

ALTO WG
Internet-Draft
Intended status: Standards Track
Expires: 5 May 2021

Q. Xiang
Y. Yang
Yale University
1 November 2020

ALTO Multi-Domain Extension: Challenges, Existing Efforts and Next Steps
[draft-xiang-alto-multidomain-extension-00.txt](#)

Abstract

The emerging multi-domain applications, such as flexible interdomain routing, distributed, federated machine learning and multi-domain collaborative dataset transfer, can benefit substantially from getting information from networks. The ALTO base protocol [[RFC7285](#)] provides a northbound interface for applications to retrieve the network information. In particular, it specifies the communication between an ALTO client and an ALTO server, where the ALTO is implicitly assumed to be able to answer any query from the ALTO client. However, it does not specify the cases when the network information are originated from multiple domains (i.e., administrative entities or geographically partitions). This document summarizes the discussion on the ALTO weekly meeting since IETF 108 on how to extend ALTO to support multi-domain applications. It identifies the key challenges for retrieving network information from multiple networks, reviews the existing efforts in the work group, and discusses the next steps to fully address the challenges.

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 <https://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 5 May 2021.

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	2
2.	Requirements Language	2
3.	Extending ALTO to Multi-Domain: Challenges	3
4.	Existing Efforts in the ALTO Working Group	4
5.	ALTO Multi-Domain Extension: Basic Ideas	4
6.	References	5
6.1.	Normative References	5
6.2.	Informative References	5
	Authors' Addresses	6

[1.](#) Introduction

The ALTO protocol [[RFC7285](#)] provides network information to applications so that applications can make network informed decisions to improve the performance. Not only traditional applications such peer-to-peer systems, many recent, novel multi-domain applications, which orchestrate resources across multiple networks, can also benefit substantially from ALTO.

Since the recharter discussion on IETF 108, there has been much interest and discussion on the ALTO weekly meetings to extend the ALTO base protocol to support multi-domain applications. This document is a summary of the discussion on these meetings. It outlines the challenges, reviews the existing efforts in the working group, and discuss the next steps to design ALTO multi-domain extensions.

[2.](#) Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [[RFC2119](#)].

3. Extending ALTO to Multi-Domain: Challenges

This section uses a motivating example to outline the key challenges for designing ALTO extensions to support multi-domain applications.

Consider the example in Figure 1, where the application wants to retrieve the network information from endpoint S to endpoint D. S and D are each attached to a different network.

- * Challenge 1: the route from S to D spans multiple networks, whose information is partitioned and stored at multiple ALTO servers. In other words, no single ALTO server has the complete information of the flow from S to D. As such, either the ALTO client needs to communicate with all ALTO servers along the route to collect and aggregate all the information, or the ALTO server that receives the ALTO client's query at the beginning needs to act as a delegate to collect and aggregate all the information before returning it to the ALTO client.
- * Challenge 2: the route from S to D may not even be fully instantiated. For example, when a network uses SDN to manage its routing and performs load balancing based on the the remainder of flows' source MAC addresses divided by 3, it is impossible to fully instantiate the forwarding table on the limited memory in a commodity OpenFlow switch. As such, the ALTO client or the delegated ALTO server needs to decide which other ALTO servers to contact to collect the needed information.
- * Challenge 3: Different ALTO servers may have information about different metrics. In particular, along the route from S to D, two ALTO servers may provide information two segments of the route with different metrics, e.g., one provides the bandwidth information while the other provides the latency information. Even if two ALTO servers provide information of the same metric, say latency, they may provide different statistics, e.g., one provide the mean latency while the other provide the median latency. As such, the ALTO client or the delegated ALTO server needs to figure out how to aggregate such information and return to the application.
- * Challenge 4: When the application wants to retrieve network information of multiple pairs of source-destination endpoints, the routes of different flows may share same link(s), leading to resource sharing. As such, collecting network information of different flows individually may yield inaccurate results.

4. Existing Efforts in the ALTO Working Group

Several documents in the ALTO group have made design proposals for extending ALTO to multi-domain settings. This section provides a review.

- * [[DRAFT-UNICORN-INFO](#)] and [[DRAFT-NFCHAIN](#)] propose broker-based architectures for collection network information from multiple networks. In particular, an ALTO client act as a broker or aggregator on behalf of applications to directly communicate with all ALTO servers in the networks to collect and aggregate information. These efforts provide initial starting point for addressing challenge 1, but does not specify how different ALTO servers can communicate with each other.
- * For challenge 2, [[RFC8686](#)] specifies a cross-domain ALTO server discovery mechanism for the ALTO client or the delegate ALTO server to discover which ALTO server possesses the information of interest. It partially address challenge 2. However, it does not specify how to handle the scenario where different ALTO servers possess different parts of the information of interest (i.e., partial information).
- * [[DRAFT-UNICORN-INFO](#)] provides an information aggregation mechanism to aggregate information from multiple networks. It is related to challenge 3, however, it assumes that all networks provide information using the same metric and the same statistics.
- * [[DRAFT-PV](#)] designs the ALTO path vector extension to support ALTO servers to capture the resource sharing among multiple flows and return to the ALTO client. However, it does not specify the scenario where network uses more complex routing strategies, such as multipath routing and on-demand routing, for source-destination endpoints.

5. ALTO Multi-Domain Extension: Basic Ideas

This section describes the basic ideas to design an ALTO multi-domain extension to address the aforementioned challenges. In particular, to address challenge 1 and 2 to identify the full route of a source-destination endpoint pair and query the corresponding ALTO servers to retrieve information, we will base on the pub-sub design of the SDN federation routing protocol [[SFP](#)] and go beyond to specify the inter-ALTO-server communication mechanism, which allows ALTO servers to exchange and aggregate information from multiple networks.

To address challenge 3 to aggregate the information from different networks, we will extend the ALTO framework to standardize single aggregation in some settings, where information are in the same metric and uses the same statistics. To aggregate the information from different networks in different metrics and statics, we will resort to the theoretical framework of routing algebra [[ROUTING-MO](#)], which aggregates routing information of different metrics as a partial order of vectors, and go beyond to specify the multi-domain ALTO server information aggregation mechanism, which supports the more flexible aggregation of ALTO information.

To address challenge 4 to retrieve the resource sharing information of multiple source-destination endpoints pairs when networks use more complex routing strategies such as multipath routing and on-demand routing, we will base on the ALTO path vector extension [[DRAFT-PV](#)], but go beyond to use generic mathematical programming constraints as a compact representation of the resource sharing information across multiple networks.

[6.](#) References

[6.1.](#) Normative References

[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>>.

[6.2.](#) Informative References

[DRAFT-NFCHAIN]

Perez, D. and C. Rothenberg, "ALTO-based Broker-assisted Multi-domain Orchestration", 2018, <<https://datatracker.ietf.org/doc/html/draft-lachosrothenberg-alto-brokermdo-01>>.

[DRAFT-PV] Bernstein, G., Lee, Y., Roome, W., Scharf, M., and Y. R. Yang, "ALTO Extension: Abstract Path Vector as a Cost Mode", 2015, <<https://tools.ietf.org/html/draft-yang-alto-path-vector-01>>.

[DRAFT-UNICORN-INFO]

Xiang, Q., Newman, H., Bernstein, G., Du, H., Gao, K., Mughal, A., Balcas, J., Zhang, J., and Y. R. Yang, "Implementation and Deployment of A Resource Orchestration System for Multi-Domain Data Analytics", 2017, <<https://datatracker.ietf.org/doc/draft-xiang-alto-exascale-network-optimization/>>.

- [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>>.
- [RFC8189] Randriamasy, S., Roome, W., and N. Schwan, "Multi-Cost Application-Layer Traffic Optimization (ALTO)", [RFC 8189](#), DOI 10.17487/RFC8189, October 2017, <<https://www.rfc-editor.org/info/rfc8189>>.
- [RFC8686] Kiesel, S. and M. Stiernerling, "Application-Layer Traffic Optimization (ALTO) Cross-Domain Server Discovery", [RFC 8686](#), DOI 10.17487/RFC8686, February 2020, <<https://www.rfc-editor.org/info/rfc8686>>.
- [ROUTING-MO] Sobrinho, J. and M. Ferreira, "Routing on Multiple Optimality Criteria", 2020, <<https://conferences.sigcomm.org/sigcomm/2020>>.
- [SFP] Xiang, Q., Guok, C., Le, F., MacAuley, J., Newman, H., and Y. R. Yang, "SFP: Toward Interdomain Routing for SDN Networks", 2018, <<https://conferences.sigcomm.org/sigcomm/2018>>.

Authors' Addresses

Qiao Xiang
Yale University
51 Prospect Street
New Haven, CT
United States of America

Email: qiao.xiang@cs.yale.edu

Y. Richard Yang
Yale University
51 Prospect Street
New Haven, CT
United States of America

Email: yry@cs.yale.edu

