

# Supporting Multi-domain Use Cases with ALTO

Danny Lachos, Christian Rothenberg  
University of Campinas

Qiao Xiang, Y. Richard Yang  
Yale University

Börje Ohlman  
Ericsson

Sabine Randriamasy  
Nokia Bell Labs

Farni Boten  
Sprint

Luis M. Contreras  
Telefonica

## ABSTRACT

Many multi-domain use cases can benefit substantially from network information exposure, but also introduce new, key requirements that existing exposure solutions, such as the Application-Layer Traffic Optimization (ALTO) protocol, do not satisfy. In this talk, we review several important multi-domain use cases, discuss the key network information exposure requirements to support these use cases, and present a unified exposure architecture as well as novel mechanisms and abstractions to substantially improve the ALTO framework in the multi-domain setting.

## CCS CONCEPTS

• **Networks** → **Network protocols**;

## KEYWORDS

ALTO, application-layer traffic optimization, multi-domain

## 1 INTRODUCTION

Many multi-domain use cases are emerging with the development of new technologies, such as software-defined networking (SDN), network function virtualization (NFV), and 5G. Examples of such use cases include multi-domain, collaborative data sciences [3, 10, 11, 20], multi-domain service function chaining (SFC) [2, 4, 7, 9], and multi-domain SDN [6, 12, 23]. Such use cases can benefit substantially from the exposure of network information, with which users can perform application-layer resource optimization to improve the performance.

The Application-Layer Optimization Protocol (ALTO) [1] already introduces basic mechanisms (e.g., modularity, dependency) and abstractions (e.g., map services) for applications to take optimized actions based on network information. For example, a major research network has adopted ALTO as a key building block. However, exposing network information to support multi-domain use cases places additional requirements that existing solutions such as the current ALTO design do not satisfy. First, abstractions that *aggregate multiple networks into a single, virtual network* are required to simplify

the application-layer optimization conducted by end users. Second, such abstractions need to provide *a unified representation of multiple resources* (e.g., networking, computation, and storage) in multiple networks.

In this talk, we review several important multi-domain use cases that can benefit substantially from network information exposure (Sec. 2). We then elaborate the key design requirements of network information exposure to support these use cases (Sec. 3). Next, we summarize the current standardization efforts in the IETF ALTO working group to support multi-domain use cases. Specifically, we present a unified network information exposure architecture to use ALTO and its extensions to support these emerging multi-domain use cases [13–15, 24–26], and novel mechanisms and abstractions based on recent research to improve the ALTO framework in the multi-domain set-up [28, 29] (Sec. 4).

## 2 USE CASES

Multi-domain network information exposure can be beneficial in supporting multiple important use cases, including multi-domain, collaborative data sciences, multi-domain SFC, and multi-domain SDN.

**Multi-domain, collaborative data sciences:** Many of today’s premier science experiments, such as the Large Hadron Collider (LHC) [11] and the Square Kilometre Array (SKA) [20], rely on finely-tuned workflows that coordinate geographically distributed resources (e.g., instrument, compute, storage) to enable scientific discoveries. One example is the movement of LHC data from Tier 0 (*i.e.*, the data center at European Organization for Nuclear Research, known as CERN) to Tier 1 (*i.e.*, national laboratories) storage sites around the world. Another example is that the Fermilab is experimenting moving the exascale LHC workflow to Amazon EC2 for more computation power [8]. The key to supporting these distributed workflows is the ability to orchestrate multiple resources across multiple network domains to facilitate predictable workflow performance (e.g., available bandwidth, packet loss rate [22]). As such, multi-domain network information exposure is a cornerstone to enable this ability.

**Multi-domain SFC:** This use case refers to building end-to-end services by composing multiple service functions in an abstract sequence across multiple network domains [7]. It is identified as an important value-added service in 5G [4, 9]. Exposing multi-domain network and resource information (e.g., link bandwidth, CPU utilization) can substantially improve the efficiency of constructing and managing such SFCs. **Multi-domain SDN:** Network providers are expanding the fine-grained capability of SDN from intradomain set-up to multi-domain setting to provide flexible interdomain routing as a valuable service [6, 12, 23]. Users of this service can specify routing actions at the provider network based

---

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

ANRW '19, July 22, 2019, Montreal, QC, Canada

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6848-3/19/07...\$15.00

<https://doi.org/10.1145/3340301.3341126>

on flexible matching conditions of flow parameters such as TCP/IP 5-tuple. This service requires provider networks to expose their available routing information to users. However, handling routing information of each network individually is too complex for users. As such, a multi-domain network exposure solution that aggregates information of multiple networks into a single abstraction can simplify the use of this service.

### 3 REQUIREMENTS ON NETWORK INFORMATION EXPOSURE

Supporting previous use cases with multi-domain network information exposure places new, key requirements, which are not satisfied by existing exposure solutions.

**Requirement 1: multi-domain aggregation abstraction.** Handling information of multiple domains may add substantial complexity to end users conducting application-layer optimization, and cause optimization errors. Exposing information on multiple networks in an aggregated view provides simplicity and robustness. As such, abstractions that aggregate multiple networks into a single, virtual network ("one-big-network") are a key requirement.

**Requirement 2: unified resource representation.** Existing exposure systems [1, 16, 21] provide separate representations of different resources. For example, BGP only exposes the routing resource information of networks, and the base ALTO protocol provides information on networking and endpoint resources, such as computation and storage, in separate abstractions. However, optimizing the performance of new multi-domain applications requires the orchestration of multiple resources in multiple networks (e.g., RAN, transport, core in 5G) where dynamics and topologies are completely different. As such, a unified representation of multiple resources is another key requirement for multi-domain network information exposure to support multi-domain use cases.

### 4 MULTI-DOMAIN INFORMATION AGGREGATION ABSTRACTION ARCHITECTURE AND MECHANISMS

ALTO already provides basic mechanisms and abstractions to expose network information. In this talk, we present the architecture unifying these mechanisms and abstractions to satisfy the requirements for multi-domain use cases (Sec. 4.1). Moreover, we present novel mechanisms and abstractions based on recent research to allow the ALTO framework supporting the important multi-domain settings (Sec. 4.2).

#### 4.1 Existing ALTO Efforts: Architecture

ALTO provides a generic framework to expose network information for applications to improve their performance. Figure 1 presents a high-level overview of key ALTO mechanisms and abstractions. In particular, ALTO introduces *generic mechanisms* such as: (i) information resource directory (IRD) [1], (ii) information consistency (tag, dependency, multi-info resources [30]), and (iii) information update model (e.g., incremental update with server-sent events [18]). ALTO also introduces *abstractions* exposing network information to the applications: (i) network and cost maps [1], (ii) the

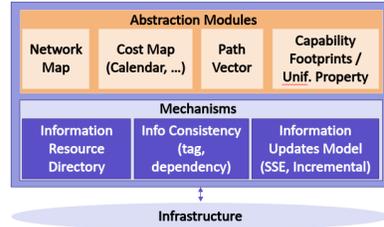


Figure 1: High level ALTO architecture.

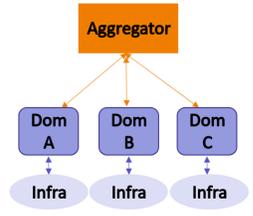


Figure 2: Multi-domain network abstraction.

path vector abstraction [5], and (iii) capability maps (e.g., CDNI [19] and unified property Map [17]).

#### 4.2 Multi-Domain Network Abstractions

We extend the preceding architecture with novel mechanisms and abstractions for multi-domain aggregation.

**Multi-domain aggregation:** Figure 2 shows a multi-domain aggregation mechanism on top of the existing single domain architecture. The new mechanism aggregates network information from ALTO servers in multiple networks to provide a single, consistent, updated, "virtual" domain abstraction. Network maps, cost maps, unified entity properties, network capabilities, and routing path abstractions (path vectors) of individual networks are consistently integrated to provide the abstraction of a single, coherent network to the users, satisfying the multi-domain aggregation requirement. This mechanism also includes a security multi-party computation protocol to protect the privacy of individual networks. In this talk, we summarize our recent research on the design options of the aggregation layer [13, 28, 29].

**Multi-resource abstraction:** Although the existing abstractions (network/cost map, unified property, and path vector) are already powerful, they cannot handle the multi-resource information requirement. We present our recent, unified resource abstraction, based on mathematical programming constraints as a generic representation of the feasible resource capability of networks [27], which users can consume via different resource management systems.

**Multi-domain programming information abstraction:** As multi-domain SDN requires multi-domain SDN resource and programming abstraction, we present our recent research on multi-domain SDN abstraction. In particular, this new abstraction computes a single, abstract, programmable network spanning multiple individual networks. Unlike existing SDN abstractions (e.g., OpenFlow and P4), it includes a built-in layer to extract and learn the interactions of inter-domain policies of individual networks (e.g., route selection preference), providing a unified abstraction.

### 5 CONCLUSIONS

In this talk, we review important multi-domain use cases that can benefit from network information exposure, and discuss the key requirements on exposure solutions. We then present a unified exposure architecture that uses ALTO to support these use cases, and introduces novel mechanisms and abstractions based on recent research to substantially improve the ALTO framework in the multi-domain setting.

## 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] 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>
- [4] ETSI. [n. d.]. Zero touch network & Service Management. <https://www.etsi.org/technologies/zero-touch-network-service-management>. ([n. d.]).
- [5] Kai Gao, Young Lee, Sabine Randriamasy, Yang Yang, and J. Zhang. 2019. *ALTO Extension: Path Vector Cost Type*. Internet-Draft draft-ietf-alto-path-vector-05. IETF Secretariat. <http://www.ietf.org/internet-drafts/draft-ietf-alto-path-vector-05.txt> <http://www.ietf.org/internet-drafts/draft-ietf-alto-path-vector-05.txt>.
- [6] 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.
- [7] J. Halpern and C. Pignataro. 2015. *Service Function Chaining (SFC) Architecture*. RFC 7665. RFC Editor.
- [8] Burt Holzman, Lothar AT Bauerdick, Brian Bockelman, Dave Dykstra, Ian Fisk, Stuart Fuess, Gabriele Garzoglio, Maria Gironi, Oliver Gutsche, Dirk Hufnagel, et al. 2017. Hepcloud, a new paradigm for hep facilities: Cms amazon web services investigation. *Computing and Software for Big Science* 1, 1 (2017), 1.
- [9] 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.
- [10] lcls [n. d.]. The Linac Coherent Light Source. <https://lcls.slac.stanford.edu/>. ([n. d.]).
- [11] lhc [n. d.]. The Large Hadron Collider (LHC) Experiment. <https://home.cern/topics/large-hadron-collider>. ([n. d.]).
- [12] 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>
- [13] Danny Perez and Christian Rothenberg. 2019. *ALTO-based Broker-assisted Multi-domain Orchestration*. Internet-Draft draft-lachosrothenberg-alto-brokermdo-02. IETF Secretariat. <http://www.ietf.org/internet-drafts/draft-lachosrothenberg-alto-brokermdo-02.txt> <http://www.ietf.org/internet-drafts/draft-lachosrothenberg-alto-brokermdo-02.txt>.
- [14] Danny Perez and Christian Rothenberg. 2019. *Multi-domain E2E Network Services*. Internet-Draft draft-lachosrothenberg-alto-md-e2e-ns-00. IETF Secretariat. <http://www.ietf.org/internet-drafts/draft-lachosrothenberg-alto-md-e2e-ns-00.txt> <http://www.ietf.org/internet-drafts/draft-lachosrothenberg-alto-md-e2e-ns-00.txt>.
- [15] Danny Perez, Qiao Xiang, Christian Rothenberg, and Y. Yang. 2018. *Multi-domain Service Function Chaining with ALTO*. Internet-Draft draft-lachos-multi-domain-sfc-alto-00. IETF Secretariat. <http://www.ietf.org/internet-drafts/draft-lachos-multi-domain-sfc-alto-00.txt> <http://www.ietf.org/internet-drafts/draft-lachos-multi-domain-sfc-alto-00.txt>.
- [16] Yakov Rekhter, Tony Li, and Susan Hares. 2005. *A border gateway protocol 4 (BGP-4)*. Technical Report.
- [17] Wendy Roome, Sabine Randriamasy, Yang Yang, and J. Zhang. 2019. *Unified Properties for the ALTO Protocol*. Internet-Draft draft-ietf-alto-unified-props-new-07. IETF Secretariat. <http://www.ietf.org/internet-drafts/draft-ietf-alto-unified-props-new-07.txt> <http://www.ietf.org/internet-drafts/draft-ietf-alto-unified-props-new-07.txt>.
- [18] Wendy Roome and Yang Yang. 2019. *ALTO Incremental Updates Using Server-Sent Events (SSE)*. Internet-Draft draft-ietf-alto-incr-update-sse-16. IETF Secretariat. <http://www.ietf.org/internet-drafts/draft-ietf-alto-incr-update-sse-16.txt> <http://www.ietf.org/internet-drafts/draft-ietf-alto-incr-update-sse-16.txt>.
- [19] Jan Seedorf, Yang Yang, Kevin Ma, Jon Peterson, and Xiao Lin. 2019. *Content Delivery Network Interconnection (CDNI) Request Routing: CDNI Footprint and Capabilities Advertisement using ALTO*. Internet-Draft draft-ietf-alto-cdni-request-routing-alto-05. IETF Secretariat. <http://www.ietf.org/internet-drafts/draft-ietf-alto-cdni-request-routing-alto-05.txt> <http://www.ietf.org/internet-drafts/draft-ietf-alto-cdni-request-routing-alto-05.txt>.
- [20] ska [n. d.]. The Square Kilometre Array. <https://www.skatelescope.org/>. ([n. d.]).
- [21] Douglas Thain, Todd Tannenbaum, and Miron Livny. 2005. Distributed computing in practice: the Condor experience. *Concurrency and computation: practice and experience* 17, 2-4 (2005), 323–356.
- [22] Qin Wu, Yang Yang, Young Lee, Dhruv Dhody, and Sabine Randriamasy. 2018. *ALTO Performance Cost Metrics*. Internet-Draft draft-ietf-alto-performance-metrics-06. IETF Secretariat. <http://www.ietf.org/internet-drafts/draft-ietf-alto-performance-metrics-06.txt> <http://www.ietf.org/internet-drafts/draft-ietf-alto-performance-metrics-06.txt>.
- [23] 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.
- [24] Qiao Xiang, Franck Le, and Yang Yang. 2019. *ALTO Extension: Unified Resource Representation*. Internet-Draft draft-xiang-alto-unified-representation-01. IETF Secretariat. <http://www.ietf.org/internet-drafts/draft-xiang-alto-unified-representation-01.txt> <http://www.ietf.org/internet-drafts/draft-xiang-alto-unified-representation-01.txt>.
- [25] Qiao Xiang, Franck Le, and Yang Yang. 2019. *ALTO for Multi-Domain Applications: A Review of Use Cases and Design Requirements*. Internet-Draft draft-xiang-alto-multidomain-usecases-00. IETF Secretariat. <http://www.ietf.org/internet-drafts/draft-xiang-alto-multidomain-usecases-00.txt> <http://www.ietf.org/internet-drafts/draft-xiang-alto-multidomain-usecases-00.txt>.
- [26] Qiao Xiang, Franck Le, Yang Yang, Harvey Newman, and Haizhou Du. 2018. *Unicorn: Resource Orchestration for Multi-Domain, Geo-Distributed Data Analytics*. Internet-Draft draft-xiang-alto-multidomain-analytics-02. IETF Secretariat. <http://www.ietf.org/internet-drafts/draft-xiang-alto-multidomain-analytics-02.txt> <http://www.ietf.org/internet-drafts/draft-xiang-alto-multidomain-analytics-02.txt>.
- [27] 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.
- [28] Qiao Xiang, Haitao Yu, James Aspnes, Franck Le, Linghe Kong, and Y. Richard Yang. 2019. Optimizing in the Dark: Learning an Optimal Solution Through a Simple Interface. AAAI 2019.
- [29] 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.
- [30] J. Zhang. 2019. *Multiple ALTO Resources Query Using Multipart Message*. Internet-Draft draft-zhang-alto-multipart-02. IETF Secretariat. <http://www.ietf.org/internet-drafts/draft-zhang-alto-multipart-02.txt> <http://www.ietf.org/internet-drafts/draft-zhang-alto-multipart-02.txt>.