Multi-Domain Information Exposure using ALTO: The Good, the Bad and the Solution

Danny Lachos
University of Campinas

Christian Rothenberg
University of Campinas

Qiao Xiang
Yale University

Y. Richard Yang
PCL/Yale University

Börje Ohlman
Ericsson

Sabine Randriamasy
Nokia Bell Labs

Luis M. Contreras
Telefonica

Kai Gao
Sichuan University

ABSTRACT

Multi-domain applications can benefit from network information exposure using the ALTO protocol framework. ALTO provides network state and capabilities to applications to be more flexible in terms of rate adaptation, transmission time, and server/path selection. However, different key issues arise when designing an ALTO solution for multi-domain environments. In this talk, we summarize such issues along with basic mechanisms to be considered to allow ALTO to expose network information across multiple domains.

CCS CONCEPTS

• **Networks → Network protocols**;

KEYWORDS

ALTO, application-layer traffic optimization, multi-domain

1 INTRODUCTION

Emerging multi-domain applications (*e.g.*, collaborative data sciences [5, 19, 28], flexible inter-domain routing [12, 20, 31], E2E network services [2, 8, 14]) require resource orchestration across multiple networks managed by different administrative domains. Such cross-domain applications can benefit substantially from network information exposure to make application-layer resource optimization and improve their performance.

The Application-Layer Traffic Optimization (ALTO) protocol [1] provides a generic framework to expose network information for applications to take optimized actions. However, exposing network information to support multi-domain use cases introduces issues to be considered in the current ALTO design.

In this talk, we identify what network information the multi-domain applications need and the benefit of using it (the *Good*, Sec. 2). Next, we discuss the ALTO design issues for gathering such multi-domain information (the *Bad*, Sec. 3). We then list a set of mechanisms to design a multi-domain ALTO framework (the *Solution*, Sec. 4).

2 WHAT INFORMATION DO APPS NEED?

Many types of network information are needed by cross-domain applications to improve their performances, including network state (*e.g.*, loss, delay, ECN bit [24], INT [16]), performance metrics (*e.g.*, throughput, max reservable Bandwidth),

ANRW '20, July 27–30, 2020, Online (Meetecho), Spain

ACM ISBN 978-1-4503-8039-3/20/07…$15.00
https://doi.org/10.1145/3404868.3406667

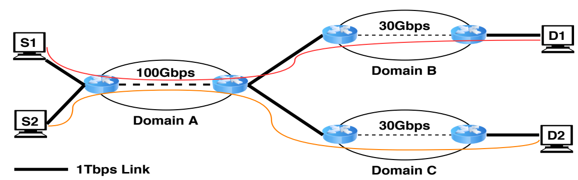

Figure 1: A collaboration network example.

capability information (*e.g.*, delivery/acquisition protocol), locality (*e.g.*, servers/domains location and paths), among others.

Motivating Use Case Example. Consider a collaborative network composed of three-member domains (See Fig. 1). An application (*e.g.*, a large data analysis system) wants to reserve bandwidth for two flows $f_1 : (S_1, D_1)$ and $f_2 : (S_2, D_2)$, across three domains A, B, and C. Before the application can run a resource allocation algorithm to execute such submitted flows, it needs to gather some information from the network. First, the *E2E cost* across multiple domains. The cost in terms of resource availability and sharing (*e.g.*, network bandwidth) for the set of requested flows to be reserved. In our presented scenario, both flows f_1 and f_2 are sharing the same network path in domain A. It means that they share a common resource, the network bandwidth. Second, each flow will consume networking resources of multiple domains. Therefore, the application needs to discover both a *sequence of domains* and *candidate paths* between source nodes and destination nodes, *i.e.*, which domains are involved for the different traffic flows and one or more paths connecting such domains. In our example, the multi-domain network paths for f_1 and f_2 are $[A, B]$, and $[A, C]$, respectively.

Basic Formulation. Consider different services, for each domain, providing previous information. Each service is defined as an object f_i with a set of network properties, such as $f_i.path$ (representing the sequence of network devices that packets of flow f_i will traverse), $f_i.abw$ (representing the available bandwidth that flow f_i can request), $f_i.delay$ (representing the average delay of packets of flow f_i), etc.

In our example, consider each domain providing the bandwidth property using a set of algebraic expressions (*i.e.*, linear inequalities) [13]: $\Pi_A : \{x_1 + x_2 \leq 100\}$, $\Pi_B : \{x_1 \leq 30\}$, and $\Pi_C : \{x_2 \leq 30\}$. Where x_1 and x_2 represent the available bandwidth that can be reserved for $(S1, D1)$, and $(S2, D2)$, respectively. Each linear inequality represents a constraint on the reservable bandwidths over different shared resources by the two flows. For example, the linear inequality Π_A indicates that both flows share a common resource and that the sum of their bandwidths can not exceed 100 Gbps.

In a multi-domain setting, a network property to a flow f_i may involve properties of multiple networks, *e.g.*:

fi.md-abw: $min(f_i.abw[A] + f_i.abw[B] + f_i.abw[C])$
fi.md-path: $f_i.path[A].f_i.path[B].f_i.path[C]$
fi.md-delay: $f_i.delay[A] + f_i.delay[B] + f_i.delay[C]$

Table 1: Issues of applying the current ALTO framework in the multi-domain setting & solutions.

Capability	Issues with the current mechanisms	Envisioned solutions & on-going efforts
Conceptual query interface and data representation		
Unified resource representation	In the current ALTO framework, each domain can have its own representation of the same network information. For example, domain A and B (See Fig. 1) may use utilization charge and available bandwidth as the path cost respectively. Even if all the domains have the same property, there may not necessarily be a uniform form, for example, domain A charges using dollars, while domains B and C take euros.	Multi-domain composition mechanisms are necessary, so that network information from ALTO servers in multiple domains can fit into a single and consistent "virtual" domain abstraction. In this talk, we present the design options of multi-domain composition mechanisms [23, 32, 34, 35].
Generic query language	Applications need to express their objectives and requirements in a query, *e.g.*, finding the bandwidth the network can provide for flow $f_1(S_1, D_1)$ subject to reachability requirements, traffic symmetry, waypoint traversal, QoS metrics, etc. The current query interfaces in ALTO (*e.g.*, filtered network/cost map) cannot express such flexible queries.	With a flexible/generic query language, the network can filter out a large number of unqualified domains. The language specification could be inspired by standard [3, 9] or pre-standard [4, 27] mechanisms, implemented with a user-friendly grammar (*e.g.*, SQL-style query).
Communication mechanisms		
ALTO Servers communication	In multi-domain scenarios, it is not possible to optimize the traffic with only locally available network information (*i.e.*, server-to-client ALTO communication). Therefore, communications among multiple ALTO servers are necessary to exchange detailed network information of multiple domains.	ALTO servers may consider either a hierarchical or mesh architectural deployment design [7, 22, 34]. In a hierarchical design, ALTO servers in domain partitions gather local information and send it to central server. In a mesh deployment, ALTO servers may be set up in each domain independently, and gathering the network information from other connected domains.
Multi-domain connectivity discovery	To find the resources shared by different source/destination pairs, an application needs to discover which domains are involved in the data movement of each node pair. Besides, a set of candidate paths need to be computed in order to know how to reach a remote destination node. The current ALTO extensions do not have this feature.	Multi-domain mechanisms combining domains sequence computation and paths computation need to be defined, or standardized computation protocols could be re-used such as BGP [25], PCE [17, 30], or BGP-LS [11].
Multi-domain ALTO Server discovery	Once the multi-domain connectivity discovery is performed, an application needs to be aware of the presence and the location of ALTO servers in order to get appropriate guidance. These ALTO servers will be located in different network domains, so that multi-domain ALTO server discovery mechanisms are needed.	The ALTO cross-domain server discovery document [15] specifies a procedure for identifying ALTO servers outside of the ALTO client's own network domain. Other mechanisms could also be leveraged, such as those based on PCE or BGP architectures [6, 26].
Computation model		
Computation complexity optimization	The optimization problems specified by the applications can be computationally expensive and time-consuming. For example, the number of available paths for each flow increases exponentially with the number of domains involved, as does the number of available configurations for a set of flows with both the network size and the number of flows.	ALTO servers need to support mechanisms (*e.g.*, pre-computation and projection) to improve the scalability and performance. Such mechanisms should effectively reduce the redundancy in the network view as much as possible while still providing the same information [10].
Security & privacy	The information provided by the ALTO protocol is considered coarse-grained in several multi-domain use cases. New ALTO extensions have been designed to provide fine-grained network information to the applications. Using these ALTO extension services for multi-domain scenarios would raise new security and privacy concerns.	ALTO needs mechanisms (with little overhead) that provide accurate sharing network information, and at the same time, protects each member domain. This privacy-preserving interdomain information process may consider, for instance, a secure multi-party computation (SMPC) protocol [33, 34].

The involved domains may exchange such multi-domain properties. They also may apply composition mechanisms to create a unified representation to reveal a compact multi-domain network resource information. For example, taking a look at the set of previous inequalities, one can conclude that the constraint π_B ($x_1 \leq 30$) and π_C ($x_2 \leq 30$) can eliminate that at domain A ($x_1 + x_2 \leq 100$). Domains may compose this information and remove the cross-domain redundancy [21, 29]. Therefore, the compressed multi-domain set of inequalities is reduced to two inequalities (*i.e.*, π_B and π_C).

3 WHAT ARE CURRENT ALTO ISSUES?

ALTO already introduces basic mechanisms (*e.g.*, modularity, dependency) and abstractions (*e.g.*, map services) for applications to improve their performance [18]. However, the current ALTO base protocol is not designed for a multi-domain setting of exposing network information. We list several key design issues of the current ALTO framework in the second column of Table 1, which can be roughly categorized in three aspects: (i) conceptual query interfaces and data representation, (ii) communication mechanisms, (iii) and computation model.

4 HOW TO DESIGN A FRAMEWORK?

Instead of building from scratch, we aim to design a multi-domain network information exposure framework on top of ALTO. The third column of Table 1 summarizes on-going efforts and potential solutions to address the aforementioned issues.

5 CONCLUSION

Many multi-domain applications are emerging with the development of new technologies (*e.g.*, 5G, SDN, NFV). ALTO emerges as a solution for exposing network information across multiple domains. In this talk, we present key issues as well as solution mechanisms in the current ALTO framework to support important multi-domain environments.

REFERENCES

[1] R. Alimi, R. Penno, Y. Yang, S. Kiesel, S. Previdi, W. Roome, S. Shalunov, and R. Woundy. 2014. *Application-Layer Traffic Optimization (ALTO) Protocol.* RFC 7285. RFC Editor. http://www.rfc-editor.org/rfc/rfc7285.txt http://www.rfc-editor.org/rfc/rfc7285.txt.

[2] NGMN Alliance. 2015. 5G white paper. *Next generation mobile networks, white paper* (2015), 1–125.

[3] GSM Association. 2019. Generic Network Slice Template. https://www.gsma.com/newsroom/wp-content/uploads//NG.116-v2.0.pdf. (October 2019).

[4] Alexander Clemm, Laurent Ciavaglia, Lisandro Granville, and Jeff Tantsura. 2020. *Intent-Based Networking - Concepts and Definitions.* Internet-Draft draft-irtf-nmrg-ibn-concepts-definitions-01. IETF Secretariat. http://www.ietf.org/internet-drafts/draft-irtf-nmrg-ibn-concepts-definitions-01.txt http://www.ietf.org/internet-drafts/draft-irtf-nmrg-ibn-concepts-definitions-01.txt.

[5] The CMS Collaboration. 2008. The CMS experiment at the CERN LHC. *Journal of Instrumentation* 3, 08 (2008). https://doi.org/10.1088/1748-0221/3/08/S08004

[6] Jie Dong, Mach Chen, Dhruv Dhody, Jeff Tantsura, Kenji Kumaki, and Tomoki Murai. 2017. *BGP Extensions for Path Computation Element (PCE) Discovery.* Internet-Draft draft-dong-pce-discovery-proto-bgp-07. IETF Secretariat. http://www.ietf.org/internet-drafts/draft-dong-pce-discovery-proto-bgp-07.txt http://www.ietf.org/internet-drafts/draft-dong-pce-discovery-proto-bgp-07.txt.

[7] Zbigniew Dulinski, Piotr Wydrych, and Rafal Stankiewicz. 2015. *Inter-ALTO Communication Problem Statement.* Internet-Draft draft-dulinski-alto-inter-problem-statement-02. IETF Secretariat. http://www.ietf.org/internet-drafts/draft-dulinski-alto-inter-problem-statement-02.txt http://www.ietf.org/internet-drafts/draft-dulinski-alto-inter-problem-statement-02.txt.

[8] ETSI. [n. d.]. Zero touch network & Service Management. https://www.etsi.org/technologies/zero-touch-network-service-management. ([n. d.]).

[9] NFVISG ETSI. [n. d.]. Network functions virtualisation (NFV); management and orchestration; network service templates specification. *ETSI, Group Specification ETSI GS NFV-IFA* 14 ([n. d.]), V2.

[10] Kai Gao, xinwang2014@hotmail.com, Qiao Xiang, Chen Gu, Yang Yang, and G.Robert Chen. 2018. *Compressing ALTO Path Vectors.* Internet-Draft draft-gao-alto-routing-state-abstraction-08. IETF Secretariat. http://www.ietf.org/internet-drafts/draft-gao-alto-routing-state-abstraction-08.txt http://www.ietf.org/internet-drafts/draft-gao-alto-routing-state-abstraction-08.txt.

[11] H. Gredler, J. Medved, S. Previdi, A. Farrel, and S. Ray. 2016. *North-Bound Distribution of Link-State and Traffic Engineering (TE) Information Using BGP.* RFC 7752. RFC Editor.

[12] Arpit Gupta, Laurent Vanbever, Muhammad Shahbaz, Sean P Donovan, Brandon Schlinker, Nick Feamster, Jennifer Rexford, Scott Shenker, Russ Clark, and Ethan Katz-Bassett. 2015. Sdx: A software defined internet exchange. *ACM SIGCOMM Computer Communication Review* 44, 4 (2015), 551–562.

[13] Victor Heorhiadi, Michael K Reiter, and Vyas Sekar. 2016. Simplifying Software-Defined Network Optimization Using {SOL}. In *13th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 16).* 223–237.

[14] Kostas Katsalis, Navid Nikaein, and Andy Edmonds. 2016. Multi-domain orchestration for NFV: Challenges and research directions. In *2016 15th International Conference on Ubiquitous Computing and Communications and 2016 International Symposium on Cyberspace and Security (IUCC-CSS).* IEEE, 189–195.

[15] S. Kiesel and M. Stiemerling. 2020. *Application-Layer Traffic Optimization (ALTO) Cross-Domain Server Discovery.* RFC 8686. RFC Editor.

[16] Changhoon Kim, Anirudh Sivaraman, Naga Katta, Antonin Bas, Advait Dixit, and Lawrence J Wobker. 2015. In-band network telemetry via programmable dataplanes. In *ACM SIGCOMM.*

[17] D. King and A. Farrel. 2012. *The Application of the Path Computation Element Architecture to the Determination of a Sequence of Domains in MPLS and GMPLS.* RFC 6805. RFC Editor.

[18] Danny Lachos, Christian Rothenberg, Qiao Xiang, Y Richard Yang, Börje Ohlman, Sabine Randriamasy, Farni Boten, and Luis M Contreras. 2019. Supporting multi-domain Use cases with ALTO. In *Proceedings of the Applied Networking Research Workshop.* 59–61.

[19] lcls [n. d.]. The Linac Coherent Light Source. https://lcls.slac.stanford.edu/. ([n. d.]).

[20] Pedro R. Marques, Jared Mauch, Nischal Sheth, Barry Greene, Robert Raszuk, and Danny R. McPherson. 2009. Dissemination of Flow Specification Rules. RFC 5575. (Aug. 2009). https://doi.org/10.17487/RFC5575

[21] S Paulraj and P Sumathi. 2010. A comparative study of redundant constraints identification methods in linear programming problems. *Mathematical Problems in Engineering* 2010 (2010).

[22] Danny Perez, Qiao Xiang, Christian Rothenberg, and Y. Yang. 2020. *Multi-domain Service Function Chaining with ALTO.* Internet-Draft draft-lachos-sfc-multi-domain-alto-00. IETF Secretariat. http://www.ietf.org/internet-drafts/draft-lachos-sfc-multi-domain-alto-00.txt http://www.ietf.org/internet-drafts/draft-lachos-sfc-multi-domain-alto-00.txt.

[23] Danny Alex Lachos Perez and Christian Esteve Rothenberg. 2018. Multi-domain Orchestration leveraging the Application-Layer Traffic Optimization Protocol. In *Anais do V Workshop Pré-IETF.* SBC.

[24] K. Ramakrishnan, S. Floyd, and D. Black. 2001. *The Addition of Explicit Congestion Notification (ECN) to IP.* RFC 3168. RFC Editor. http://www.rfc-editor.org/rfc/rfc3168.txt http://www.rfc-editor.org/rfc/rfc3168.txt.

[25] Y. Rekhter, T. Li, and S. Hares. 2006. *A Border Gateway Protocol 4 (BGP-4).* RFC 4271. RFC Editor. http://www.rfc-editor.org/rfc/rfc4271.txt http://www.rfc-editor.org/rfc/rfc4271.txt.

[26] J.L. Le Roux. 2006. *Requirements for Path Computation Element (PCE) Discovery.* RFC 4674. RFC Editor.

[27] Philipp S. Schmidt, T. Enghardt, R. Khalili, and A. Feldmann. 2013. Socket Intents: Leveraging Application Awareness for Multi-Access Connectivity. In *Proceedings of CoNEXT '13.* ACM, New York, NY, USA.

[28] ska [n. d.]. The Square Kilometre Array. https://www.skatelescope.org/. ([n. d.]).

[29] Jan Telgen. 1983. Identifying redundant constraints and implicit equalities in systems of linear constraints. *Management Science* 29, 10 (1983), 1209–1222.

[30] JP. Vasseur, R. Zhang, N. Bitar, and JL. Le Roux. 2009. *A Backward-Recursive PCE-Based Computation (BRPC) Procedure to Compute Shortest Constrained Inter-Domain Traffic Engineering Label Switched Paths.* RFC 5441. RFC Editor.

[31] Qiao Xiang, Chin Guok, Franck Le, John MacAuley, Harvey Newman, and Y Richard Yang. 2018. SFP: Toward Interdomain Routing for SDN Networks. In *Proceedings of the ACM SIGCOMM 2018 Conference on Posters and Demos.* ACM, 87–89.

[32] Qiao Xiang, X Tony Wang, J Jensen Zhang, Harvey Newman, Y Richard Yang, and Y Jace Liu. 2019. Unicorn: Unified resource orchestration for multi-domain, geo-distributed data analytics. *Future Generation Computer Systems* 93 (2019), 188–197.

[33] Qiao Xiang, Jensen Zhang, Franck Le, Y. Yang, and Harvey Newman. 2020. *Resource Orchestration for Multi-Domain, Exascale, Geo-Distributed Data Analytics.* Internet-Draft draft-xiang-alto-multidomain-analytics-03. IETF Secretariat. http://www.ietf.org/internet-drafts/draft-xiang-alto-multidomain-analytics-03.txt http://www.ietf.org/internet-drafts/draft-xiang-alto-multidomain-analytics-03.txt.

[34] Qiao Xiang, J Jensen Zhang, X Tony Wang, Y Jace Liu, Chin Guok, Franck Le, John MacAuley, Harvey Newman, and Y Richard Yang. 2018. Fine-grained, multi-domain network resource abstraction as a fundamental primitive to enable high-performance, collaborative data sciences. In *SC18: International Conference for High Performance Computing, Networking, Storage and Analysis.* IEEE, 58–70.

[35] Qiao Xiang, Jingxuan Jensen Zhang, Xin Tony Wang, Yang Jace Liu, Chin Guok, Franck Le, John MacAuley, Harvey Newman, and Y Richard Yang. 2019. Toward Fine-Grained, Privacy-Preserving, Efficient Multi-Domain Network Resource Discovery. *IEEE Journal on Selected Areas in Communications* 37, 8 (2019), 1924–1940.