

ALTO IETF 107 Virtual Interim

Supporting Multi-domain Use Cases with ALTO

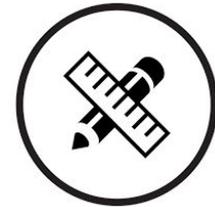
Danny Alex Lachos Perez

Christian E. Rothenberg, Qiao Xiang, Richard yang, Börje Ohlman, Sabine Randriamasy, Farni Boten, Luis M. Contreras, Jensen Zhang, Kai Gao.

April 21, 2020

Presentation in a Nutshell

- Key questions for ALTO to support multi-domain* use cases:
 - **What information** do multi-domain applications need?
 - **What are the issues** of gathering that information?
 - **How to design** a complete framework?

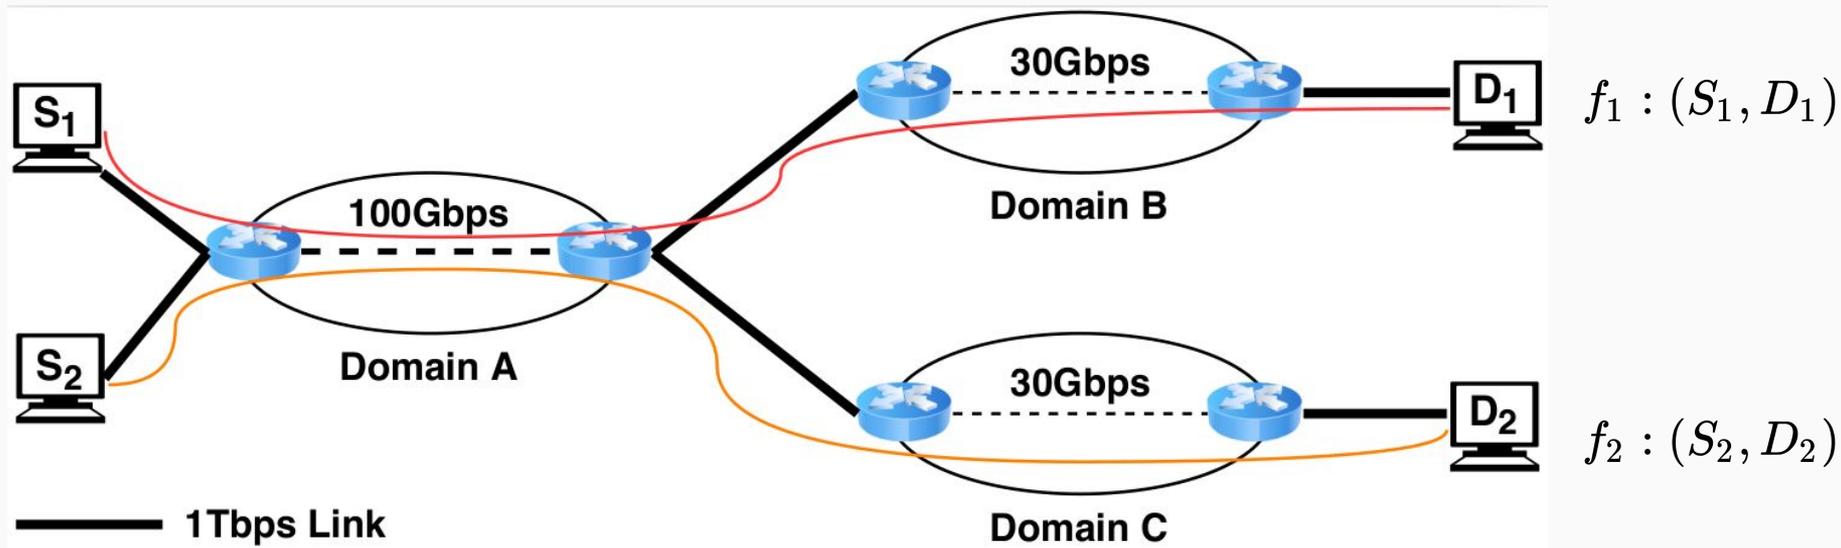


*A domain is considered to be a separate administrative environment.

What information do multi-domain applications need?

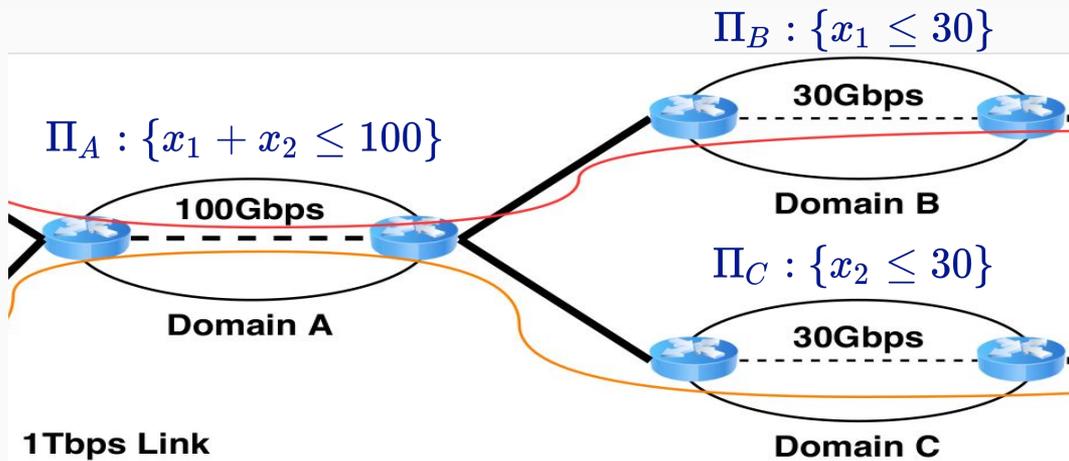
Motivating use case example

Collaboration Network



- **End-to-End cost** across multiple domains
 - Resources availability (e.g., bandwidth) and sharing
- **Sequence of domains and candidate paths**
 - Which domains are involved for the different traffic flows
 - One or more potential paths connecting such domains

Single/Multi domain Network Properties



- Consider services provided by the domains as an object f_i .
- f_i has a set of **network properties**:
 - $f_i.abw$
 - $f_i.path$
 - $f_i.delay$,
 - $f_i.UtilizationCharge$, etc.

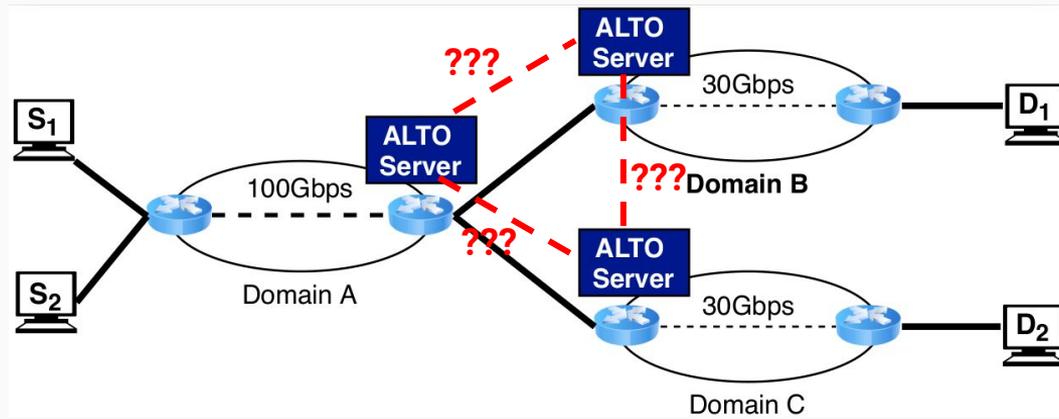
- **Multi-domain Network Properties:** A network property to a flow f_i may involve properties of multiple component domains:
 - $f_i.md-abw: \min(f_i.abw[A] + f_i.abw[B] + f_i.abw[C])$
 - $f_i.md-path: f_i.path[A] \cdot f_i.path[B] \cdot f_i.path[C]$
 - $f_i.md-delay: f_i.delay[A] + f_i.delay[B] + f_i.delay[C]$
 - etc.

What are the issues of gathering that information?

ALTO Key Issues

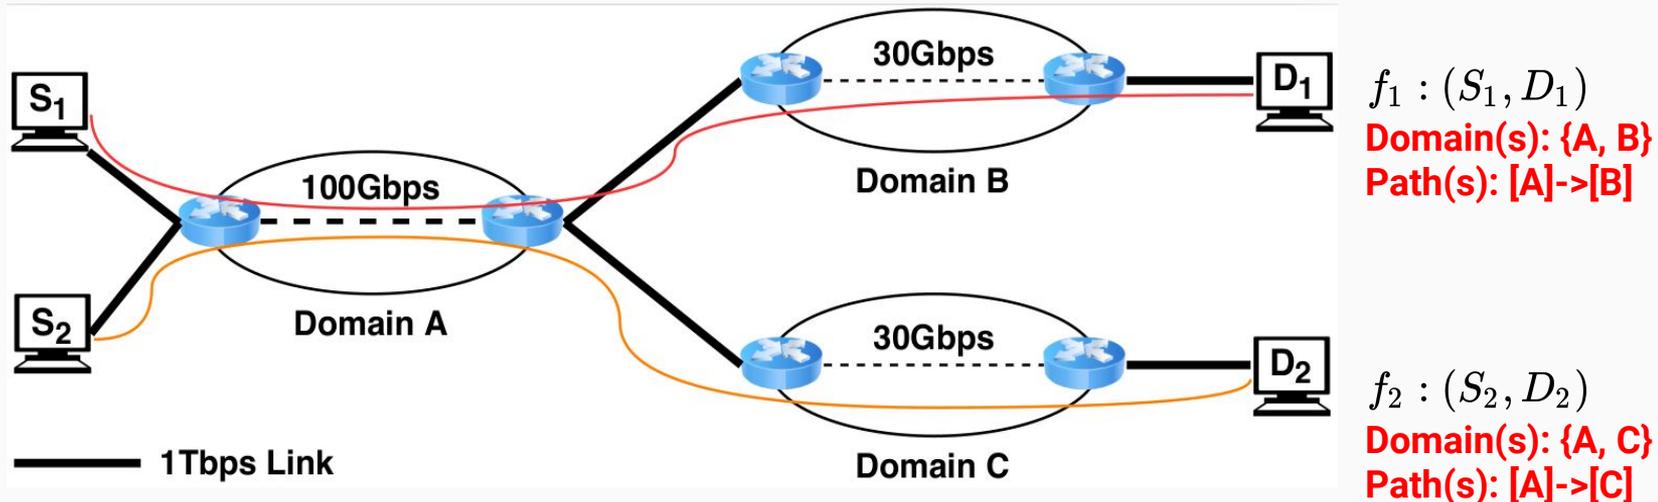
ALTO Servers Communication

- Server-to-Client ALTO communication is not enough.
- It is necessary **multi-ALTO server communication** to allow exchanging network information from multiple domains.
- The ALTO protocol specification states [[RFC7285](#)]:
 - "It may also be possible for an ALTO server to exchange network information with other ALTO servers (either within the same administrative domain or another administrative domain with the consent of both parties) ...".
 - However, such a **protocol is outside the scope** of the specification.



Connectivity Discovery

- The connectivity information is the reachability between source nodes and the destination nodes. It involves:
 - Discover **which domains are involved** in the data movement of each node pair.
 - Discover a **set of candidate paths** in order to know how to reach a remote destination node.
- Multi-domain mechanisms combining **domains sequence computation** and **paths computation** need to be defined.



ALTO Server Discovery

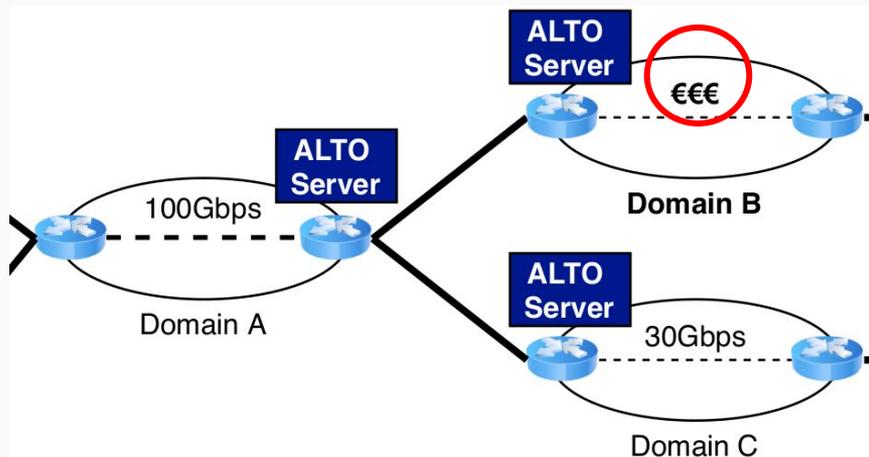
- An application (as an ALTO client) needs to be **aware of the presence and the location of ALTO servers** in order to get appropriate guidance.
- ALTO servers will be located in different network domains, so that **multi-domain ALTO server discovery** mechanisms are needed.



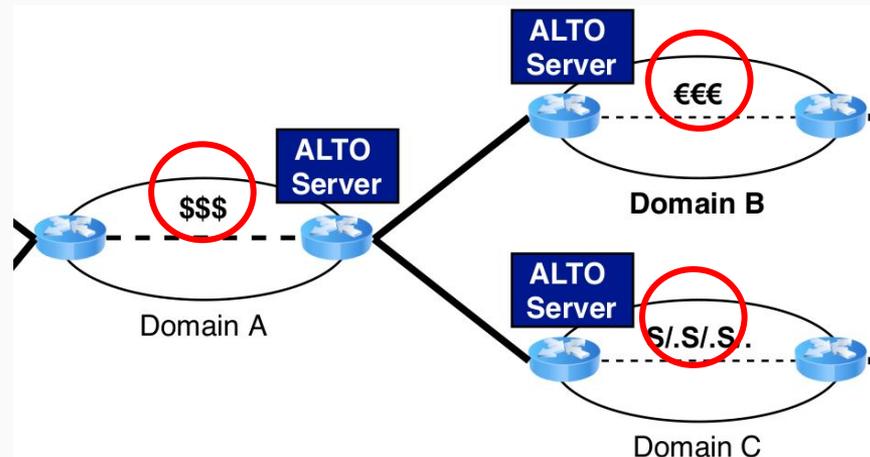
Single-domain Composition

- Each domain may have its **own representation** of the same network. E.g.:

Property values may not be comparable together
(utilization charge and available bandwidth)



Same utilization charge property but there would not necessarily be a uniform form of billing



Simple Resource Query Language

- Applications need a query to express all **common resource requirements** to the network.
 - E.g., flow f_1 may provide application's requirements:
 - **Reachability requirements:** "from S1 to D1"
 - **Bi-direction symmetry:** "Data traffic from S1 to D1 and from D1 to S1"
 - **Waypoint traversal:** "f1 must traverse one middlebox m1"
 - **QoS metrics:** "the bandwidth of the flow f1 needs to be at least 30 Gbps"
 - Etc.
- ALTO provides a very **simple ALTO resource query** (e.g., filtered network/cost map).

Scalability / Privacy

SCALABILITY

- Solving the optimization problems, specified by the applications' requirements, can be **computationally expensive and time-consuming**.
 - E.g., the number of available paths for each flow is increased exponentially with the number of domains involved.
 - As such, the number of available configurations for a set of flows would also increase exponentially with both the network size and the number of flows.

PRIVACY

- The information provided by the ALTO base protocol is considered coarse-grained in several recent multi-domain use cases.
- As a result, several ALTO extension services have been designed to provide **fine-grained network information** to the application.
 - Using these ALTO extension services for multi-domain scenarios would **raise new security and privacy concerns**.

How to design a complete framework?

ALTO Architecture

From Single to Multi -domain: Mechanisms

From	To	Mechanism(s)
Server-to-Client ALTO communication	ALTO Servers Communication	<ul style="list-style-type: none">ALTO may consider either a hierarchical or mesh architectural deployment design [INTER-ALTO][MERCATOR][SFC-MD]<ul style="list-style-type: none">Hierarchical design, ALTO servers in domain partitions gather local information and send it to central server.Mesh deployment, ALTO servers may be set up in each domain independently, and gathering the network information from other connected domains. <div data-bbox="761 653 1850 1005"><p>The diagram illustrates two architectural models for ALTO server communication. On the left, labeled 'Hierarchical', three separate domain clouds are shown at the top. Each cloud contains a server icon and a router icon connected by a dashed line. Solid arrows from each of these three clouds point downwards to a single central server icon located in a larger cloud at the bottom. On the right, labeled 'Mesh', three domain clouds are arranged in a triangle. Each cloud contains a server icon and a router icon connected by a dashed line. Solid arrows connect the server icons of each cloud to the server icons of the other two clouds, forming a mesh topology.</p></div>

From Single to Multi -domain: Mechanisms

From	To	Mechanism(s)
Connectivity Discovery	Multi-domain Connectivity discovery	<ul style="list-style-type: none">• Multi-domain mechanisms combining domains sequence computation and paths computation need to be defined.• Standardized computation protocols may be considered:<ul style="list-style-type: none">○ BGP [RFC4271]: Provides multi-domain sequence computation (It does not advertise multiple alternative routes).○ BGP-LS [RFC7752]: Allows visibility of the network topology and export traffic engineering information with external domains using the BGP routing protocol.○ PCE-based [RFC5441, RFC6805]: Define mechanisms where a PCE entity cooperates either with other PCE entities in adjacent domains or with a parent PCE entity.

From Single to Multi -domain: Mechanisms

From	To	Mechanism(s)
ALTO Server Discovery	Multi-domain ALTO Server Discovery	<ul style="list-style-type: none">● [ALTO-RFC8686] specifies a procedure for identifying ALTO servers outside of the ALTO client's own network domain.● [RFC4674] proposes a set of functional requirements to allow a path computation client (PCC) to automatically and dynamically discover the location of PCEs entities (including additional information about supported capabilities) for each controller domain.● [PROTO-BGP] is defining extensions to BGP to also carry PCE discovery information. Specifically, this document extends BGP to allow a PCE entities to advertise its location and some information useful to a PCC for the PCE selection.

From Single to Multi -domain: Mechanisms

From	To	Mechanism(s)
Single-domain Composition	Unified resource Representation	<ul style="list-style-type: none">Multi-domain composition mechanisms are necessary<ul style="list-style-type: none">Network information from ALTO servers in multiple domains need to be adapted together to a single and consistent "virtual" domain abstraction [UNICORN][MERCATOR]. <p>$\Pi_A : \{x_1 + x_2 \leq 100\}$ $\Pi_B : \{x_1 \leq 30\}$ $\Pi_C : \{x_2 \leq 30\}$  $\{x_1 \leq 30\}, \{x_2 \leq 30\}$</p>
Simple resource query Language	Generic/Flexible Query Language	<ul style="list-style-type: none">With a flexible/generic query language:<ul style="list-style-type: none">ALTO can filter out a large number of unqualified configurationsALTO can selectively send the resource information of a small number of qualified domains to the applications. <p><i>f1: {src_ip = 10.0.0.1 and dst_ip = 10.0.0.2};</i> <i>flow_set: {f1};</i> <i>req1: f1.bandwidth >= 30 Gbps;</i> SELECT bandwidth from flow_set WHERE req1;</p>

From Single to Multi -domain: Mechanisms

From	To	Mechanism(s)
Scalability	Scalability++	<ul style="list-style-type: none">● ALTO servers need to support mechanisms (e.g., pre-computation and projection [DRAFT-RSA]) 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.
Privacy	Privacy Preserving	<ul style="list-style-type: none">● ALTO needs mechanisms (with little overhead) that provide accurate sharing network information, and at the same time protects each member domain.● [MD-ANALYTICS] presents a privacy-preserving, multi-domain extension of ALTO.<ul style="list-style-type: none">○ It supports the privacy-preserving discovery of networking resource sharing among a set of candidate endpoint pairs.

Next Steps

Next Steps

- Gather **feedback** from the WG
 - Comments, questions, suggestions are greatly appreciated.
- **Engage** more people from the academy/industry
- **Next steps**
 - Systematic investigation of **extensions** to address corresponding Key design issues
 - Systematic investigation of **deployment concerns** of ALTO for **multi-domain** applications
- **Additional Information**
 - SIGCOMM'20 Network-Application Integration/CoDesign Workshop
 - Submission deadline: April 30, 2020
 - Workshop date: August 14, 2020
 - Link: <https://conferences.sigcomm.org/sigcomm/2020/workshop-nai.html>
 - Internal meetings
 - Wednesday weekly meetings (9:30 US ET)
 - Bridge: <https://yale.zoom.us/j/8423318713>