



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

Danny A. Lachos*

Christian E. Rothenberg*

Qiao Xiang†

Y. Richard Yang‡

Börje Ohlman#

Sabine Randriamasy§

Luis M. Contreras¶

Kai Gao&



*Unicamp



†Yale



‡Ericsson

NOKIA

§Nokia Bell Labs

Telefonica

¶Telefónica



&Sichuan

Introduction

The basics first

- A logical entity that provides interfaces (e.g., REST-ful APIs) to consult the ALTO information services

ALTO Server

- A logical entity that sends ALTO queries to obtain guiding information from the ALTO server

ALTO Client

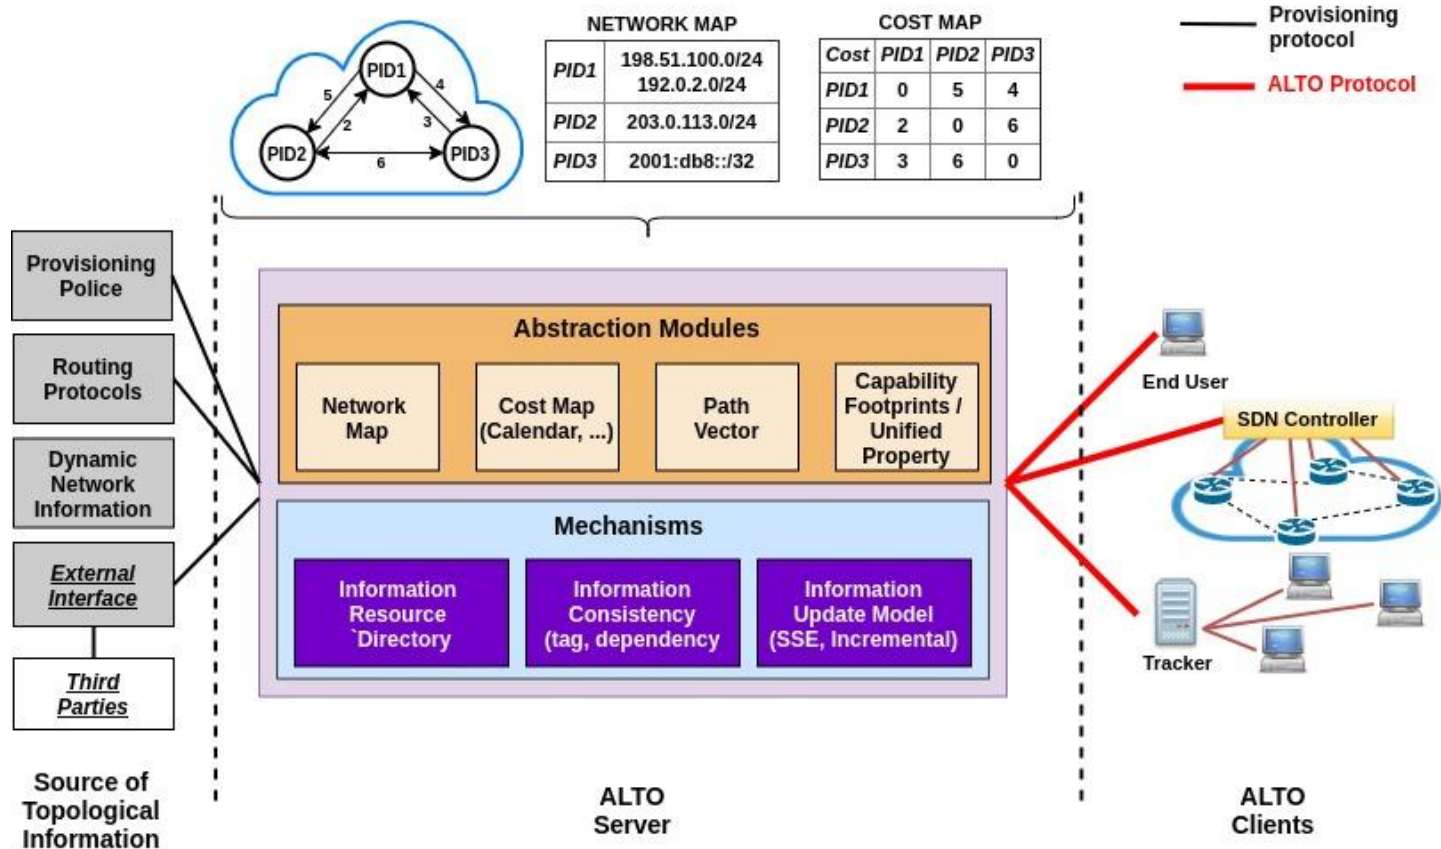
- Communication protocol between ALTO Server and ALTO Client(s)
- Based on existing HTTP implementations (RESTful and JSON)

ALTO Protocol

- Chartered in 2008
- First RFC - RFC5693 (2009)
- Recharter discussions (**In progress**)

ALTO WG

ALTO architecture (RFC7285)



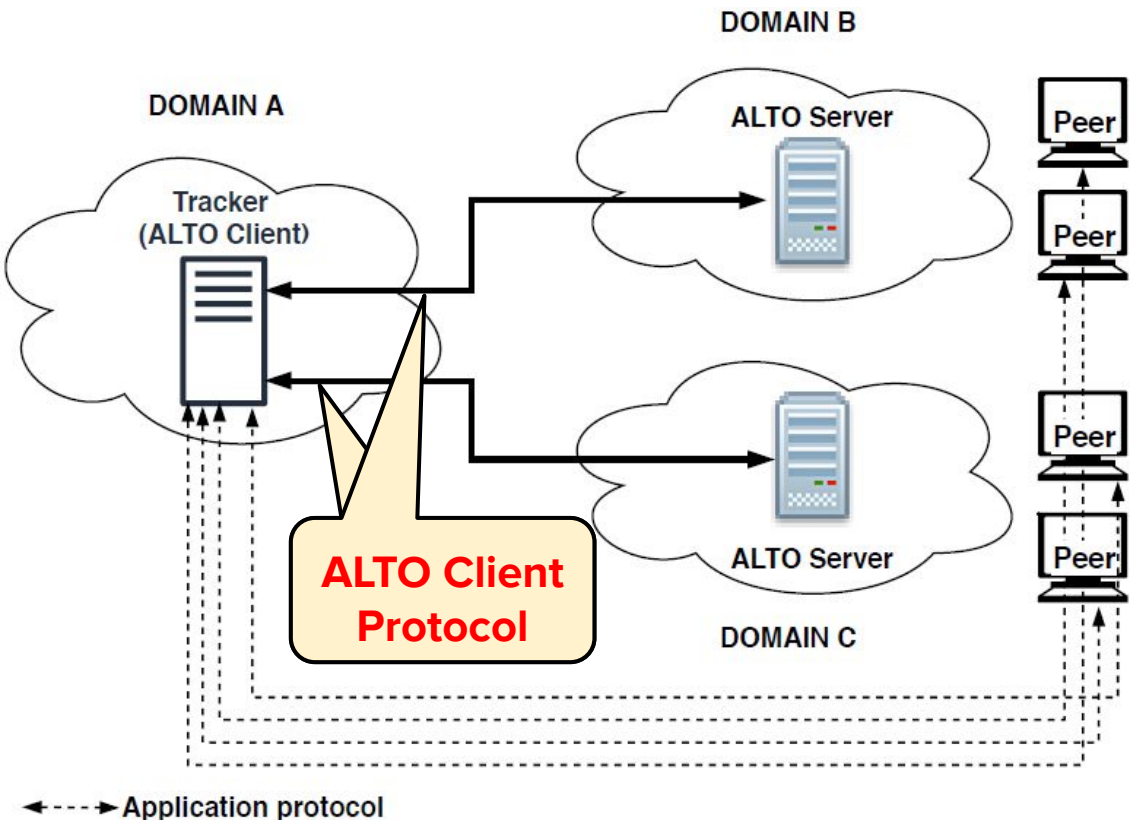
What does “multi-domain” mean?

Multi-domain: Scope

- A **domain** is considered to be a **network region** in the global Internet.
 - Each domain has a **(partial) network view** from the perspective of the network region.
 - **Network region examples:** An ISP/AS, a set of ASes/ISPs, transport/access/science networks, mobile edge clouds, etc.
- The **multi-domain** approach involves **multiple network regions** with different technology, administration or ownership.
- A common setting for **many novel use cases** is that the traffic from a source to a destination traverses multiple domains.
 - **Use case examples:** data intensive science applications, multi-domain SFC, flexible inter-domain routing control, etc.

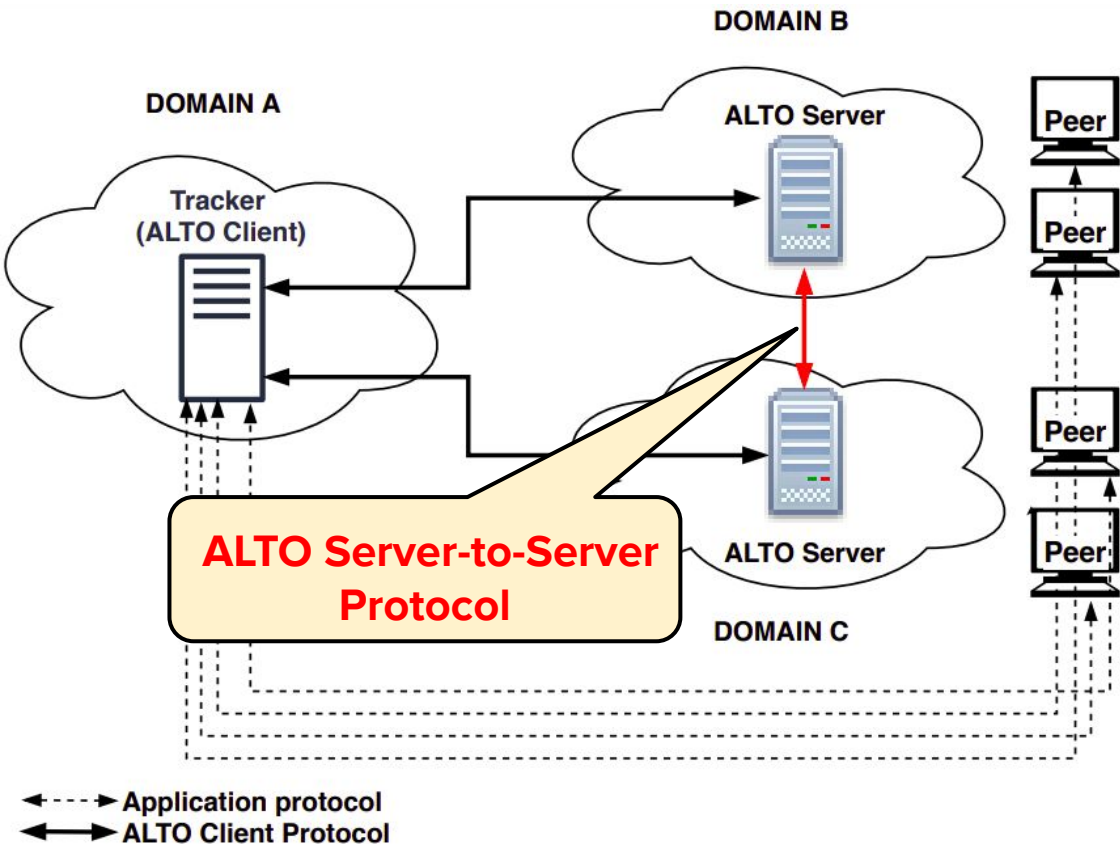
**From single- to multi-domain
information exposure using ALTO**

Single-domain info. exposure using ALTO



- The ALTO server in each domain will **provide only local information** to ALTO clients.
- The ALTO client (Tracker), in domain A, will **receive partial network information** from domain B or domain C.

Multi-domain info. exposure using ALTO



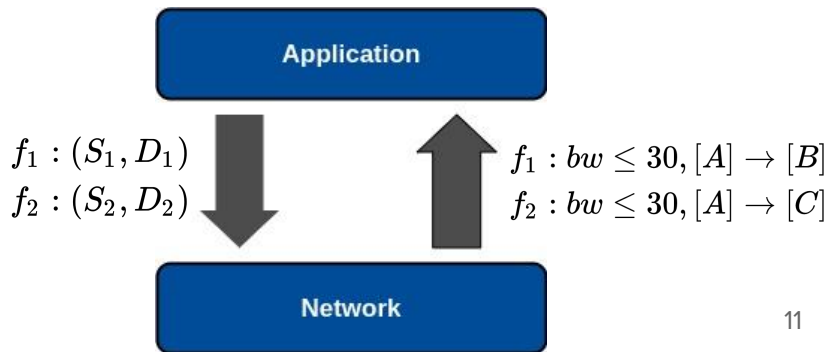

- ALTO servers will **exchange information** and the ALTO client will **receive merged information** from multiple domains
- In the example, the tracker will receive merged information from domain A and domain B.

What information do multi-domain applications need?

and how does the network provide such information?

Application/Network interaction

- Application interacts with networks by asking the networks to **carry traffic for a set of flows**: $[f_1, f_2, \dots, f_n]$
- The network **provide resource/topology information** for applications:
 - **End-to-End (E2E) cost** across multiple domains
 - Resource 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




What are the issues of gathering multi-domain information?

Current ALTO design


Server-to-Client ALTO communication

- Server-to-Client ALTO communication **is not enough**
- 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.



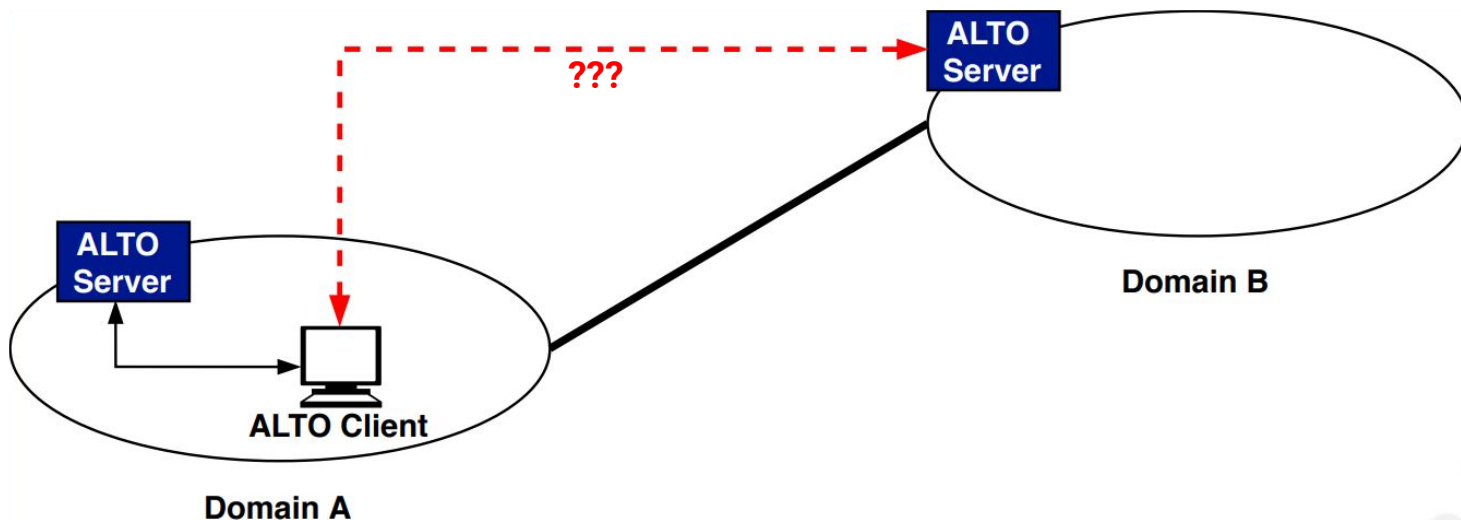
Domain connectivity discovery

- 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.
- The current ALTO extensions do not have this feature.



ALTO server discovery

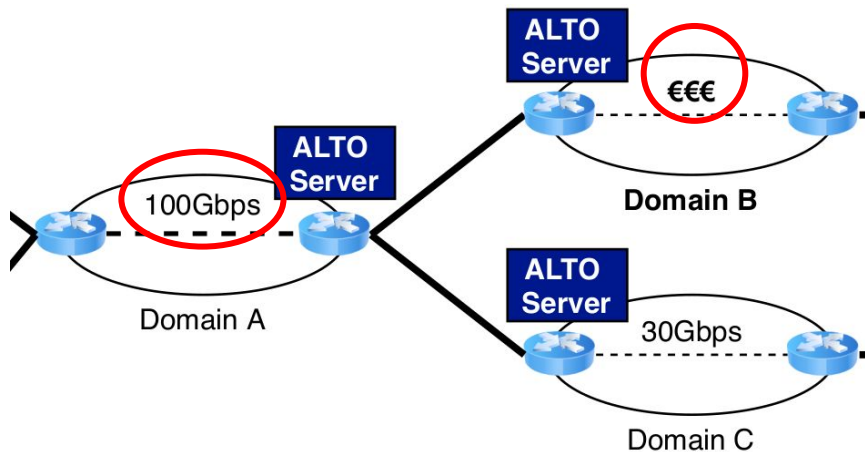
- An application (as an ALTO client) needs to be **aware of the presence and the location of ALTO servers** in order to get appropriate guidance.
- ALTO servers will be located in different network domains, so that **multi-domain ALTO server discovery** mechanisms are needed.



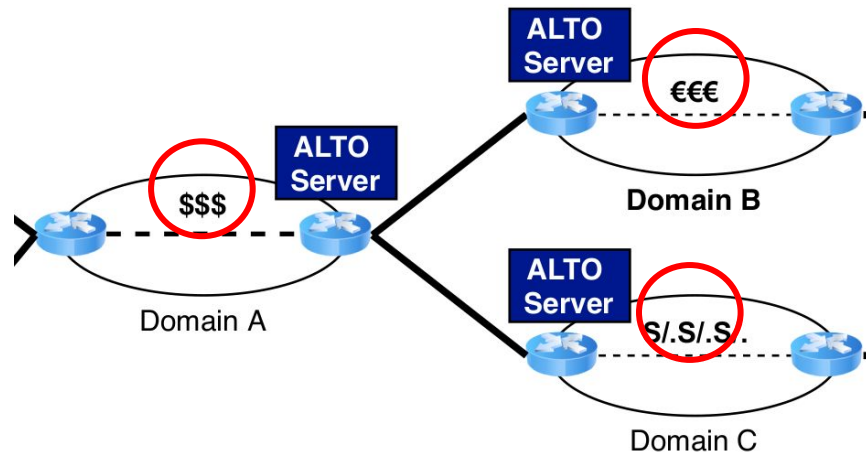
Single-domain composition

- Each domain can have its **own representation** of the same network information

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



Same utilization charge property but the form of
billing **may not be uniform**



Simple resource query language

- Applications need a query to express all **common resource requirements** to the network.
 - E.g., A 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 query interface** (e.g., filtered network/cost map).

Scalability & Privacy

Scalability

- The optimization problems, specified by the applications' requirements, can be **computationally expensive and time-consuming**.
 - The number of available paths for each flow increases exponentially with the number of domains involved.
 - The number of available configurations for a set of flows increase exponentially with both the network size and the number of flows.

Privacy & Security

- The information provided by the ALTO base protocol is considered coarse-grained in several recent multi-domain use cases.
- New ALTO extension services 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**.

How to design a whole ALTO framework?


Envisioned solutions & on-going efforts

Relationship: ALTO issues & solutions


FROM	TO	Related RFCs/Drafts/Papers
Server-to-Client ALTO communication	Server-to-Server ALTO communication	draft-dulinski-alto-inter-problem-statement draft-xiang-alto-multidomain-analytics draft-lachosrothenberg-alto-brokermdo draft-lachos-sfc-multi-domain-alto
Domain connectivity discovery	Multi-domain connectivity discovery	BGP (RFC4271) PCE (RFC5441, RFC6805) BGP-LS (RFC7752)
ALTO server discovery	Multi-domain ALTO server discovery	cross-domain ALTO server (RFC8686) PCE (RFC4674) draft-dong-pce-discovery-proto-bgp-07
Single-domain composition	Unified Resource Representation	draft-xiang-alto-unified-representation UNICORN MERCATOR
Simple resource query language	Generic/Flexible query language	Standard (GSM, NFV-NSD) Pre-standard (SOCKET-INTENTS, IBN)
Scalability	Computation complexity optimization	draft-gao-alto-routing-state-abstraction
Security & Privacy	Security/Privacy preserving	draft-xiang-alto-multidomain-analytics MERCATOR

Server-to-Server ALTO communication

- ALTO may consider either a **hierarchical or mesh** architectural deployment design [[INTER-ALTO](#)][[MD-ANALY](#)][[BROKER](#)][[SFC-MD](#)].
 - 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.



Hierarchical



Mesh

Multi-domain connectivity discovery

- Multi-domain mechanisms combining **domains sequence computation** and **paths computation** need to be defined.
- Standardized computation protocols could be leveraged:
 - **BGP-based** [[RFC4271](#)]: Provides multi-domain sequence computation (It does not advertise multiple alternative routes).
 - **BGP-LS-based** [[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.

- ☐ Intra-domain node
- ☒ Border Node




Multi-domain ALTO server discovery

- ALTO cross-domain server discovery [[RFC8686](#)]
 - It specifies a procedure for **identifying ALTO servers outside of the ALTO client's** own network domain.
- PCE Discovery [[RFC4674](#)]
 - It 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.
- BGP extension for PCE discovery [[PROTO-BGP](#)]
 - It 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.

Unified Resource Representation

- Aggregate and expose network information from multiple ALTO servers into a **single and consistent “virtual” network view**.
 - [\[UNI-REPRE\]](#)[\[UNICORN\]](#)[\[MERCATOR\]](#) use **mathematical programming constraints** for multi-domain composition.



$$Domain_A : \{x_1 + x_2 \leq 100\}$$


$$Domain_B : \{x_1 \leq 30\}$$

$$Domain_C : \{x_2 \leq 30\}$$

Linear Inequalities

Unified Resource Representation (2)

- Aggregate and expose network information from multiple ALTO servers into a **single and consistent “virtual” network view**.
 - [[UNI-REPRE](#)][[UNICORN](#)][[MERCATOR](#)]: Design options of multi-domain composition mechanisms using **mathematical programming constraints**.




$$Domain_A : \{x_1 + x_2 \leq 100\}$$

$$Domain_B : \{x_1 \leq 30\}$$

$$Domain_C : \{x_2 \leq 30\}$$

Linear Inequalities



~~$$Domain_A : \{x_1 + x_2 \leq 100\}$$~~


$$Domain_B : \{x_1 \leq 30\}$$

$$Domain_C : \{x_2 \leq 30\}$$

Remove redundancy


Unified Resource Representation (3)

- Aggregate and expose network information from multiple ALTO servers into a **single and consistent “virtual” network view**.
 - [\[UNI-REPRE\]](#)[\[UNICORN\]](#)[\[MERCATOR\]](#): Design options of multi-domain composition mechanisms using **mathematical programming constraints**.




$Domain_A : \{x_1 + x_2 \leq 100\}$
 $Domain_B : \{x_1 \leq 30\}$
 $Domain_C : \{x_2 \leq 30\}$

Linear Inequalities



~~$Domain_A : \{x_1 + x_2 \leq 100\}$~~
 $Domain_B : \{x_1 \leq 30\}$
 $Domain_C : \{x_2 \leq 30\}$

Remove redundancy



$\{x_1 \leq 30\},$
 $\{x_2 \leq 30\}$

Unified/Single Repres.

Generic/Flexible query language

- With a flexible/generic query language:
 - The network can **filter out a large number of unqualified** domains.
 - The network can **selectively send the resource information** of a small number of qualified domains.
- Language specification:
 - Inspired by standard [[GSM](#)], [[NFV-NSD](#)] or pre-standard [[SOCKET-INTENTS](#)][[IBN](#)] mechanisms, implemented with a user-friendly grammar (e.g., SQL-style query).

Generic/Flexible query language (2)

```
"sg":{  
  "nfs": [ "NF1", "NF2", "NF3" ],  
  "saps": [ "SAP1", "SAP2" ],  
  "sg_links":[  
    {"id": 2, "src-node": "SAP1", "dst-node": "NF1"},  
    {"id": 2, "src-node": "NF1", "dst-node": "NF2"},  
    {"id": 3, "src-node": "NF2", "dst-node": "NF3"},  
    {"id": 4, "src-node": "NF3", "dst-node": "SAP2"}],  
  "reqs": [  
    {"id": 1, "src-node": "SAP1", "dst-node": "SAP2",  
      "sg-path": [ 1, 2, 3, 4 ]}]  
}
```

[[BROKER](#)][[MD-E2E](#)]

```
f1: {src_ip = 10.0.0.1 and dst_ip = 10.0.0.2};  
f2: {src_ip = 10.0.0.1 and dst_ip = 10.0.0.3};  
flow_set: {f1, f2};  
req1: f1.bandwidth >= 30 Gbps;  
req2: f2.bandwidth >= 30 Gbps;  
SELECT bandwidth from flow_set  
WHERE req1 and req2;
```

[PED*]

Computation complexity optimization

- ALTO servers need to support mechanisms such as **pre-computation, projection and/or compression** 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**.
 - [[DRAFT-RSA](#)] describes equivalent transformation algorithms that identify/remove redundant information to obtain a more compact view.
 - [[MERCATOR](#)] proactively discovers network resource information for a set of flows, and project the pre-computed result to get the information when receiving actual requests from applications.

Security/Privacy preserving

- ALTO needs mechanisms (with little overhead) that **provide accurate sharing network information**, and at the same time **protects each member domain**.
- [[MD-ANALYTICS](#)][[MERCATOR](#)] present a privacy-preserving, multi-domain extension of ALTO.
 - ALTO servers in all member domains **use a secure multi-party computation (SMPC) protocol** to collectively send the responses to the ALTO client without revealing the source of any entry.

Conclusion

Summary

- ALTO emerges as a solution for exposing network information across multiple domains
- Key issues & solution mechanisms in the current ALTO framework were presented to support important multi-domain environments.

Next Steps

- Continue the discussions on feasibility and deployment concerns

Additional information

- ALTO Internal meetings
 - Wednesday weekly meetings (9:00 US ET)
 - Bridge: <https://yale.zoom.us/j/8423318713>
- SIGCOMM'20 Network-Application Integration/CoDesign Workshop (NAI'20)
 - Link: <https://conferences.sigcomm.org/sigcomm/2020/workshop-nai.html>
 - Workshop date: August 14, 2020

Thanks!

(more) Questions?



[https://www.linkedin.com/in/dannylachos/
dlachosp@dca.fee.unicamp.br](https://www.linkedin.com/in/dannylachos/dlachosp@dca.fee.unicamp.br)