ALTO Extensions to Support Application and Network Resource Information Exchange for High Bandwidth Applications

<draft-lee-alto-app-net-info-exchange-00.txt>

ALTO Extensions for Collecting Data Center Resource Information

<draft-lee-alto-ext-dc-resource-00.txt>

Young Lee, Huawei
Greg Bernstein, Grotto Networking
Tae Sang Choi, ETRI
Sreekanth Madhavan, Huawei
Dhruv Dhody, Huawei
Luis M. Contreras, Telefonica

84th IETF – Vancouver DC, July 29 - August 3, 2012

Large Bandwidth Use Case

<draft-bernstein-alto-large-bandwidth-cases-01>

Case 1: End System Aggregation: Many clients using services offered at two or more "data centers"

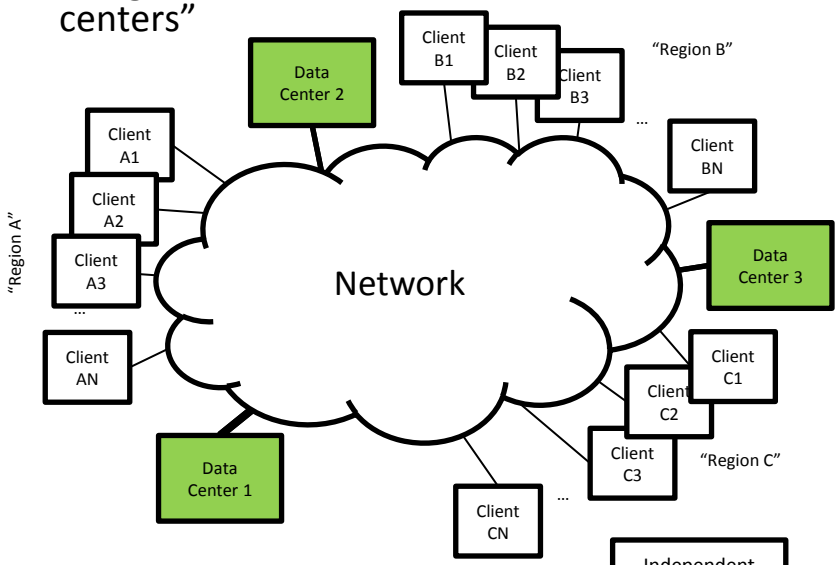

Case 2: Traffic engineered "express lanes" between data centers and end user regions

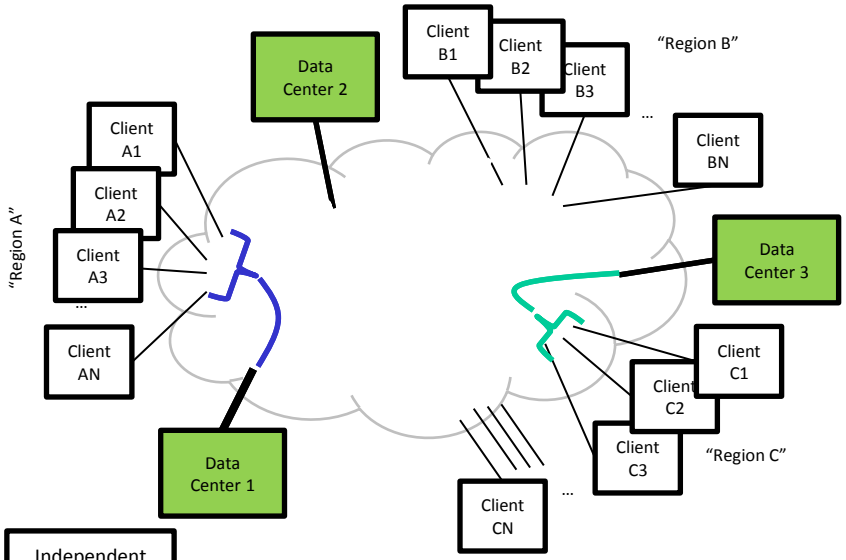

Case 3: Data Center to Data Center Communications: Application Overlays, Recovery

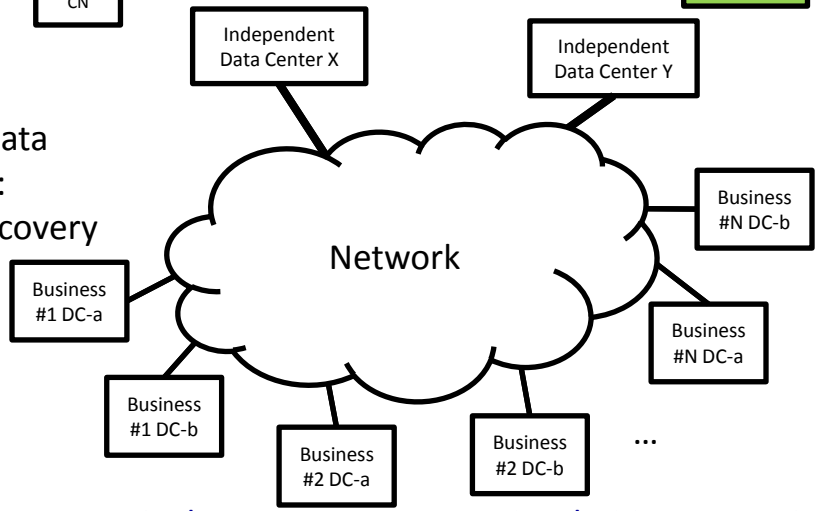

ALTO i2aex

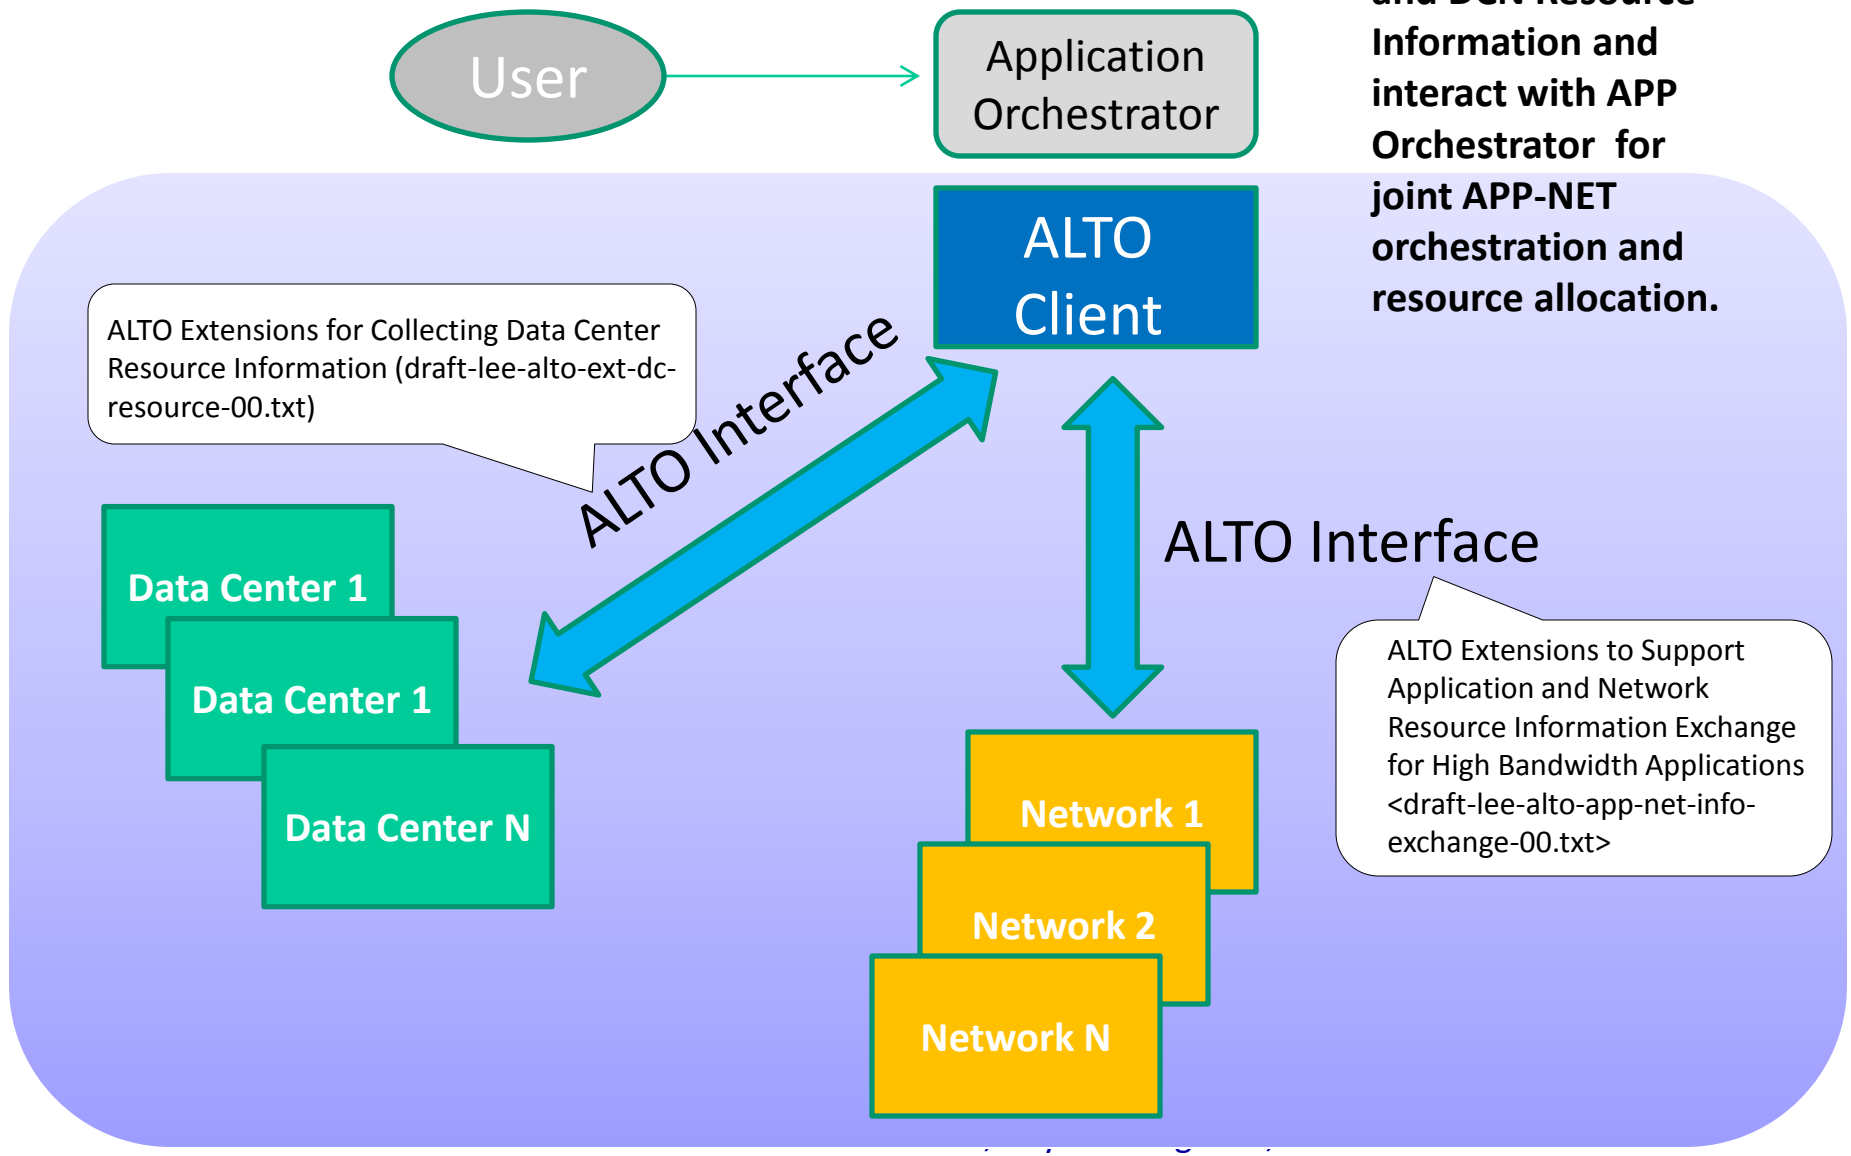

ALTO Client can collect Net Resource and DCN Resource Information and interact with APP Orchestrator for joint APP-NET orchestration and resource allocation.

ALTO Query Mechanism that filters the request constraints

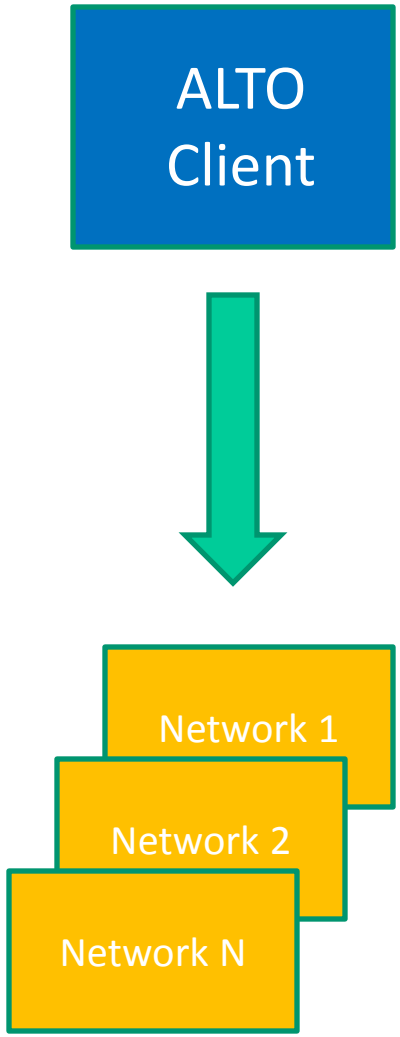

- End Point Source Address(es)
- End Point Destination Address(es)
- Cost Type :={summary, graph}
- Constraints /* For instance, constraints can be the minimum bandwidth, maximum latency, maximum hop counts, maximum packet loss, etc. */
- Parameters /* a set of result parameters that each result (summary or a link in graph) should have. For instance, latency, cost, etc.) */
- Objective-function: The summary or the graph should be computed based on optimizing which parameter – IGP cost, latency, residual bandwidth, etc.

ALTO Response Mechanism with Reduction of Data Sets

- The list of feasible Source-Destination pair and its Cost Type
- For each feasible S-D pair, indicate the following:
- Constraints Values /* indicate the actual values of each constraint requested */
- Administration Domain ID /* For each network administration domain, the domain ID needs to be conveyed */

Info Model and Encoding Suggestions

Alto query:

```
object {
    TypedEndpointAddr src;
    TypedEndpointAddr dsts<1..*>;
  } EndpointFilterExt;
  object {
    CostMode        cost-mode;
    CostType        cost-type;
    JSONString      constraints<0..*>;   [OPTIONAL]
    EndpointFilterExt  endpoints;
  } CsoReqEndpointCostMap;
```

Alto response:

```
object {
     JSONNumber hopcount;
     JSONNumber latency;
     JSONNumber pktloss;
} DstCostsConstraints;

object EndpointDstCosts {
    DstCostsConstraints[TypedEndpointAddr];     ...
  };
  object {
    EndpointDstCosts [TypedEndpointAddr]<0..*>;
    ...
  } EndpointCostMapData;
  object {
    CostMode          cost-mode;
    CostType          cost-type;
    EndpointCostMapData map;
  } CsoInfoResourceEndpointCostMap;
```

```
POST /endpointcost/lookup HTTP/1.1
 Host: alto.example.com
 Content-Length: [TODO]
 Content-Type: application/alto-csoendpointcostparams+json
 Accept: application/alto-csoendpointsummary+json,application/alto-error+json
 {
   "cost-mode" : “numerical",
   "cost-type" : “summary",
   "constraints": ["bw gt 20", "latency lt 10“, “hopcount lt 5”, “pktloss lt 0.03”],
   "endpoints" : {
     "srcs": [ "ipv4:192.0.2.2" ],
     "dsts": [  "ipv4:192.0.2.89",  "ipv4:198.51.100.34",  "ipv4:203.0.113.45"  ]
   }
 }

HTTP/1.1 200 OK
Content-Length: [TODO]
Content-Type: application/alto-csoendpointsummary+json
 {
   "meta" : {},
   "data" : {
     "cost-mode" : “numerical",
     "cost-type" : “summary",
     "map" : {
       "ipv4:192.0.2.2": {
       "ipv4:192.0.2.89"    : [ “latency eq 5”,  “hopcount eq 8”, “pktloss eq 0.01” ],
       "ipv4:18.51.100.34"  : [ “latency eq 9”,  “hopcount eq 10”, “pktloss eq 0.02” ],
       "ipv4:203.0.113.45"  : [ “latency eq 40”, “hopcount eq 12”, “pktloss eq 0.02” ]
       }
     }
   }
 }
```

ALTO Collection of DC Resource Information

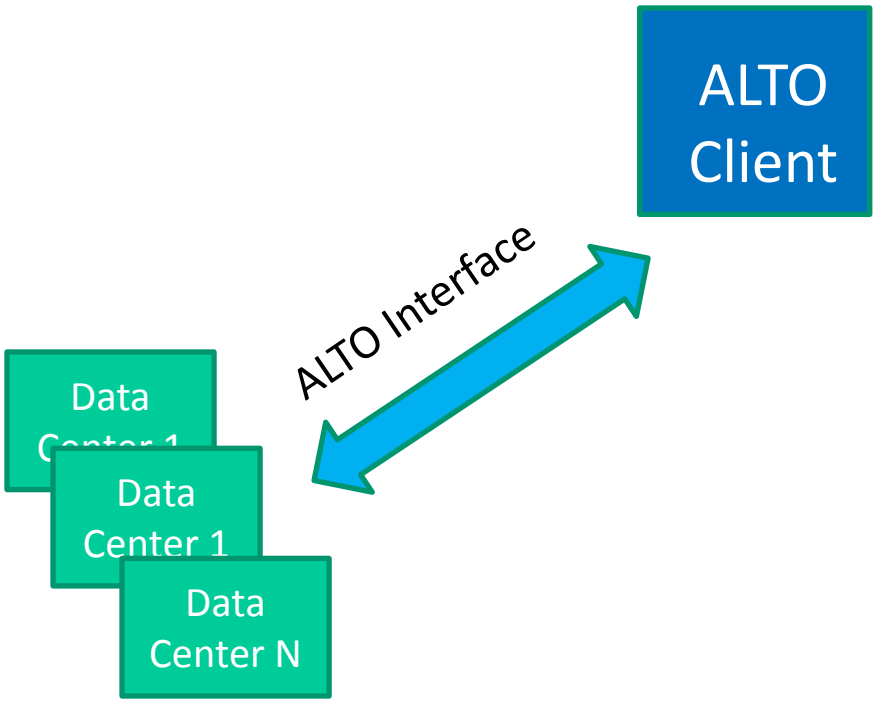

- Data Center Identifier (DCI)
- Data Center Location Identifier (e.g., IP address of the gateway node)
- Time Stamp
- Abstracted Memory Usage
- Abstracted Available Memory
- Abstracted CPU usage
- Abstracted Available CPU
- Abstracted Power Consumption Level & Cost
- DC Network Cost Models (Reserve, On-demand, Spot)
- DC Network Cost
- DC Network Resource Constraints

Encoding Suggestions

Pull based Query:

Based on GET URL /getdcinfo

Pull based response:

object

{

VersionTag vtag; [OPTIONAL]

TypedEndpointAddrall addr;

JSONNumber srvload;

JSONNUmber ramusage;

...

} InfoDCProperty**;**

GET /dcinfo HTTP/1.1

Host: alto.example.com

Accept: application/alto-dcinfo+json

HTTP/1.1 200 OK

Content-Length: [TODO]

Content-Type: application/alto-dcinfo+json

{

"meta" : {},

"data" : {

"vtag" : "1266506139",,

addr : “ipv4: 10.18.51.151:5060”,

srvload: 25,

ramusage: 60

}

}

Push based Query:

object

{

VersionTag vtag; [OPTIONAL]

TypedEndpointAddrall addr;

JSONNumber srvload;

JSONNUmber ramusage;

...

} InfoDCProperty**;**

Push based response:

200 OK with body NUl

POST /dcinfo HTTP/1.1

Host: alto.example.com

Content-Length: [TODO]

Content-Type: application/alto-dcinfo+json

{

"meta" : {},

"data" : {

"vtag" : "1266506139",

addr : “ipv4: 10.18.51.151:5060”,

srvload: 25,

ramusage: 60

}

}

HTTP/1.1 200 OK

Host: alto.example.com

Multi-Domain Issues

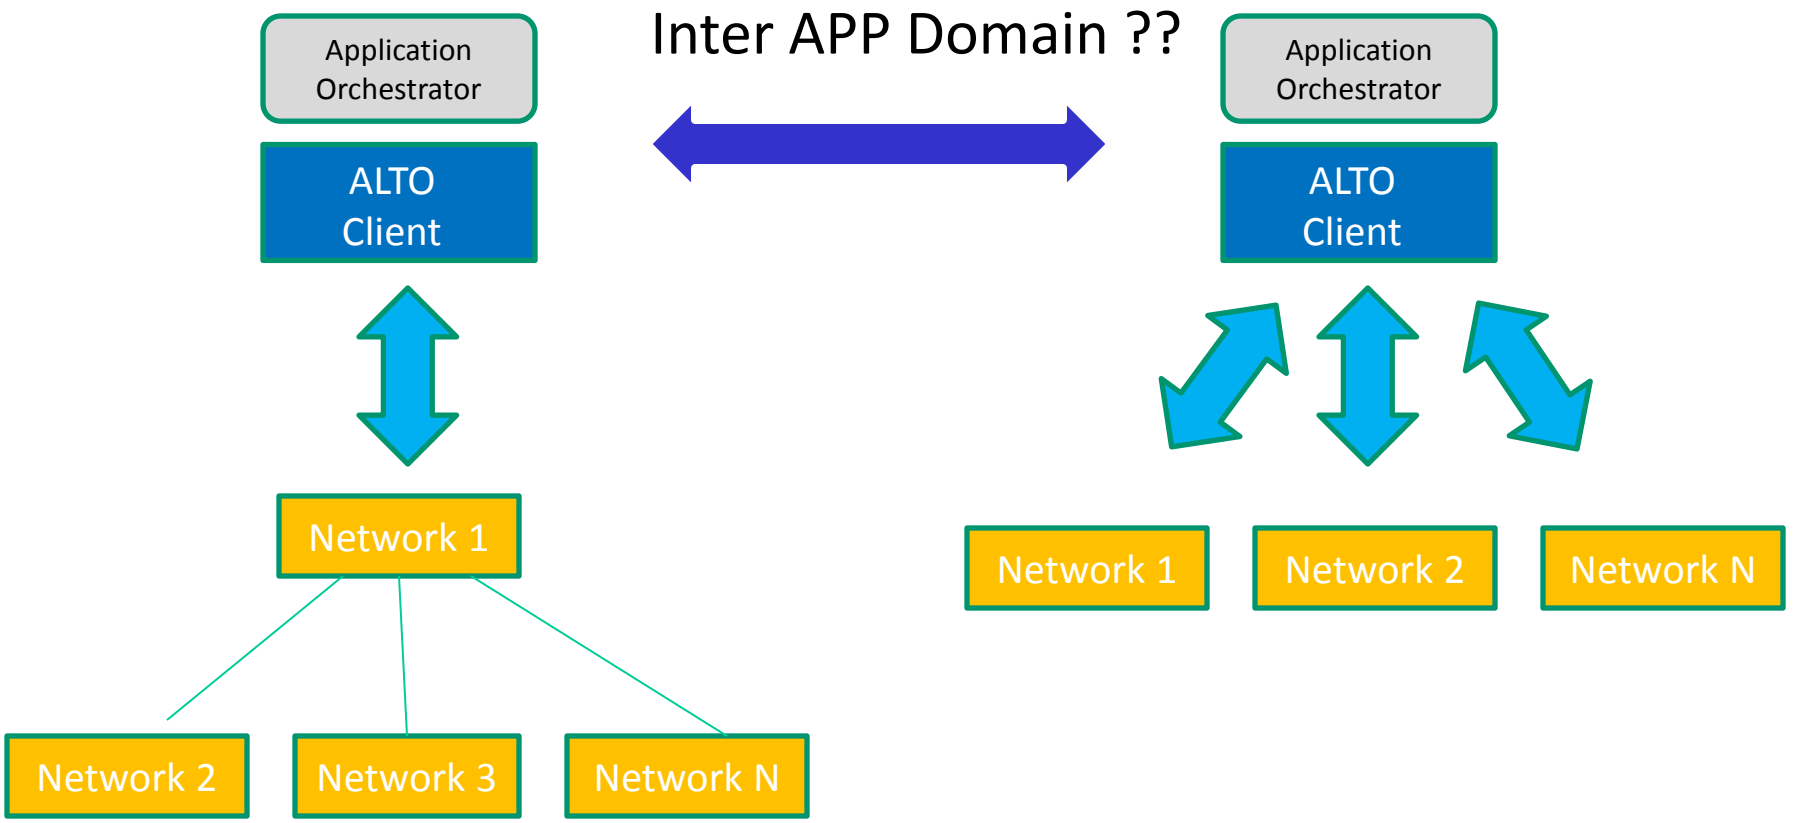

Hierarchical Network Inter-domain

Peer Network Inter-domain

Summary & Next Steps

- We presented two ALTO extension drafts that would help joint APP-NET orchestration and resource allocation for large bandwidth use case in the data center environments.
- Future work would include multi-domain issues.
- Greg's next presentation "Bandwidth Constraint Representation" to reduce amount of information shared while promoting optimization
  - Abstract paths with abstract shared bottlenecks
  - Abstract cost-constraint graphs