overview
3. ALTO Cross-Domain Server Discovery Procedure: Specification
 - 3.1. Interface
 - 3.2. Step 1: Prepare Domain Name for Reverse DNS Lookup
 - 3.3. Step 2: Prepare Shortened Domain Names
 - 3.4. Step 3: Perform DNS U-NAPTR Lookups
 - 3.5. Error Handling
4. Using the ALTO Protocol with Cross-Domain Server Discovery
 - 4.1. Network and Cost Map Service
 - 4.2. Map-Filtering Service
 - 4.3. Endpoint Property Service
 - 4.4. Endpoint Cost Service
 - 4.5. Summary and Further Extensions
5. Implementation, Deployment, and Operational Considerations
 - 5.1. Considerations for ALTO Clients
 - 5.2. Considerations for Network Operators
6. Security Considerations
 - 6.1. Integrity of the ALTO Server's URI
 - 6.2. Availability of the ALTO Server Discovery Procedure
 - 6.3. Confidentiality of the ALTO Server's URI
 - 6.4. Privacy for ALTO Clients
7. IANA Considerations
8. References
 - 8.1. Normative References
 - 8.2. Informative References
- Appendix A. Solution Approaches for Partitioned ALTO Knowledge
 - A.1. Classification of Solution Approaches
 - A.2. Discussion of Solution Approaches
 - A.3. The Need for Cross-Domain ALTO Server Discovery
 - A.4. Our Solution Approach
 - A.5. Relation to the ALTO Requirements
- Appendix B. Requirements for Cross-Domain Server Discovery
 - B.1. Discovery Client Application Programming Interface
 - B.2. Data Storage and Authority Requirements
 - B.3. Cross-Domain Operations Requirements
 - B.4. Protocol Requirements
 - B.5. Further Requirements
- Appendix C. ALTO and Tracker-Based Peer-to-Peer Applications
 - C.1. A Generic Tracker-Based Peer-to-Peer Application
 - C.2. Architectural Options for Placing the ALTO Client
 - C.3. Evaluation
 - C.4. Example
- Acknowledgments
- Authors' Addresses

1. Introduction

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 capable of providing a desired resource [RFC5693]. ALTO is realized by an HTTP-based client-server protocol [RFC7285], which can be used in various scenarios [RFC7971].

The ALTO base protocol document [RFC7285] specifies the communication between an ALTO client and one ALTO server. In principle, the client may send any ALTO query. For example, it might ask for the routing cost between any two IP addresses, or it might request network and cost maps for the whole network, which might be the worldwide Internet. It is assumed that the server can answer any query, possibly with some kind of default value if no exact data is known.

No special provisions were made for deployment scenarios with multiple ALTO servers, with some servers having more accurate information about some parts of the network topology while others have better information about other parts of the network ("partitioned knowledge"). Various ALTO use cases have been studied in the context of such scenarios. In some cases, one cannot assume that a topologically nearby ALTO server (e.g., a server discovered with the procedure specified in [RFC7286]) will always provide useful information to the client. One such scenario is detailed in Appendix C. Several solution approaches, such as redirecting a

client to a server that has more accurate information or forwarding the request to such a server on behalf of the client, have been proposed and analyzed (see Appendix A), but no solution has been specified so far.

Section 3 of this document specifies the "ALTO Cross-Domain Server Discovery Procedure" for client-side usage in these scenarios. An ALTO client that wants to send an ALTO query related to a specific IP address or prefix X may call this procedure with X as a parameter. It will use Domain Name System (DNS) lookups to find one or more ALTO servers that can provide a competent answer. The above wording "related to" was intentionally kept somewhat unspecific, as the exact semantics depends on the ALTO service to be used; see Section 4.

Those who are in control of the "reverse DNS" for a given IP address or prefix (i.e., the corresponding subdomain of "in-addr.arpa." or "ip6.arpa.") -- typically an Internet Service Provider (ISP), a corporate IT department, or a university's computing center -- may add resource records to the DNS that point to one or more relevant ALTO servers. In many cases, it may be an ALTO server run by that ISP or IT department, as they naturally have good insight into routing costs from and to their networks. However, they may also refer to an ALTO server provided by someone else, e.g., their upstream ISP.

1.1. Terminology and Requirements Language

This document makes use of the ALTO terminology defined in RFC 5693 [RFC5693].

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. ALTO Cross-Domain Server Discovery Procedure: Overview

This section gives a non-normative overview of the ALTO Cross-Domain Server Discovery Procedure. The detailed specification will follow in the next section.

This procedure was inspired by "Location Information Server (LIS) Discovery Using IP Addresses and Reverse DNS" [RFC7216] and reuses parts of the basic ALTO Server Discovery Procedure [RFC7286].

The basic idea is to use the Domain Name System (DNS), more specifically the "in-addr.arpa." or "ip6.arpa." trees, which are mostly used for "reverse mapping" of IP addresses to host names by means of PTR resource records. There, URI-enabled Naming Authority Pointer (U-NAPTR) resource records [RFC4848], which allow the mapping of domain names to Uniform Resource Identifiers (URIs), are installed as needed. Thereby, it is possible to store a mapping from an IP address or prefix to one or more ALTO server URIs in the DNS.

The ALTO Cross-Domain Server Discovery Procedure is called with one IP address or prefix and a U-NAPTR Service Parameter [RFC4848] as parameters.

The service parameter is usually set to "ALTO:https". However, other parameter values may be used in some scenarios -- e.g., "ALTO:http" to search for a server that supports unencrypted transmission for debugging purposes, or other application protocol or service tags if applicable.

The procedure performs DNS lookups and returns one or more URIs of information resources related to said IP address or prefix, usually the URIs of one or more ALTO Information Resource Directories (IRDs; see Section 9 of [RFC7285]). The U-NAPTR records also provide preference values, which should be considered if more than one URI is returned.

The discovery procedure sequentially tries two different lookup strategies. First, an ALTO-specific U-NAPTR record is searched in the "reverse tree" -- i.e., in subdomains of "in-addr.arpa." or "ip6.arpa." corresponding to the given IP address or prefix. If this lookup does not yield a usable result, the procedure tries further lookups with truncated domain names, which correspond to shorter prefix lengths. The goal is to allow deployment scenarios that require fine-grained discovery on a per-IP basis, as well as large-scale scenarios where discovery is to be enabled for a large number of IP addresses with a small number of additional DNS resource records.

3. ALTO Cross-Domain Server Discovery Procedure: Specification

3.1. Interface

The procedure specified in this document takes two parameters, X and SP, where X is an IP address or prefix and SP is a U-NAPTR Service Parameter.

The parameter X may be an IPv4 or an IPv6 address or prefix in Classless Inter-Domain Routing (CIDR) notation (see [RFC4632] for the IPv4 CIDR notation and [RFC4291] for IPv6). Consequently, the address type AT is either "IPv4" or "IPv6". In both cases, X consists of an IP address A and a prefix length L. From the definitions of IPv4 and IPv6, it follows that syntactically valid values for L are $0 \leq L \leq 32$ when AT=IPv4 and $0 \leq L \leq 128$ when AT=IPv6. However, not all syntactically valid values of L are actually supported by this procedure; Step 1 (see below) will check for unsupported values and report an error if necessary.

For example, for X=198.51.100.0/24, we get AT=IPv4, A=198.51.100.0, and L=24. Similarly, for X=2001:0DB8::20/128, we get AT=IPv6, A=2001:0DB8::20, and L=128.

In the intended usage scenario, the procedure is normally always called with the parameter SP set to "ALTO:https". However, for general applicability and in order to support future extensions, the procedure MUST support being called with any valid U-NAPTR Service Parameter (see Section 4.5 of [RFC4848] for the syntax of U-NAPTR Service Parameters and Section 5 of the same document for information about the IANA registries).

The procedure performs DNS lookups and returns one or more URIs of information resources related to that IP address or prefix, usually the URIs of one or more ALTO Information Resource Directories (IRDs; see Section 9 of [RFC7285]). For each URI, the procedure also returns order and preference values (see Section 4.1 of [RFC3403]), which should be considered if more than one URI is returned.

During execution of this procedure, various error conditions may occur and have to be reported to the caller; see Section 3.5.

For the remainder of the document, we use the following notation for calling the ALTO Cross-Domain Server Discovery Procedure: `IRD_URI_X = XDOMDISC(X, "ALTO:https")`

3.2. Step 1: Prepare Domain Name for Reverse DNS Lookup

First, the procedure checks the prefix length L for unsupported values: If AT=IPv4 (i.e., if A is an IPv4 address) and $L < 8$, the procedure aborts and indicates an "unsupported prefix length" error to the caller. Similarly, if AT=IPv6 and $L < 32$, the procedure aborts and indicates an "unsupported prefix length" error to the caller. Otherwise, the procedure continues.

If AT=IPv4, the procedure will then produce a DNS domain name, which will be referred to as R32. This domain name is constructed according to the rules specified in Section 3.5 of [RFC1035], and it is rooted in the special domain "IN-ADDR.ARPA."

1: Address Type AT	2: Prefix Length L	3: MUST do 1st lookup	4: SHOULD do further lookups in that order
IPv4	32	R32	R24, R16, R8
IPv4	24 .. 31	R24	R16, R8
IPv4	16 .. 23	R16	R8
IPv4	8 .. 15	R8	(none)
IPv4	0 .. 7	(none, abort: unsupported prefix length)	
IPv6	128	R128	R64, R56, R48, R40, R32
IPv6	64 .. 127	R64	R56, R48, R40, R32
IPv6	56 .. 63	R56	R48, R40, R32
IPv6	48 .. 55	R48	R40, R32
IPv6	40 .. 47	R40	R32
IPv6	32 .. 39	R32	(none)
IPv6	0 .. 31	(none, abort: unsupported prefix length)	

Table 1: Perform DNS U-NAPTR lookups

Then, the domain name given in the 3rd column and the U-NAPTR Service Parameter SP with which the procedure was called (usually "ALTO:https") MUST be used for a U-NAPTR [RFC4848] lookup, in order to obtain one or more URIs (indicating protocol, host, and possibly path elements) for the ALTO server's Information Resource Directory (IRD). If such URIs can be found, the ALTO Cross-Domain Server Discovery Procedure returns that information to the caller and terminates successfully.

For example, the following two U-NAPTR resource records can be used for mapping "100.51.198.IN-ADDR.ARPA." (i.e., R24 from the example in the previous step) to the HTTPS URIs "https://alto1.example.net/ird" and "https://alto2.example.net/ird", with the former being preferred.

```
100.51.198.IN-ADDR.ARPA. IN NAPTR 100 10 "u" "ALTO:https"
"!.*!https://alto1.example.net/ird!" ""

100.51.198.IN-ADDR.ARPA. IN NAPTR 100 20 "u" "ALTO:https"
"!.*!https://alto2.example.net/ird!" ""
```

If no matching U-NAPTR records can be found, the procedure SHOULD try further lookups, using the domain names from the fourth column in the indicated order, until one lookup succeeds. If no IRD URI can be found after looking up all domain names from the 3rd and 4th columns, the procedure terminates unsuccessfully, returning an empty URI list.

3.5. Error Handling

The ALTO Cross-Domain Server Discovery Procedure may fail for several reasons.

If the procedure is called with syntactically invalid parameters or unsupported parameter values (in particular, the prefix length L; see Section 3.2), the procedure aborts, no URI list will be returned, and the error has to be reported to the caller.

The procedure performs one or more DNS lookups in a well-defined order (corresponding to descending prefix lengths, see Section 3.4) until one produces a usable result. Each of these DNS lookups might fail to produce a usable result, due to either a normal condition (e.g., a domain name exists, but no ALTO-specific NAPTR resource records are associated with it), a permanent error (e.g., nonexistent domain name), or a temporary error (e.g., timeout). In all three cases, and as long as there are further domain names that can be looked up, the procedure SHOULD immediately try to look up the next domain name (from Column 4 in the table given in Section 3.4). Only after all domain names have been tried at least once, the procedure MAY retry those domain names that had caused temporary lookup errors.

Generally speaking, ALTO provides advisory information for the optimization of applications (peer-to-peer applications, overlay networks, etc.), but applications should not rely on the availability of such information for their basic functionality (see Section 8.3.4.3 of [RFC7285]). Consequently, the speedy detection of an ALTO server, even though it may give less accurate answers than other servers, or the quick realization that there is no suitable ALTO server, is in general preferable to causing long delays by retrying failed queries. Nevertheless, if DNS queries have failed due to temporary errors, the ALTO Cross-Domain Server Discovery Procedure SHOULD inform its caller that DNS queries have failed for that reason and that retrying the discovery at a later point in time might give more accurate results.

4. Using the ALTO Protocol with Cross-Domain Server Discovery

Based on a modular design principle, ALTO provides several ALTO services, each consisting of a set of information resources that can be accessed using the ALTO protocol. The information resources that are available at a specific ALTO server are listed in its Information Resource Directory (IRD, see Section 9 of [RFC7285]). The ALTO protocol specification defines the following ALTO services and their corresponding information resources:

- * Network and Cost Map Service, see Section 11.2 of [RFC7285]
- * Map-Filtering Service, see Section 11.3 of [RFC7285]
- * Endpoint Property Service, see Section 11.4 of [RFC7285]
- * Endpoint Cost Service, see Section 11.5 of [RFC7285]

The ALTO Cross-Domain Server Discovery Procedure is most useful in conjunction with the Endpoint Property Service and the Endpoint Cost Service. However, for the sake of completeness, possible interaction with all four services is discussed below. Extension documents may specify further information resources; however, these are out of scope of this document.

4.1. Network and Cost Map Service

An ALTO client may invoke the ALTO Cross-Domain Server Discovery Procedure (as specified in Section 3) for an IP address or prefix X and get a list of one or more IRD URIs, including order and preference values: IRD_URI_X = XDOMDISC(X, "ALTO:https"). The IRD(s) referenced by these URIs will always contain a network and a cost map, as these are mandatory information resources (see Section 11.2 of [RFC7285]). However, the cost matrix may be very sparse. If, according to the network map, PID_X is the Provider-defined Identifier (PID; see Section 5.1 of [RFC7285]) that contains the IP address or prefix X, and PID_1, PID_2, PID_3, ... are other PIDs, the cost map may look like this:

From	To PID_1	PID_2	PID_X	PID_3
PID_1			92	

PID_2			6	
PID_X	46	3	1	19
PID_3			38	

Table 2: Cost Map

In this example, all cells outside Column X and Row X are unspecified. A cost map with this structure contains the same information as what could be retrieved using the Endpoint Cost Service, Cases 1 and 2 in Section 4.4. Accessing cells that are neither in Column X nor Row X may not yield useful results.

Trying to assemble a more densely populated cost map from several cost maps with this very sparse structure may be a nontrivial task, as different ALTO servers may use different PID definitions (i.e., network maps) and incompatible scales for the costs, in particular for the "routingcost" metric.

4.2. Map-Filtering Service

An ALTO client may invoke the ALTO Cross-Domain Server Discovery Procedure (as specified in Section 3) for an IP address or prefix X and get a list of one or more IRD URIs, including order and preference values: `IRD_URIS_X = XDOMDISC(X, "ALTO:https")`. These IRDs may provide the optional Map-Filtering Service (see Section 11.3 of [RFC7285]). This service returns a subset of the full map, as specified by the client. As discussed in Section 4.1, a cost map may be very sparse in the envisioned deployment scenario. Therefore, depending on the filtering criteria provided by the client, this service may return results similar to the Endpoint Cost Service, or it may not return any useful result.

4.3. Endpoint Property Service

If an ALTO client wants to query an Endpoint Property Service (see Section 11.4 of [RFC7285]) about an endpoint with IP address X or a group of endpoints within IP prefix X, respectively, it has to invoke the ALTO Cross-Domain Server Discovery Procedure (as specified in Section 3): `IRD_URIS_X = XDOMDISC(X, "ALTO:https")`. The result, `IRD_URIS_X`, is a list of one or more URIs of Information Resource Directories (IRDs, see Section 9 of [RFC7285]). Considering the order and preference values, the client has to check these IRDs for a suitable Endpoint Property Service and query it.

If the ALTO client wants to do a similar Endpoint Property query for a different IP address or prefix "Y", the whole procedure has to be repeated, as `IRD_URIS_Y = XDOMDISC(Y, "ALTO:https")` may yield a different list of IRD URIs. Of course, the results of individual DNS queries may be cached as indicated by their respective time-to-live (TTL) values.

4.4. Endpoint Cost Service

The optional ALTO Endpoint Cost Service (ECS; see Section 11.5 of [RFC7285]) provides information about costs between individual endpoints and also supports ranking. The ECS allows endpoints to be denoted by IP addresses or prefixes. The ECS is called with a list of one or more source IP addresses or prefixes, which we will call (S1, S2, S3, ...), and a list of one or more destination IP addresses or prefixes, called (D1, D2, D3, ...).

This specification distinguishes several cases, regarding the number of elements in the list of source and destination addresses, respectively:

1. Exactly one source address S1 and more than one destination addresses (D1, D2, D3, ...). In this case, the ALTO client has to invoke the ALTO Cross-Domain Server Discovery Procedure (as

specified in Section 3) with that single source address as a parameter: `IRD_URIS_S1 = XDOMDISC(S1,"ALTO:https")`. The result, `IRD_URIS_S1`, is a list of one or more URIs of Information Resource Directories (IRDs, see Section 9 of [RFC7285]). Considering the order and preference values, the client has to check these IRDs for a suitable Endpoint Cost Service and query it. The ECS is an optional service (see Section 11.5.1 of [RFC7285]), and therefore, it may well be that an IRD does not refer to an ECS.

Calling the Cross-Domain Server Discovery Procedure only once with the single source address as a parameter -- as opposed to multiple calls, e.g., one for each destination address -- is not only a matter of efficiency. In the given scenario, it is advisable to send all ECS queries to the same ALTO server. This ensures that the results can be compared (e.g., for sorting candidate resource providers), even when cost metrics lack a well-defined base unit -- e.g., the "routingcost" metric.

2. More than one source address (`S1, S2, S3, ...`) and exactly one destination address `D1`. In this case, the ALTO client has to invoke the ALTO Cross-Domain Server Discovery Procedure with that single destination address as a parameter:
`IRD_URIS_D1 = XDOMDISC(D1,"ALTO:https")`. The result, `IRD_URIS_D1`, is a list of one or more URIs of IRDs. Considering the order and preference values, the client has to check these IRDs for a suitable ECS and query it.
3. Exactly one source address `S1` and exactly one destination address `D1`. The ALTO client may perform the same steps as in Case 1, as specified above. As an alternative, it may also perform the same steps as in Case 2, as specified above.
4. More than one source address (`S1, S2, S3, ...`) and more than one destination address (`D1, D2, D3, ...`). In this case, the ALTO client should split the list of desired queries based on source addresses and perform separately for each source address the same steps as in Case 1, as specified above. As an alternative, the ALTO client may also group the list based on destination addresses and perform separately for each destination address the same steps as in Case 2, as specified above. However, comparing results between these subqueries may be difficult, in particular if the cost metric is a relative preference without a well-defined base unit (e.g., the "routingcost" metric).

See Appendix C for a detailed example showing the interaction of a tracker-based peer-to-peer application, the ALTO Endpoint Cost Service, and the ALTO Cross-Domain Server Discovery Procedure.

4.5. Summary and Further Extensions

Considering the four services defined in the ALTO base protocol specification [RFC7285], the ALTO Cross-Domain Server Discovery Procedure works best with the Endpoint Property Service (EPS) and the Endpoint Cost Service (ECS). Both the EPS and the ECS take one or more IP addresses as a parameter. The previous sections specify how the parameter for calling the ALTO Cross-Domain Server Discovery Procedure has to be derived from these IP addresses.

In contrast, the ALTO Cross-Domain Server Discovery Procedure seems less useful if the goal is to retrieve network and cost maps that cover the whole network topology. However, the procedure may be useful if a map centered at a specific IP address is desired (i.e., a map detailing the vicinity of said IP address or a map giving costs from said IP address to all potential destinations).

The interaction between further ALTO services (and their corresponding information resources) needs to be investigated and defined once such further ALTO services are specified in an extension document.

5. Implementation, Deployment, and Operational Considerations

5.1. Considerations for ALTO Clients

5.1.1. Resource-Consumer-Initiated Discovery

Resource-consumer-initiated ALTO server discovery (cf. ALTO requirement AR-32 [RFC6708]) can be seen as a special case of cross-domain ALTO server discovery. To that end, an ALTO client embedded in a resource consumer would have to perform the ALTO Cross-Domain Server Discovery Procedure with its own IP address as a parameter. However, due to the widespread deployment of Network Address Translators (NATs), additional protocols and mechanisms such as Session Traversal Utilities for NAT (STUN) [RFC5389] are usually needed to detect the client's "public" IP address before it can be used as a parameter for the discovery procedure. Note that a different approach for resource-consumer-initiated ALTO server discovery, which is based on DHCP, is specified in [RFC7286].

5.1.2. IPv4/v6 Dual Stack, Multihoming and Host Mobility

The procedure specified in this document can discover ALTO server URIs for a given IP address or prefix. The intention is that a third party (e.g., a resource directory) that receives query messages from a resource consumer can use the source address in these messages to discover suitable ALTO servers for this specific resource consumer.

However, resource consumers (as defined in Section 2 of [RFC5693]) may reside on hosts with more than one IP address -- for example, due to IPv4/v6 dual stack operation and/or multihoming. IP packets sent with different source addresses may be subject to different routing policies and path costs. In some deployment scenarios, it may even be required to ask different sets of ALTO servers for guidance. Furthermore, source addresses in IP packets may be modified en route by Network Address Translators (NATs).

If a resource consumer queries a resource directory for candidate resource providers, the locally selected (and possibly en-route-translated) source address of the query message -- as observed by the resource directory -- will become the basis for the ALTO server discovery and the subsequent optimization of the resource directory's reply. If, however, the resource consumer then selects different source addresses to contact returned resource providers, the desired better-than-random "ALTO effect" may not occur.

One solution approach for this problem is that a dual-stack or multihomed resource consumer could always use the same address for contacting the resource directory and all resource providers, thus overriding the operating system's automatic selection of source IP addresses. For example, when using the BSD socket API, one could always bind() the socket to one of the local IP addresses before trying to connect() to the resource directory or the resource providers, respectively. Another solution approach is to perform ALTO-influenced resource provider selection (and source-address selection) locally in the resource consumer, in addition to, or instead of, performing it in the resource directory. See Section 5.1.1 for a discussion of how to discover ALTO servers for local usage in the resource consumer.

Similarly, resource consumers on mobile hosts SHOULD query the resource directory again after a change of IP address, in order to get a list of candidate resource providers that is optimized for the new IP address.

5.1.3. Interaction with Network Address Translation

The ALTO Cross-Domain Server Discovery Procedure has been designed to enable the ALTO-based optimization of applications such as large-scale overlay networks, that span -- on the IP layer -- multiple administrative domains, possibly the whole Internet. Due to the widespread usage of Network Address Translators (NATs), it may well

be that nodes of the overlay network (i.e., resource consumers or resource providers) are located behind a NAT, maybe even behind several cascaded NATs.

If a resource directory is located in the public Internet (i.e., not behind a NAT) and receives a message from a resource consumer behind one or more NATs, the message's source address will be the public IP address of the outermost NAT in front of the resource consumer. The same applies if the resource directory is behind a different NAT than the resource consumer. The resource directory may call the ALTO Cross-Domain Server Discovery Procedure with the message's source address as a parameter. In effect, not the resource consumer's (private) IP address, but the public IP address of the outermost NAT in front of it, will be used as a basis for ALTO optimization. This will work fine as long as the network behind the NAT is not too big (e.g., if the NAT is in a residential gateway).

If a resource directory receives a message from a resource consumer and the message's source address is a "private" IP address [RFC1918], this may be a sign that both of them are behind the same NAT. An invocation of the ALTO Cross-Domain Server Discovery Procedure with this private address may be problematic, as this will only yield usable results if a DNS "split horizon" and DNSSEC trust anchors are configured correctly. In this situation, it may be more advisable to query an ALTO server that has been discovered using [RFC7286] or any other local configuration. The interaction between intradomain ALTO for large private domains (e.g., behind a "carrier-grade NAT") and cross-domain, Internet-wide optimization, is beyond the scope of this document.

5.2. Considerations for Network Operators

5.2.1. Flexibility vs. Load on the DNS

The ALTO Cross-Domain Server Discovery Procedure, as specified in Section 3, first produces a list of domain names (Steps 1 and 2) and then looks for relevant NAPTR records associated with these names, until a useful result can be found (Step 3). The number of candidate domain names on this list is a compromise between flexibility when installing NAPTR records and avoiding excess load on the DNS.

A single invocation of the ALTO Cross-Domain Server Discovery Procedure, with an IPv6 address as a parameter, may cause up to, but no more than, six DNS lookups for NAPTR records. For IPv4, the maximum is four lookups. Should the load on the DNS infrastructure caused by these lookups become a problem, one solution approach is to populate the DNS with ALTO-specific NAPTR records. If such records can be found for individual IP addresses (possibly installed using a wildcarding mechanism in the name server) or long prefixes, the procedure will terminate successfully and not perform lookups for shorter prefix lengths, thus reducing the total number of DNS queries. Another approach for reducing the load on the DNS infrastructure is to increase the TTL for caching negative answers.

On the other hand, the ALTO Cross-Domain Server Discovery Procedure trying to look up truncated domain names allows for efficient configuration of large-scale scenarios, where discovery is to be enabled for a large number of IP addresses with a small number of additional DNS resource records. Note that it expressly has not been a design goal of this procedure to give clients a means of understanding the IP prefix delegation structure. Furthermore, this specification does not assume or recommend that prefix delegations should preferably occur at those prefix lengths that are used in Step 2 of this procedure (see Section 3.3). A network operator that uses, for example, an IPv4 /18 prefix and wants to install the NAPTR records efficiently could either install 64 NAPTR records (one for each of the /24 prefixes contained within the /18 prefix), or they could try to team up with the owners of the other fragments of the enclosing /16 prefix, in order to run a common ALTO server to which only one NAPTR would point.

5.2.2. BCP 20 and Missing Delegations of the Reverse DNS

[RFC2317], also known as BCP 20, describes a way to delegate the "reverse DNS" (i.e., subdomains of "in-addr.arpa.") for IPv4 address ranges with fewer than 256 addresses (i.e., less than a whole /24 prefix). The ALTO Cross-Domain Server Discovery Procedure is compatible with this method.

In some deployment scenarios -- e.g., residential Internet access -- where customers often dynamically receive a single IPv4 address (and/or a small IPv6 address block) from a pool of addresses, ISPs typically will not delegate the "reverse DNS" to their customers. This practice makes it impossible for these customers to populate the DNS with NAPTR resource records that point to an ALTO server of their choice. Yet, the ISP may publish NAPTR resource records in the "reverse DNS" for individual addresses or larger address pools (i.e., shorter prefix lengths).

While ALTO is by no means technologically tied to the Border Gateway Protocol (BGP), it is anticipated that BGP will be an important source of information for ALTO and that the operator of the outermost BGP-enabled router will have a strong incentive to publish a digest of their routing policies and costs through ALTO. In contrast, an individual user or an organization that has been assigned only a small address range (i.e., an IPv4 prefix with a prefix length longer than /24) will typically connect to the Internet using only a single ISP, and they might not be interested in publishing their own ALTO information. Consequently, they might wish to leave the operation of an ALTO server up to their ISP. This ISP may install NAPTR resource records, which are needed for the ALTO Cross-Domain Server Discovery Procedure, in the subdomain of "in-addr.arpa." that corresponds to the whole /24 prefix (cf. R24 in Section 3.3 of this document), even if delegations in the style of BCP 20 or no delegations at all are in use.

6. Security Considerations

A high-level discussion of security issues related to ALTO is part of the ALTO problem statement [RFC5693]. A classification of unwanted information disclosure risks, as well as specific security-related requirements, can be found in the ALTO requirements document [RFC6708].

The remainder of this section focuses on security threats and protection mechanisms for the Cross-Domain ALTO Server Discovery Procedure as such. Once the ALTO server's URI has been discovered, and the communication between the ALTO client and the ALTO server starts, the security threats and protection mechanisms discussed in the ALTO protocol specification [RFC7285] apply.

6.1. Integrity of the ALTO Server's URI

Scenario Description

An attacker could compromise the ALTO server discovery procedure or the underlying infrastructure in such a way that ALTO clients would discover a "wrong" ALTO server URI.

Threat Discussion

The Cross-Domain ALTO Server Discovery Procedure relies on a series of DNS lookups, in order to produce one or more URIs. If an attacker were able to modify or spoof any of the DNS records, the resulting URIs could be replaced by forged URIs. This is probably the most serious security concern related to ALTO server discovery. The discovered "wrong" ALTO server might not be able to give guidance to a given ALTO client at all, or it might give suboptimal or forged information. In the latter case, an attacker could try to use ALTO to affect the traffic distribution in the network or the performance of applications (see also Section 15.1 of [RFC7285]). Furthermore, a hostile ALTO server could threaten user privacy (see also Case (5a) in Section 5.2.1 of [RFC6708]).

Protection Strategies and Mechanisms

The application of DNS security (DNSSEC) [RFC4033] provides a means of detecting and averting attacks that rely on modification of the DNS records while in transit. All implementations of the Cross-Domain ALTO Server Discovery Procedure MUST support DNSSEC or be able to use such functionality provided by the underlying operating system. Network operators that publish U-NAPTR resource records to be used for the Cross-Domain ALTO Server Discovery Procedure SHOULD use DNSSEC to protect their subdomains of "in-addr.arpa." and/or "ip6.arpa.", respectively. Additional operational precautions for safely operating the DNS infrastructure are required in order to ensure that name servers do not sign forged (or otherwise "wrong") resource records. Security considerations specific to U-NAPTR are described in more detail in [RFC4848].

In addition to active protection mechanisms, users and network operators can monitor application performance and network traffic patterns for poor performance or abnormalities. If it turns out that relying on the guidance of a specific ALTO server does not result in better-than-random results, the usage of the ALTO server may be discontinued (see also Section 15.2 of [RFC7285]).

Note

The Cross-Domain ALTO Server Discovery Procedure finishes successfully when it has discovered one or more URIs. Once an ALTO server's URI has been discovered and the communication between the ALTO client and the ALTO server starts, the security threats and protection mechanisms discussed in the ALTO protocol specification [RFC7285] apply.

A threat related to the one considered above is the impersonation of an ALTO server after its correct URI has been discovered. This threat and protection strategies are discussed in Section 15.1 of [RFC7285]. The ALTO protocol's primary mechanism for protecting authenticity and integrity (as well as confidentiality) is the use of HTTPS-based transport -- i.e., HTTP over TLS [RFC2818]. Typically, when the URI's host component is a host name, a further DNS lookup is needed to map it to an IP address before the communication with the server can begin. This last DNS lookup (for A or AAAA resource records) does not necessarily have to be protected by DNSSEC, as the server identity checks specified in [RFC2818] are able to detect DNS spoofing or similar attacks after the connection to the (possibly wrong) host has been established. However, this validation, which is based on the server certificate, can only protect the steps that occur after the server URI has been discovered. It cannot detect attacks against the authenticity of the U-NAPTR lookups needed for the Cross-Domain ALTO Server Discovery Procedure, and therefore, these resource records have to be secured using DNSSEC.

6.2. Availability of the ALTO Server Discovery Procedure

Scenario Description

An attacker could compromise the Cross-Domain ALTO Server Discovery Procedure or the underlying infrastructure in such a way that ALTO clients would not be able to discover any ALTO server.

Threat Discussion

If no ALTO server can be discovered (although a suitable one exists), applications have to make their decisions without ALTO guidance. As ALTO could be temporarily unavailable for many reasons, applications must be prepared to do so. However, the resulting application performance and traffic distribution will correspond to a deployment scenario without ALTO.

Protection Strategies and Mechanisms

Operators should follow best current practices to secure their DNS and ALTO servers (see Section 15.5 of [RFC7285]) against Denial-of-Service (DoS) attacks.

6.3. Confidentiality of the ALTO Server's URI

Scenario Description

An unauthorized party could invoke the Cross-Domain ALTO Server Discovery Procedure or intercept discovery messages between an authorized ALTO client and the DNS servers, in order to acquire knowledge of the ALTO server URI for a specific IP address.

Threat Discussion

In the ALTO use cases that have been described in the ALTO problem statement [RFC5693] and/or discussed in the ALTO working group, the ALTO server's URI as such has always been considered as public information that does not need protection of confidentiality.

Protection Strategies and Mechanisms

No protection mechanisms for this scenario have been provided, as it has not been identified as a relevant threat. However, if a new use case is identified that requires this kind of protection, the suitability of this ALTO server discovery procedure as well as possible security extensions have to be re-evaluated thoroughly.

6.4. Privacy for ALTO Clients

Scenario Description

An unauthorized party could eavesdrop on the messages between an ALTO client and the DNS servers and thereby find out the fact that said ALTO client uses (or at least tries to use) the ALTO service in order to optimize traffic from/to a specific IP address.

Threat Discussion

In the ALTO use cases that have been described in the ALTO problem statement [RFC5693] and/or discussed in the ALTO working group, this scenario has not been identified as a relevant threat. However, pervasive surveillance [RFC7624] and DNS privacy considerations [RFC7626] have seen significant attention in the Internet community in recent years.

Protection Strategies and Mechanisms

DNS over TLS [RFC7858] and DNS over HTTPS [RFC8484] provide means for protecting confidentiality (and integrity) of DNS traffic between a client (stub) and its recursive name servers, including DNS queries and replies caused by the ALTO Cross-Domain Server Discovery Procedure.

7. IANA Considerations

This document has no IANA actions.

8. References

8.1. Normative References

- [RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, DOI 10.17487/RFC1035, November 1987, <<https://www.rfc-editor.org/info/rfc1035>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC3403] Mealling, M., "Dynamic Delegation Discovery System (DDDS) Part Three: The Domain Name System (DNS) Database", RFC 3403, DOI 10.17487/RFC3403, October 2002, <<https://www.rfc-editor.org/info/rfc3403>>.
- [RFC3596] Thomson, S., Huitema, C., Ksinant, V., and M. Souissi, "DNS Extensions to Support IP Version 6", STD 88, RFC 3596, DOI 10.17487/RFC3596, October 2003, <<https://www.rfc-editor.org/info/rfc3596>>.

[RFC4848] Daigle, L., "Domain-Based Application Service Location Using URIs and the Dynamic Delegation Discovery Service (DDDS)", RFC 4848, DOI 10.17487/RFC4848, April 2007, <<https://www.rfc-editor.org/info/rfc4848>>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

8.2. Informative References

[ALTO-ANYCAST]

Kiesel, S. and R. Penno, "Application-Layer Traffic Optimization (ALTO) Anycast Address", Work in Progress, Internet-Draft, draft-kiesel-alto-ip-based-srv-disc-03, 1 July 2014, <<https://tools.ietf.org/html/draft-kiesel-alto-ip-based-srv-disc-03>>.

[ALTO4ALTO]

Kiesel, S., "Using ALTO for ALTO server selection", Work in Progress, Internet-Draft, draft-kiesel-alto-alto4alto-00, 5 July 2010, <<https://tools.ietf.org/html/draft-kiesel-alto-alto4alto-00>>.

[RFC1918] Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G. J., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, DOI 10.17487/RFC1918, February 1996, <<https://www.rfc-editor.org/info/rfc1918>>.

[RFC2317] Eidnes, H., de Groot, G., and P. Vixie, "Classless IN-ADDR.ARPA delegation", BCP 20, RFC 2317, DOI 10.17487/RFC2317, March 1998, <<https://www.rfc-editor.org/info/rfc2317>>.

[RFC2818] Rescorla, E., "HTTP Over TLS", RFC 2818, DOI 10.17487/RFC2818, May 2000, <<https://www.rfc-editor.org/info/rfc2818>>.

[RFC4033] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "DNS Security Introduction and Requirements", RFC 4033, DOI 10.17487/RFC4033, March 2005, <<https://www.rfc-editor.org/info/rfc4033>>.

[RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, DOI 10.17487/RFC4291, February 2006, <<https://www.rfc-editor.org/info/rfc4291>>.

[RFC4632] Fuller, V. and T. Li, "Classless Inter-domain Routing (CIDR): The Internet Address Assignment and Aggregation Plan", BCP 122, RFC 4632, DOI 10.17487/RFC4632, August 2006, <<https://www.rfc-editor.org/info/rfc4632>>.

[RFC5389] Rosenberg, J., Mahy, R., Matthews, P., and D. Wing, "Session Traversal Utilities for NAT (STUN)", RFC 5389, DOI 10.17487/RFC5389, October 2008, <<https://www.rfc-editor.org/info/rfc5389>>.

[RFC5693] Seedorf, J. and E. Burger, "Application-Layer Traffic Optimization (ALTO) Problem Statement", RFC 5693, DOI 10.17487/RFC5693, October 2009, <<https://www.rfc-editor.org/info/rfc5693>>.

[RFC6708] Kiesel, S., Ed., Previdi, S., Stiernerling, M., Woundy, R., and Y. Yang, "Application-Layer Traffic Optimization (ALTO) Requirements", RFC 6708, DOI 10.17487/RFC6708, September 2012, <<https://www.rfc-editor.org/info/rfc6708>>.

[RFC7216] Thomson, M. and R. Bellis, "Location Information Server (LIS) Discovery Using IP Addresses and Reverse DNS", RFC 7216, DOI 10.17487/RFC7216, April 2014, <<https://www.rfc-editor.org/info/rfc7216>>.

- [RFC7285] Alimi, R., Ed., Penno, R., Ed., Yang, Y., Ed., Kiesel, S., Previdi, S., Roome, W., Shalunov, S., and R. Woundy, "Application-Layer Traffic Optimization (ALTO) Protocol", RFC 7285, DOI 10.17487/RFC7285, September 2014, <<https://www.rfc-editor.org/info/rfc7285>>.
- [RFC7286] Kiesel, S., Stiemerling, M., Schwan, N., Scharf, M., and H. Song, "Application-Layer Traffic Optimization (ALTO) Server Discovery", RFC 7286, DOI 10.17487/RFC7286, November 2014, <<https://www.rfc-editor.org/info/rfc7286>>.
- [RFC7624] Barnes, R., Schneier, B., Jennings, C., Hardie, T., Trammell, B., Huitema, C., and D. Borkmann, "Confidentiality in the Face of Pervasive Surveillance: A Threat Model and Problem Statement", RFC 7624, DOI 10.17487/RFC7624, August 2015, <<https://www.rfc-editor.org/info/rfc7624>>.
- [RFC7626] Bortzmeyer, S., "DNS Privacy Considerations", RFC 7626, DOI 10.17487/RFC7626, August 2015, <<https://www.rfc-editor.org/info/rfc7626>>.
- [RFC7858] Hu, Z., Zhu, L., Heidemann, J., Mankin, A., Wessels, D., and P. Hoffman, "Specification for DNS over Transport Layer Security (TLS)", RFC 7858, DOI 10.17487/RFC7858, May 2016, <<https://www.rfc-editor.org/info/rfc7858>>.
- [RFC7971] Stiemerling, M., Kiesel, S., Scharf, M., Seidel, H., and S. Previdi, "Application-Layer Traffic Optimization (ALTO) Deployment Considerations", RFC 7971, DOI 10.17487/RFC7971, October 2016, <<https://www.rfc-editor.org/info/rfc7971>>.
- [RFC8484] Hoffman, P. and P. McManus, "DNS Queries over HTTPS (DoH)", RFC 8484, DOI 10.17487/RFC8484, October 2018, <<https://www.rfc-editor.org/info/rfc8484>>.

Appendix A. Solution Approaches for Partitioned ALTO Knowledge

The ALTO base protocol document [RFC7285] specifies the communication between an ALTO client and a single ALTO server. It is implicitly assumed that this server can answer any query, possibly with some kind of default value if no exact data is known. No special provisions were made for the case that the ALTO information originates from multiple sources, which are possibly under the control of different administrative entities (e.g., different ISPs) or that the overall ALTO information is partitioned and stored on several ALTO servers.

A.1. Classification of Solution Approaches

Various protocol extensions and other solutions have been proposed to deal with multiple information sources and partitioned knowledge. They can be classified as follows:

1. Ensure that all ALTO servers have the same knowledge.
 - 1.1 Ensure data replication and synchronization within the provisioning protocol (cf. [RFC5693], Figure 1).
 - 1.2 Use an inter-ALTO-server data replication protocol. Possibly, the ALTO protocol itself -- maybe with some extensions -- could be used for that purpose; however, this has not been studied in detail so far.
2. Accept that different ALTO servers (possibly operated by different organizations, e.g., ISPs) do not have the same knowledge.
 - 2.1 Allow ALTO clients to send arbitrary queries to any ALTO

server (e.g., the one discovered using [RFC7286]). If this server cannot answer the query itself, it will fetch the data on behalf of the client, using the ALTO protocol or a to-be-defined inter-ALTO-server request forwarding protocol.

- 2.2 Allow ALTO clients to send arbitrary queries to any ALTO server (e.g., the one discovered using [RFC7286]). If this server cannot answer the query itself, it will redirect the client to the "right" ALTO server that has the desired information, using a small to-be-defined extension of the ALTO protocol.
- 2.3 ALTO clients need to use some kind of "search engine" that indexes ALTO servers and redirects and/or gives cached results.
- 2.4 ALTO clients need to use a new discovery mechanism to discover the ALTO server that has the desired information and contact it directly.

A.2. Discussion of Solution Approaches

The provisioning or initialization protocol for ALTO servers (cf. [RFC5693], Figure 1) is currently not standardized. It was a conscious decision not to include this in the scope of the IETF ALTO working group. The reason is that there are many different kinds of information sources. This implementation-specific protocol will adapt them to the ALTO server, which offers a standardized protocol to the ALTO clients. However, adding the task of synchronization between ALTO servers to this protocol (i.e., Approach 1.1) would overload this protocol with a second functionality that requires standardization for seamless multidomain operation.

For Approaches 1.1 and 1.2, in addition to general technical feasibility and issues like overhead and caching efficiency, another aspect to consider is legal liability. Operator "A" might prefer not to publish information about nodes in, or paths between, the networks of operators "B" and "C" through A's ALTO server, even if A knew that information. This is not only a question of map size and processing load on A's ALTO server. Operator A could also face legal liability issues if that information had a bad impact on the traffic engineering between B's and C's networks or on their business models.

No specific actions to build a solution based on a "search engine" (Approach 2.3) are currently known, and it is unclear what could be the incentives to operate such an engine. Therefore, this approach is not considered in the remainder of this document.

A.3. The Need for Cross-Domain ALTO Server Discovery

Approaches 1.1, 1.2, 2.1, and 2.2 require more than just the specification of an ALTO protocol extension or a new protocol that runs between ALTO servers. A large-scale, maybe Internet-wide, multidomain deployment would also need mechanisms by which an ALTO server could discover other ALTO servers, learn which information is available where, and ideally also who is authorized to publish information related to a given part of the network. Approach 2.4 needs the same mechanisms, except that they are used on the client side instead of the server side.

It is sometimes questioned whether there is a need for a solution that allows clients to ask arbitrary queries, even if the ALTO information is partitioned and stored on many ALTO servers. The main argument is that clients are supposed to optimize the traffic from and to themselves, and that the information needed for that is most likely stored on a "nearby" ALTO server -- i.e., the one that can be discovered using [RFC7286]. However, there are scenarios where the ALTO client is not co-located with an endpoint of the to-be-optimized data transmission. Instead, the ALTO client is located at a third party that takes part in the application signaling -- e.g., a so-called "tracker" in a peer-to-peer application. One such scenario,

where it is advantageous to place the ALTO client not at an endpoint of the user data transmission, is analyzed in Appendix C.

A.4. Our Solution Approach

Several solution approaches for cross-domain ALTO server discovery have been evaluated, using the criteria documented in Appendix B. One of them was to use the ALTO protocol itself for the exchange of information availability [ALTO4ALTO]. However, the drawback of that approach is that a new registration administration authority would have to be established.

This document specifies a DNS-based procedure for cross-domain ALTO server discovery, which was inspired by "Location Information Server (LIS) Discovery Using IP Addresses and Reverse DNS" [RFC7216]. The primary goal is that this procedure can be used on the client side (i.e., Approach 2.4), but together with new protocols or protocol extensions, it could also be used to implement the other solution approaches itemized above.

A.5. Relation to the ALTO Requirements

During the design phase of the overall ALTO solution, two different server discovery scenarios were identified and documented in the ALTO requirements document [RFC6708]. The first scenario, documented in Req. AR-32, can be supported using the discovery mechanisms specified in [RFC7286]. An alternative approach, based on IP anycast [ALTO-ANYCAST], has also been studied. This document, in contrast, tries to address Req. AR-33.

Appendix B. Requirements for Cross-Domain Server Discovery

This appendix itemizes requirements that were collected before the design phase and are reflected in the design of the ALTO Cross-Domain Server Discovery Procedure.

B.1. Discovery Client Application Programming Interface

The discovery client will be called through some kind of application programming interface (API), and the parameters will be an IP address and, for purposes of extensibility, a service identifier such as "ALTO". The client will return one or more URIs that offer the requested service ("ALTO") for the given IP address.

In other words, the client would be used to retrieve a mapping:

(IP address, "ALTO") -> IRD-URI(s)

where IRD-URI(s) is one or more URIs of Information Resource Directories (IRDs, see Section 9 of [RFC7285]) of ALTO servers that can give reasonable guidance to a resource consumer with the indicated IP address.

B.2. Data Storage and Authority Requirements

The information for mapping IP addresses and service parameters to URIs should be stored in a -- preferably distributed -- database. It must be possible to delegate administration of parts of this database. Usually, the mapping from a specific IP address to a URI is defined by the authority that has administrative control over this IP address -- e.g., the ISP in residential access networks or the IT department in enterprise, university, or similar networks.

B.3. Cross-Domain Operations Requirements

The cross-domain server discovery mechanism should be designed in such a way that it works across the public Internet and also in other IP-based networks. This, in turn, means that such mechanisms cannot rely on protocols that are not widely deployed across the Internet or protocols that require special handling within participating networks. An example is multicast, which is not generally available

across the Internet.

The ALTO Cross-Domain Server Discovery Protocol must support gradual deployment without a network-wide flag day. If the mechanism needs some kind of well-known "rendezvous point", reusing an existing infrastructure (such as the DNS root servers or the WHOIS database) should be preferred over establishing a new one.

B.4. Protocol Requirements

The protocol must be able to operate across middleboxes, especially NATs and firewalls.

The protocol shall not require any preknowledge from the client other than any information that is known to a regular IP host on the Internet.

B.5. Further Requirements

The ALTO cross-domain server discovery cannot assume that the server-discovery client and the server-discovery responding entity are under the same administrative control.

Appendix C. ALTO and Tracker-Based Peer-to-Peer Applications

This appendix provides a complete example of using ALTO and the ALTO Cross-Domain Server Discovery Procedure in one specific application scenario -- namely, a tracker-based peer-to-peer application. First, in Appendix C.1, we introduce a generic model of such an application and show why ALTO optimization is desirable. Then, in Appendix C.2, we introduce two architectural options for integrating ALTO into the tracker-based peer-to-peer application; one option is based on the "regular" ALTO server discovery procedure [RFC7286], and one relies on the ALTO Cross-Domain Server Discovery Procedure. In Appendix C.3, a simple mathematical model is used to show that the latter approach is expected to yield significantly better optimization results. The appendix concludes with Appendix C.4, which details an exemplary complete walk-through of the ALTO Cross-Domain Server Discovery Procedure.

C.1. A Generic Tracker-Based Peer-to-Peer Application

The optimization of peer-to-peer (P2P) applications such as BitTorrent was one of the first use cases that lead to the inception of the IETF ALTO working group. Further use cases have been identified as well, yet we will use this scenario to illustrate the operation and usefulness of the ALTO Cross-Domain Server Discovery Procedure.

For the remainder of this chapter, we consider a generic, tracker-based peer-to-peer file-sharing application. The goal is the dissemination of a large file, without using one large server with a correspondingly high upload bandwidth. The file is split into chunks. So-called "peers" assume the role of both a client and a server. That is, they may request chunks from other peers, and they may serve the chunks they already possess to other peers at the same time, thereby contributing their upload bandwidth. Peers that want to share the same file participate in a "swarm". They use the peer-to-peer protocol to inform each other about the availability of chunks and request and transfer chunks from one peer to another. A swarm may consist of a very large number of peers. Consequently, peers usually maintain logical connections to only a subset of all peers in the swarm. If a new peer wants to join a swarm, it first contacts a well-known server, the "tracker", which provides a list of IP addresses of peers in the swarm.

A swarm is an overlay network on top of the IP network. Algorithms that determine the overlay topology and the traffic distribution in the overlay may consider information about the underlying IP network, such as topological distance, link bandwidth, (monetary) costs for sending traffic from one host to another, etc. ALTO is a protocol

for retrieving such information. The goal of such "topology-aware" decisions is to improve performance or Quality of Experience in the application while reducing the utilization of the underlying network infrastructure.

C.2. Architectural Options for Placing the ALTO Client

The ALTO protocol specification [RFC7285] details how an ALTO client can query an ALTO server for guiding information and receive the corresponding replies. However, in the considered scenario of a tracker-based P2P application, there are two fundamentally different possible locations for where to place the ALTO client:

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

In the following, both scenarios are compared in order to explain the need for ALTO queries on behalf of remote resource consumers.

In the first scenario (see Figure 2), the resource consumer queries the 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.

In the second scenario (see Figure 4), the resource directory has an embedded ALTO client. After receiving a query for a given resource (F1), the resource directory invokes this ALTO client to evaluate all resource providers it knows (F2/F3). Then it returns a list, possibly shortened, containing the "best" resource providers to the resource consumer (F4).

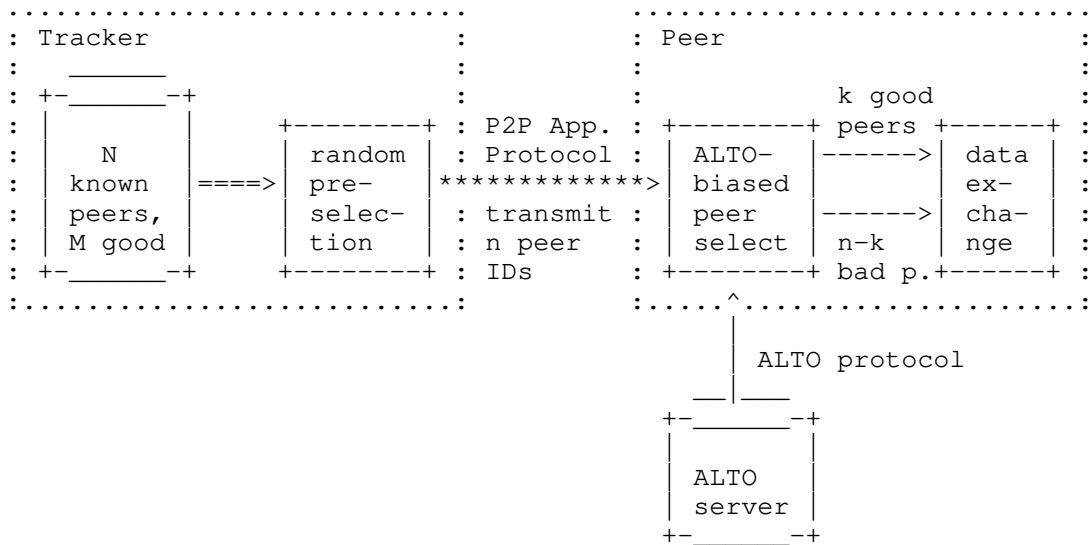
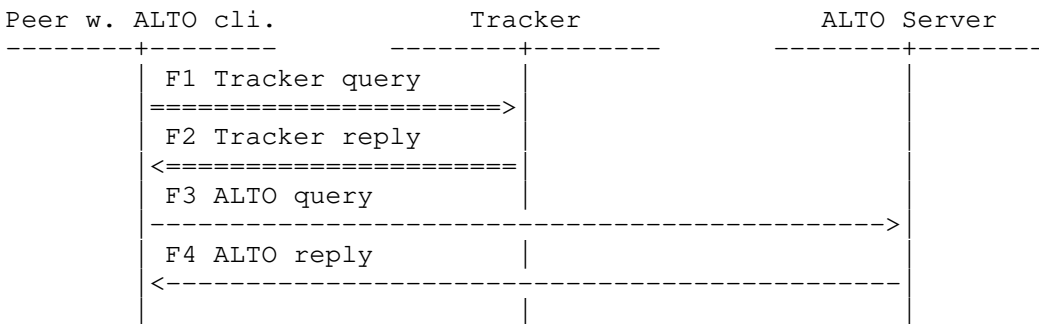


Figure 1: Tracker-Based P2P Application with Random Peer Preselection



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

Figure 2: Basic Message Sequence Chart for Resource Consumer-Initiated ALTO Query

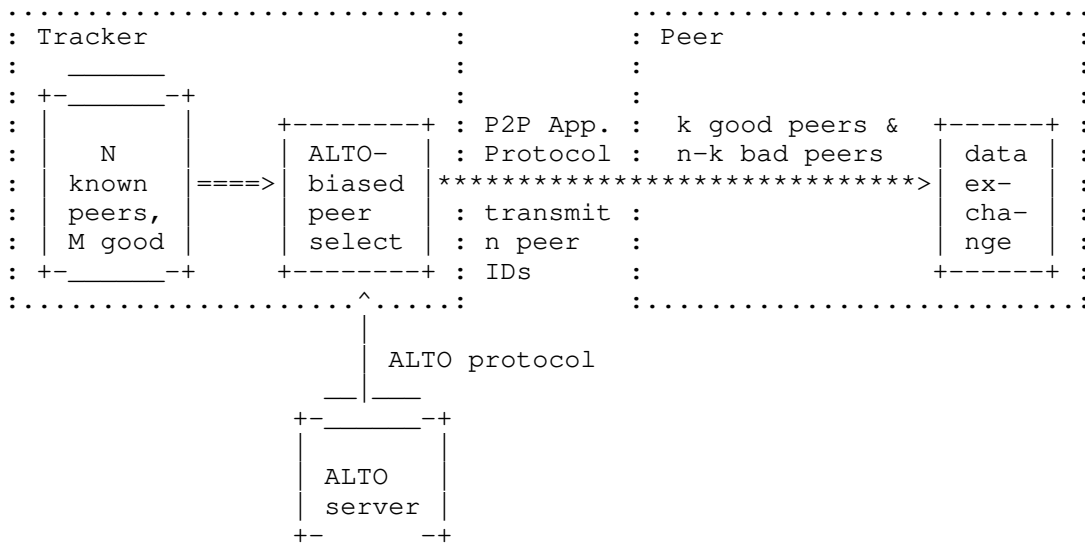
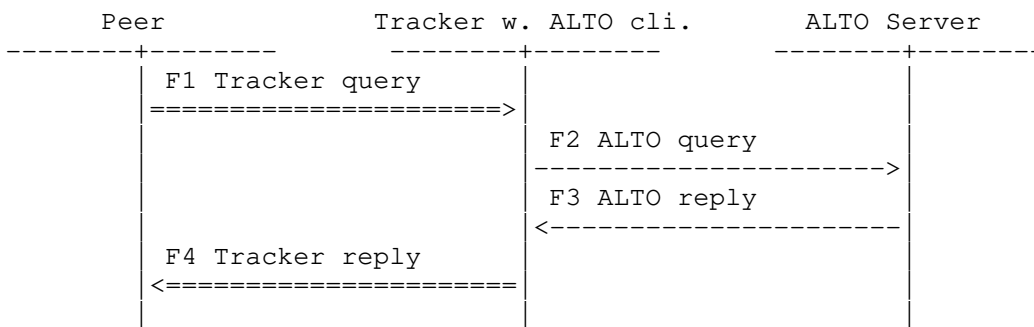


Figure 3: Tracker-Based P2P Application with ALTO Client in Tracker



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

Figure 4: Basic Message Sequence Chart for ALTO Query on Behalf of Remote Resource Consumer

Note: The message sequences depicted in Figures 2 and 4 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.

C.3. Evaluation

The problem with the first approach is that 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.

For illustration, consider a simple model of a swarm, in which all peers fall into one of only two categories: assume that there are only "good" (in the sense of ALTO's better-than-random peer selection, based on an arbitrary desired rating criterion) and "bad" peers. Having more different categories makes the math more complex but does not change anything about the basic outcome of this analysis. Assume that the swarm has a total number of N peers, out of which there are M "good" and N-M "bad" peers, which are all known to the tracker. A new peer wants to join the swarm and therefore asks the tracker for a list of peers.

If, according to the first approach, the tracker randomly picks n peers from the N known peers, the result can be described with the hypergeometric distribution. The probability that the tracker reply contains exactly k "good" peers (and $n-k$ "bad" peers) is:

$$P(X=k) = \frac{\binom{M}{k} \binom{N-M}{n-k}}{\binom{N}{n}}$$

with $\binom{n}{k} = \frac{n!}{k!(n-k)!}$ and $n! = n * (n-1) * (n-2) * .. * 1$

The probability that the reply contains at most k "good" peers is:
 $P(X \leq k) = P(X=0) + P(X=1) + .. + P(X=k)$.

For example, consider a swarm with $N=10,000$ peers known to the tracker, out of which $M=100$ are "good" peers. If the tracker randomly selects $n=100$ peers, the formula yields for the reply: $P(X=0)=36\%$, $P(X \leq 4)=99\%$. That is, with a probability of approximately 36%, this list does not contain a single "good" peer, and with 99% probability, there are only four or fewer of the "good" peers on the list. Processing this list with the guiding ALTO information will ensure that the few favorable peers are ranked to the top of the list; however, the benefit is rather limited as the number of favorable peers in the list is just too small.


Much better traffic optimization could be achieved if the tracker would evaluate all known peers using ALTO and return a list of 100 peers afterwards. This list would then include a significantly higher fraction of "good" peers. (Note that 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 ALTO queries on behalf of remote resource consumers. That is, as the tracker issues ALTO queries on behalf of the peer that contacted the tracker, the tracker must be able to discover an ALTO server that can give guidance suitable for that peer. This task can be solved using the ALTO Cross-Domain Server Discovery Procedure.

C.4. Example

This section provides a complete example of the ALTO Cross-Domain Server Discovery Procedure in a tracker-based peer-to-peer scenario.

The example is based on the network topology shown in Figure 5. Five access networks -- Networks a, b, c, x, and t -- are operated by five different network operators. They are interconnected by a backbone structure. Each network operator runs an ALTO server in their network -- i.e., ALTO_SRV_A, ALTO_SRV_B, ALTO_SRV_C, ALTO_SRV_X, and ALTO_SRV_T, respectively.






Figure 5: Example Network Topology

A new peer of a peer-to-peer application wants to join a specific swarm (overlay network), in order to access a specific resource. This new peer will be called "Resource Consumer X", in accordance with the terminology of [RFC6708], and is located in Network x. It contacts the tracker ("Resource Directory"), which is located in Network t. The mechanism by which the new peer discovers the tracker is out of the scope of this document. The tracker maintains a list of peers that take part in the overlay network, and hence it can determine that Resource Providers A, B, and C are candidate peers for Resource Consumer X.

As shown in the previous section, a tracker-side ALTO optimization (cf. Figures 3 and 4) is more efficient than a client-side optimization. Consequently, the tracker wants to use the ALTO Endpoint Cost Service (ECS) to learn the routing costs between X and A, X and B, and X and C, in order to sort A, B, and C by their respective routing costs to X.

In theory, there are many options for how the ALTO Cross-Domain Server Discovery Procedure could be used. For example, the tracker could do the following steps:

```

IRD_URIS_A = XDOMDISC(A,"ALTO:https")
COST_X_A   = query the ECS(X,A,routingcost) found in IRD_URIS_A

IRD_URIS_B = XDOMDISC(B,"ALTO:https")
COST_X_B   = query the ECS(X,B,routingcost) found in IRD_URIS_B

IRD_URIS_C = XDOMDISC(C,"ALTO:https")
COST_X_C   = query the ECS(X,C,routingcost) found in IRD_URIS_C
  
```

In this scenario, the ALTO Cross-Domain Server Discovery Procedure queries might yield: IRD_URIS_A = ALTO_SRV_A, IRD_URIS_B = ALTO_SRV_B, and IRD_URIS_C = ALTO_SRV_C. That is, each ECS query would be sent to a different ALTO server. The problem with this approach is that we are not necessarily able to compare COST_X_A, COST_X_B, and COST_X_C with each other. The specification of the routingcost metric mandates that "A lower value indicates a higher preference", but "an ISP may internally compute routing cost using any method that it chooses" (see Section 6.1.1.1 of [RFC7285]). Thus, COST_X_A could be 10 (milliseconds round-trip time), while COST_X_B could be 200 (kilometers great circle distance between the approximate geographic locations of the hosts) and COST_X_C could be 3 (router hops, corresponding to a decrease of the TTL field in the IP header). Each of these metrics fulfills the "lower value is more preferable" requirement on its own, but they obviously cannot be compared with each other. Even if there were a reasonable formula to compare, for example, kilometers with milliseconds, we could not use it, as the units of measurement (or any other information about the computation method for the routingcost) are not sent along with the value in the ECS reply.

To avoid this problem, the tracker tries to send all ECS queries to the same ALTO server. As specified in Section 4.4 of this document, Case 2, it uses the IP address of Resource Consumer x as a parameter of the discovery procedure:

```

IRD_URIS_X = XDOMDISC(X,"ALTO:https")
COST_X_A   = query the ECS(X,A,routingcost) found in IRD_URIS_X
  
```

COST_X_B = query the ECS(X,B,routingcost) found in IRD_URIS_X
COST_X_C = query the ECS(X,C,routingcost) found in IRD_URIS_X

This strategy ensures that COST_X_A, COST_X_B, and COST_X_C can be compared with each other.

As discussed above, the tracker calls the ALTO Cross-Domain Server Discovery Procedure with IP address X as a parameter. For the remainder of this example, we assume that X = 2001:DB8:1:2:227:eff:fe6a:de42. Thus, the procedure call is IRD_URIS_X = XDOMDISC(2001:DB8:1:2:227:eff:fe6a:de42,"ALTO:https").

The first parameter, 2001:DB8:1:2:227:eff:fe6a:de42, is a single IPv6 address. Thus, we get AT = IPv6, A = 2001:DB8:1:2:227:eff:fe6a:de42, L = 128, and SP = "ALTO:https".

The procedure constructs (see Step 1 in Section 3.2)

```
R128 = "2.4.E.D.A.6.E.F.F.F.E.0.7.2.2.0.2.0.0.0.1.0.0.0.  
      8.B.D.0.1.0.0.2.IP6.ARPA."
```

as well as the following (see Step 2 in Section 3.2):

```
R64 = "2.0.0.0.1.0.0.0.8.B.D.0.1.0.0.2.IP6.ARPA."  
R56 = "0.0.1.0.0.0.8.B.D.0.1.0.0.2.IP6.ARPA."  
R48 = "1.0.0.0.8.B.D.0.1.0.0.2.IP6.ARPA."  
R40 = "0.0.8.B.D.0.1.0.0.2.IP6.ARPA."  
R32 = "8.B.D.0.1.0.0.2.IP6.ARPA."
```

In order to illustrate the third step of the ALTO Cross-Domain Server Discovery Procedure, we use the "dig" (domain information groper) DNS lookup utility that is available for many operating systems (e.g., Linux). A real implementation of the ALTO Cross-Domain Server Discovery Procedure would not be based on the "dig" utility but instead would use appropriate libraries and/or operating-system APIs. Please note that the following steps have been performed in a controlled lab environment with an appropriately configured name server. A suitable DNS configuration will be needed to reproduce these results. Please also note that the rather verbose output of the "dig" tool has been shortened to the relevant lines.

Since AT = IPv6 and L = 128, in the table given in Section 3.4, the sixth row (not counting the column headers) applies.

As mandated by the third column, we start with a lookup of R128, looking for NAPTR resource records:

```
| user@labpc:~$ dig -tNAPTR 2.4.E.D.A.6.E.F.F.F.E.0.7.2.2.0.\  
| 2.0.0.0.1.0.0.0.8.B.D.0.1.0.0.2.IP6.ARPA.  
|  
| ;; Got answer:  
| ;; ->>HEADER<<- opcode: QUERY, status: NXDOMAIN, id: 26553  
| ;; flags: qr aa rd ra; QUERY: 1, ANSWER: 0, AUTHORITY: 1, ADD'L: 0
```

The domain name R128 does not exist (status: NXDOMAIN), so we cannot get a useful result. Therefore, we continue with the fourth column of the table and do a lookup of R64:

```
| user@labpc:~$ dig -tNAPTR 2.0.0.0.1.0.0.0.8.B.D.0.1.0.0.2.IP6.ARPA.  
|  
| ;; Got answer:  
| ;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 33193  
| ;; flags: qr aa rd ra; QUERY: 1, ANSWER: 0, AUTHORITY: 1, ADD'L: 0
```

The domain name R64 could be looked up (status: NOERROR), but there are no NAPTR resource records associated with it (ANSWER: 0). There may be some other resource records such as PTR, NS, or SOA, but we are not interested in them. Thus, we do not get a useful result, and we continue with looking up R56:


```

user@labpc:~$ dig -tNAPTR 0.0.1.0.0.0.8.B.D.0.1.0.0.2.IP6.ARPA.
;; Got answer:
;; ->HEADER<<- opcode: QUERY, status: NOERROR, id: 35966
;; flags: qr aa rd ra; QUERY: 1, ANSWER: 2, AUTHORITY: 1, ADD'L: 2
;; ANSWER SECTION:
0.0.1.0.0.0.8.B.D.0.1.0.0.2.IP6.ARPA. 604800 IN NAPTR 100 10 "u"
  "LIS:HELD" "!.*!https://lis1.example.org:4802/?c=ex!" .
0.0.1.0.0.0.8.B.D.0.1.0.0.2.IP6.ARPA. 604800 IN NAPTR 100 20 "u"
  "LIS:HELD" "!.*!https://lis2.example.org:4802/?c=ex!" .

```

The domain name R56 could be looked up, and there are NAPTR resource records associated with it. However, each of these records has a service parameter that does not match our SP = "ALTO:https" (see [RFC7216] for "LIS:HELD"), and therefore we have to ignore them. Consequently, we still do not have a useful result and continue with a lookup of R48:

```

user@labpc:~$ dig -tNAPTR 1.0.0.0.8.B.D.0.1.0.0.2.IP6.ARPA.
;; Got answer:
;; ->HEADER<<- opcode: QUERY, status: NOERROR, id: 50459
;; flags: qr aa rd ra; QUERY: 1, ANSWER: 2, AUTHORITY: 1, ADD'L: 2
;; ANSWER SECTION:
1.0.0.0.8.B.D.0.1.0.0.2.IP6.ARPA. 604800 IN NAPTR 100 10 "u"
  "ALTO:https" "!.*!https://alto1.example.net/ird!" .
1.0.0.0.8.B.D.0.1.0.0.2.IP6.ARPA. 604800 IN NAPTR 100 10 "u"
  "LIS:HELD" "!.*!https://lis.example.net:4802/?c=ex!" .

```

This lookup yields two NAPTR resource records. We have to ignore the second one as its service parameter does not match our SP, but the first NAPTR resource record has a matching service parameter. Therefore, the procedure terminates successfully and the final outcome is: IRD_URIS_X = "https://alto1.example.net/ird".

The ALTO client that is embedded in the tracker will access the ALTO Information Resource Directory (IRD, see Section 9 of [RFC7285]) at this URI, look for the Endpoint Cost Service (ECS, see Section 11.5 of [RFC7285]), and query the ECS for the costs between A and X, B and X, and C and X, before returning an ALTO-optimized list of candidate resource providers to resource consumer X.

Acknowledgments

The initial draft version of this document was co-authored by Marco Tomsu (Alcatel-Lucent).

This document borrows some text from [RFC7286], as historically, it was part of the draft that eventually became said RFC. Special thanks to Michael Scharf and Nico Schwan.

Authors' Addresses

Sebastian Kiesel
 University of Stuttgart Information Center
 Allmandring 30
 70550 Stuttgart
 Germany

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

Martin Stiernerling
 University of Applied Sciences Darmstadt, Computer Science Dept.
 Haardtring 100
 64295 Darmstadt
 Germany

Phone: +49 6151 16 37938
Email: mls.ietf@gmail.com
URI: <https://danet.fbi.h-da.de>