

ALTO for Multi-Domain Applications: Use Cases and Design Requirements

draft-xiang-alto-multidomain-usecases-00
draft-xiang-alto-unified-representation-01

Qiao Xiang¹, Franck Le², Y. Richard Yang¹,

¹ Yale University,

² IBM Watson Research Center,

March 26, 2019, IETF 104 ALTO Meeting

Overview

- Introduce important multi-domain applications and how they can benefit from ALTO
- Describe a generic framework for multi-domain applications to use ALTO to improve the performance
- Discuss new design requirements for ALTO to better support these applications

ALTO RFC7285

- Provide network information (e.g., basic network location structure and preferences of network paths) with the goal of modifying network resource consumption patterns while maintaining or improving application performance.
 - Exposed information: abstract map of network.
- Example use case: P2P peer selection.

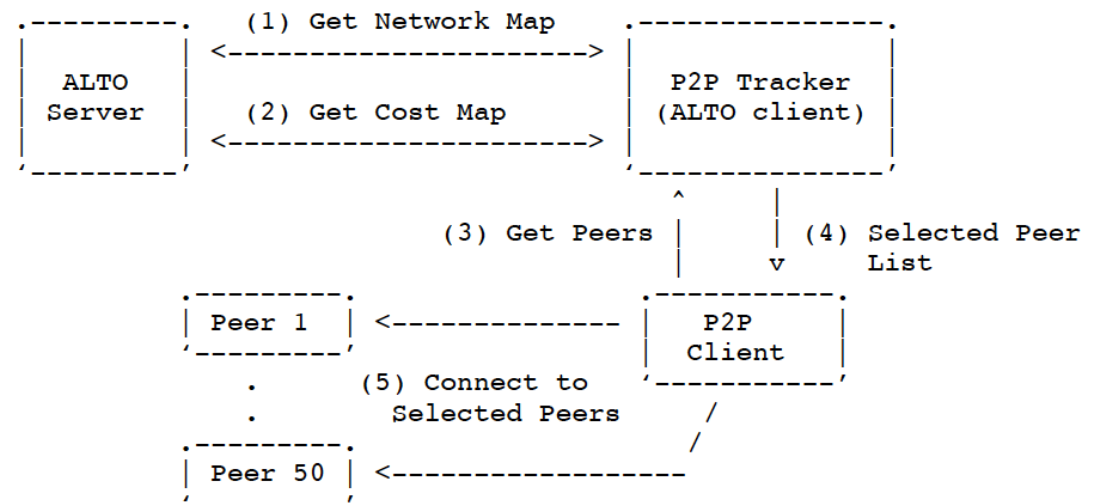


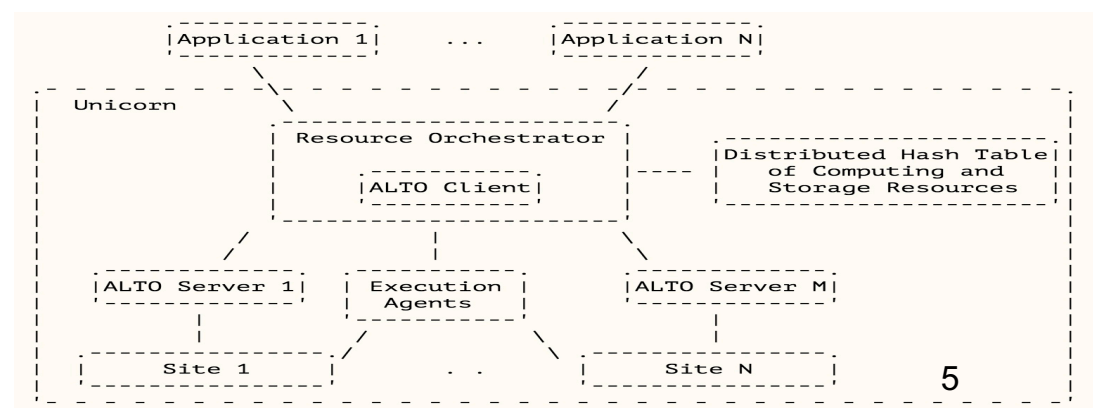
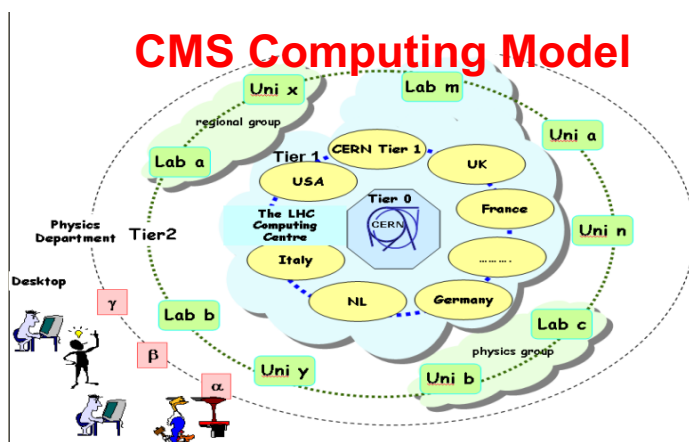
Figure 4: ALTO Client Embedded in P2P Tracker

ALTO in Multi-Domain

- New network architectures, e.g., SFC, edge clouds and SD-WAN, are deployed in multi-domain networks. As such, new use cases emerge, which can benefit from ALTO.
 - Multi-domain e2e network service deployment
 - draft-lachosrothenberg-alto-md-e2e-ns
 - draft-lachosrothenberg-alto-brokermdo
 - Multi-domain, collaborative data sciences
 - draft-xiang-alto-multidomain-analytics
 - draft-xiang-alto-exascale-network-optimization
 - Flexible interdomain routing control
 - ...

Multi-Domain, Collaborative Data Sciences

- **Settings:** Different organizations contribute various resources (e.g., sensing, computation, storage and networking resources) to collaboratively collect, share and analyze extremely large amounts of data.
 - Example: the CMS experiment in Large Hardon Collider.
 - A resource orchestration system (Mercator) was designed to orchestrate large-scale datasets between different CMS sites.
 - ALTO cost map / ECS services with path-vector extension are deployed for resource discovery.
 - Full demonstration at SC'18 to orchestrate the transmission of a set of scientific workflows from Dallas to Pasadena at 100 Gbps




Flexible Interdomain Routing Control

- **Settings:** An autonomous system (AS) providing such a service (the "provider") allows other ASes ("clients") to specify routing actions at the provider based on flexible matching conditions (e.g. , match on TCP/IP 5-tuple).
 - Use cases: DDoS and congestion mitigation, application-specific inbound traffic control, ...
 - Enabling flexible interdomain routing control requires the exposure of network's routing capability.
 - An ALTO server can be deployed to return **not a single cost (FIB), but a set of different costs (RIB)**.


Example 1: DDoS and Congestion Mitigation

- AS A is compromised and being used to send DDoS traffic to E.
- W/o flexible interdomain routing, E can setup a firewall locally, but normal traffic from B to E will still be congested at C-D-E.
- With flexible interdomain routing, E can specify such a firewall at C to simultaneously block DDoS traffic from A, and avoid congestion of normal traffic from B to E.



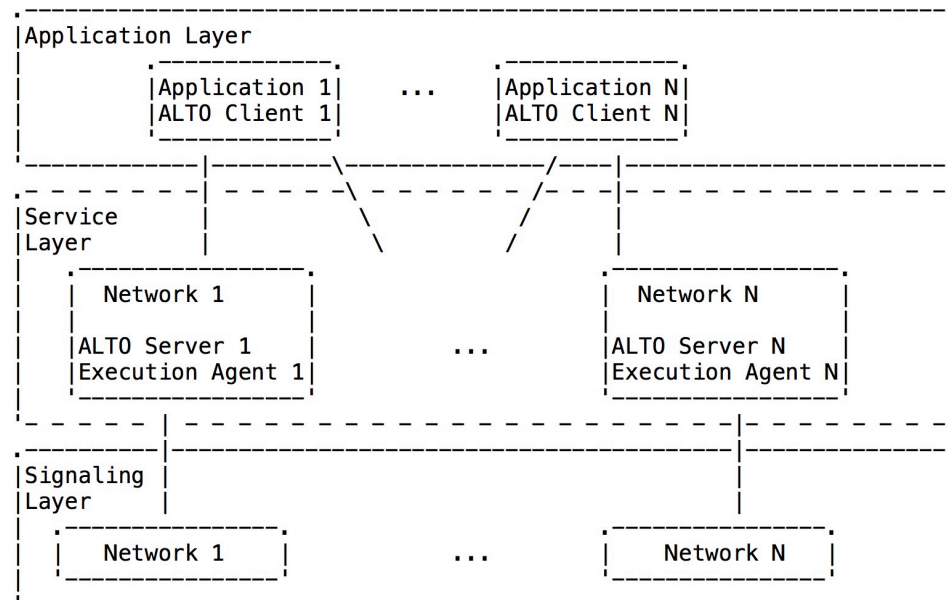
Example 2: Application-Specific Inbound Traffic Control

- AS E wants to receive HTTP traffic from A along route A-B-C-E, and non-HTTP traffic along route A-B-D-E.
- W/o flexible interdomain routing, E has limited control on routes for inbound traffic.
 - BGP selective announcement / AS-path prepending are only destination IP prefixed based, and are not scalable.
- With flexible interdomain routing, E can specify application-specific route selection at AS B.



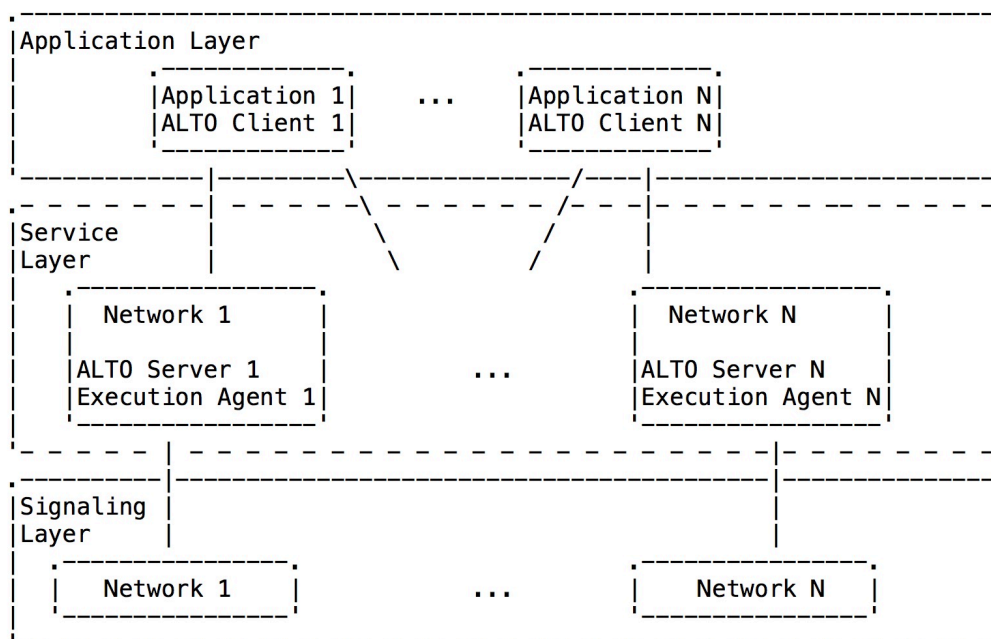
A Generic Framework for Using ALTO to Improve Multi-Domain Applications' Performance

- **Application layer:** each application deploys one or more **ALTO clients** to query information provided by networks.
- **Service layer:** each network deploys one more more **ALTO servers** to respond the queries sent by the ALTO clients from applications, and one or more **execution agents** to respond to the applications' resource consumption actions.
- **Signaling layer:** each network deploys **interdomain protocols / systems**, such as routing protocol BGP and resource reservation system OSCARS.



Basic Workflow

1. An application identifies the networks whose resources it may want to consume, and invokes its ALTO clients to query the ALTO servers;
2. Upon receiving a query from an ALTO client, an ALTO server checks its local information, contacts the underlying signaling layer protocol / system of its residing network if needed, and returns the latest resource information to the querying ALTO client;
3. The applications uses the resource information collected from ALTO servers to make resource allocation decisions (e.g., route selection, resource reservation, etc.), and send such decisions to corresponding execute agents.



Requirements for ALTO to Support New Multi-Domain Cases

- Expose accurate information while preserving network privacy
- Expose information of alternative resources
 - Provide users more flexibility choosing different resources, and provide networks with additional business opportunities (e.g., flexible interdomain routing)
- Provide interfaces for more flexible query
 - New networking architectures (e.g., SDN) support finer-grained management of network
 - New applications (e.g., service function chaining) have finer-grained requirement of network information (e.g., precedence, link-disjointness)
- Provide a unified, accurate representation of multiple types of resources
 - New multi-domain applications (e.g., data analytics) often consume multiple types of resources across multiple networks

Existing Efforts in WG

- Expose accurate, bottleneck information while preserving network privacy
 - draft-ietf-alto-path-vector
 - draft-gao-alto-routing-state-abstraction
 - draft-xiang-alto-exascale-network-optimization
- Expose information of alternative resources
 - draft-ietf-alto-cost-calendar
 - draft-randriamasy-alto-cost-context
 - draft-lachosrothenberg-alto-brokermdo
 - draft-lachosrothenberg-alto-md-e2e-ns
- Provide interfaces for more flexible query
 - draft-ietf-alto-unified-props
 - draft-gao-alto-fcs

Existing Efforts in WG

- Provide a unified, accurate representation of multiple types of resources
 - draft-xiang-alto-unified-representation
 - -00 version proposed at IETF 102
 - **Driving question:** Can we design a unified resource representation framework in ALTO to provides accurate, compact resource information to applications, who may have a wide range of requirements / objectives?
 - **Basic idea:** Use mathematical programming constraints to represent the capacity region for a set of flows.
 - **Missing piece:** how an ALTO client expresses the requirements on what information is needed?
 - Status: initial abstract language syntax specified

draft-xiang-alto-unified-representation: First Step

- Introduce a new cost type (cost mode: "array", cost metric: "variable-list") to allow the ALTO server to send a cost map:
(source, destination) pair -> a list of decision variables related to this pair
- Introduce a new entity domain "cstr" (short for constraint), and use a cstr property map to send the set of mathematical constraints.
- Lots of issues need to be addressed
 - Obviously, this design is generic, but it may be too generic ...
 - Security/privacy ...

```
"cost-type": {  
  "cost-mode": "array",  
  "cost-metric": "variable-list"  
},
```

```
"property-map": {  
  "cstr:001": { "bw-cstr": "[0][0] add [0][1] leq 100"},  
  "cstr:002": { "bw-cstr": "[0][0] eq [0][1]"},  
  "cstr:003": { "bw-cstr": "[1][0] add [0][0] leq 100"},  
  "cstr:004": { "srlg-cstr": "[0][2] intersect [1][1] eq {2, 3, 4}"},  
}
```

```
"cost-map": {  
  "PID1": {  
    "PID2": ["f1:b2:p1", "f2:bw:p2", "f1:srlg"],  
    "PID3": [ "f2:bw:p1", "f2:srlg" ]  
  }  
}
```

$f1:bw:p1 + f1:bw:p2 \leq 100$

$f1:bw:p1 = f1:bw:p2$

$f2:bw:p1 + f1:bw:p1 \leq 100$

$f1.srlg \cap f2.srlg = \{2, 3, 4\}$

Summary

- Multiple important multi-domain applications can benefit substantially from ALTO.
- Corresponding new design requirements also emerge
- Different drafts have been proposed to address some of the design requirements

Next Steps

- Solicit feedback and from WG members and industrial partners
- Systematic investigation of deployment concern of ALTO for multi-domain applications
 - Incentive, stability, scalability, privacy, etc.
- Systematic design of extensions to address corresponding design requirements