

Side meeting on network and application integration

Introduction and agenda

Börje Ohlman, Sabine Randriamasy, Richard Yang

IETF-106 Singapore – Wed. Nov 20th 2019

Context and purpose

- Look at application and network integration, including both
 - application-aware networking (AAN) and
 - network-aware applications (NAA)
- Identify the major missing pieces (information, mechanisms)
 - to allow better network and application integration
 - based on ALTO but also beyond
- Given today's network usage trends
 - Identify problems that would be solved with the help of such a cooperation

Objectives

- Gather input, helping to better identify how ALTO evolutions can complement with other IETF protocols or research activities.
- In particular, insights, experience and expectations from participants involved in network operations, applications, or cooperation thereof
- Feedbacks that this meeting would like to gather include
 - AAN or NAA: what is it?
 - What problems did we address?
 - What did we try? What worked? What did not? What is missing?
 - What are the changes in networks, applications, in research, technologies that we need to take into account?
- A list of key questions has been proposed and may be updated

IETF ALTO WG - <https://tools.ietf.org/wg/alto/>

- Defines protocols for network-awareness in applications
 - Network topology abstraction is exposed to applications by providers
- Addresses unnecessary network costs or performance loss due to
 - network topology agnostic selection of application endpoints such as peers or CDN surrogates, [RFC7285] []
 - poor information on end-to-end network paths [id-aperf] [RFC8189] [id-pvect] [id-sse]
 - single routing cost metric | no intermediate insight | outdated | too voluminous
 - bad timing for connection, [id-acal]
- Addresses need to expose properties on relevant path steps [id-uprop]
- What misses is extensions to support applications spanning multiple network domains, network technologies, multiple paths, ...
- What else ?

ALTO exploration w.r.t network trends - examples

- Network usage trends currently considered for future ALTO work include
 - Prevalence of wireless and cellular networks,
 - 5G multi-technology e2e topologies,
 - Multi-domain topologies
 - Orchestration in single and multiple domains
- Multi-domain topologies & orchestration
 - <https://dl.acm.org/citation.cfm?id=3341126>
 - <https://tools.ietf.org/html/draft-lachosrothenberg-alto-brokermdo-02>
 - <https://tools.ietf.org/id/draft-xiang-alto-multidomain-analytics-02.txt>
- ALTO support to deploy applications and functions in edge computing context
 - <https://tools.ietf.org/id/draft-contreras-alto-service-edge-00.txt>
- ALTO for cellular networks
 - <https://tools.ietf.org/html/draft-bertz-alto-mobilitynets-00>
 - <https://tools.ietf.org/html/draft-randriamasy-alto-cellular-adresses-02>
- ALTO context-dependent costs for multi-access or highly dynamic networks
 - <https://tools.ietf.org/id/draft-randriamasy-alto-cost-context-02.txt>

Agenda

- **Session 1: 8h30 - 9h45 in Room VIP A**
 - **08:30-08:40** Meeting context - Brief ALTO introduction
 - **08:40-08:55** Speaker1: Yunfei Zhang (Tencent)
 - **08:55-09:10** Speaker2: Luis Contreras (Telefonica)
 - **09:10-09:25** Speaker3: Zhenbin Li (Huawei)
 - **09:25-08:40** Speaker4: Richard Yang (Yale University)
 - **09:40-09:45** Conclusion
- **Session 2: 10h - 11h in Room Bras Basah**
 - Open discussion

Key questions – 1/2

- What is Application Aware Networking (AAN)?
 - What properties of application should network be aware of, RSpec, Dspec, ..., connectivity, diffserv, connection request specifying intent? What is the relevant time scale, long term, per flow aware, per packet?
 - What has been tried? What worked? What did not work in your view? Why?
 - What changes or potential changes may cause more or less application -aware networking?
 - What problems do we need to solve now?
 - If you were to design today, what AAN features/mechanisms would you do (context-dependent, for IoT, for 5G)?

Key questions – 2/2

- What is Network Aware Application (NAA)?
 - What properties of networks should applications be aware of?
 - What have been tried? What worked? What did not work in your view? Why?
 - What changes or potential changes may cause more or less network- aware application?
 - What problems do we need to solve now?
 - If you were to design today, what NAA features/mechanisms would you do?
- How to harmonize decision making by both networks and applications?
 - How do you think the decision process should be harmonized?

References

- ALTO Status Pages
 - <https://tools.ietf.org/wg/alto/>
- [RFC 7285] - Application-Layer Traffic Optimization (ALTO) Protocol
 - <https://tools.ietf.org/html/rfc7285>
- [id-pvect] « ALTO Extension: Path Vector »
 - <https://tools.ietf.org/wg/alto/draft-ietf-alto-path-vector/>
- [RFC 8189] « Multi-Cost ALTO »
 - <https://tools.ietf.org/html/rfc8189>
- [id-aperf] « ALTO Performance Cost Metrics »
 - <https://tools.ietf.org/wg/alto/draft-ietf-alto-performance-metrics/>
- [id-acal] « ALTO Cost Calendar »
 - <https://tools.ietf.org/wg/alto/draft-ietf-alto-cost-calendar/>
- [id-sse] « ALTO Incremental Updates Using Server-Sent Events (SSE)»
 - <https://tools.ietf.org/wg/alto/draft-ietf-alto-incr-update-sse/>
- [id-uprop] “Unified Properties for the ALTO Protocol”
 - <https://tools.ietf.org/wg/alto/draft-ietf-alto-unified-props-new/>

Back-up

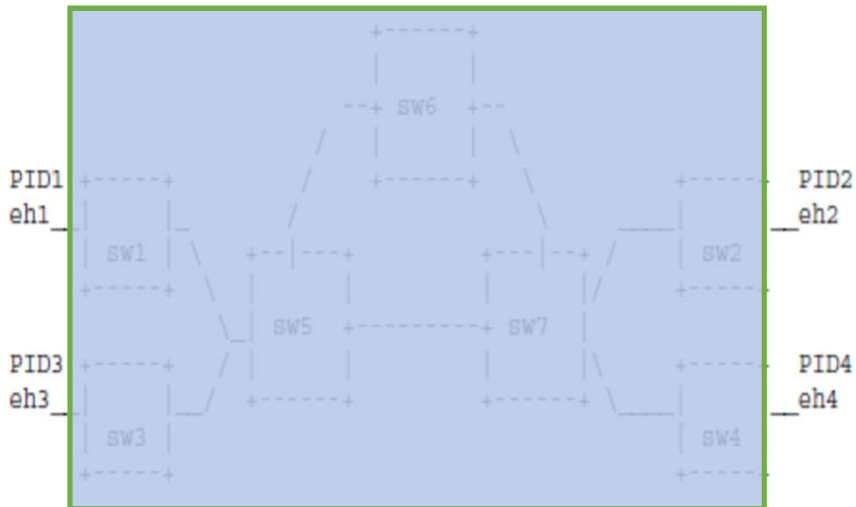
ALTO highlights

The IETF ALTO protocol

- Exposes abstracted operator-centric network view to applications and end hosts
- Goal: guide applications for a topology-aware selection among several endpoints
 - Trading operator cost efficiency with equal or better application performance
- To this end, ALTO offers RESTful APIs to convey provider-defined
 - **ALTO network map**: set of network location groupings with Provider-Defined Identifier (PIDs) and enumerated endpoints in each group.
 - **PID**: indirect and network agnostic manner to aggregate network endpoints that share some characteristic: subnet, POP, autonomous system, central office, ...
 - **ALTO cost map**: pairwise e2e path costs amongst sets of source and destination PIDs or endpoints.
- ALTO thus hides complexity and confidentiality
 - Network can protect confidential network state information, by abstracting real metric values into non-real numerical scores or ordinal ranking
- ALTO information assumed not available to 3rd parties by other means
- Requires mutual trust between operator and applications

ALTO – base protocol – RFC 7285

- « single node » topology abstraction
- 1 path per destination
- Specifies and conveys 1 single « routingcost » metric between src and dest



```
GET /costmap/num/routingcost HTTP/1.1
```

```
...
```

```
HTTP/1.1 200 OK
```

```
...
```

```
{  
  "meta" : {  
    ...  
    "cost-type" :  
      {"cost-mode": "numerical",  
       "cost-metric": "routingcost"},  
  }  
  "cost-map" : {  
    "PID1": {"PID1": 1, "PID2": 5, "PID3": 8 "PID4": 6},  
    "PID2": {"PID1": 5, "PID2": 1, "PID3": 1, "PID4": 8},  
    "PID3": {"PID1": 5, "PID3": 8, "PID3": 1},  
    "PID4": {"PID1": 6, "PID2": 10, "PID3": 1},  
  }  
}
```

PID = City, Region, any name, ...


Main ALTO extensions to enrich connection decision

- Base protocol: [RFC 7285] = **WHERE to connect**
- The ALTO WG specifies protocol extensions for deeper insight in paths
- **HOW to connect** given $N \geq 1$ paths / destination
 - **ALTO Path Vector Extension** = multi-switched, multi-step path [id-pvect]
 - Exposes abstraction of some intermediate steps of available paths
 - **Multi-Cost ALTO** [RFC 8189]
 - Exposes costs IF path is feasible : w.r.t. constraints on cost values – path filtering
- **ALTO Performance Cost Metrics** [id-aperf]
 - Abstracts estimated/SLA network costs (delay, jitter, packet loss, hop count, bandwidth..)
- **WHEN to connect** = ALTO Cost Calendars [id-acal]
 - 1 or N destinations, one path for each

HOW to connect - ALTO path vector

Provides abstracted details on paths

- Abstracted Network Elements (ANE)
 - Set of $N \geq 1$ switches, links, networks, ...
 - ANE properties may be exposed in a separate « ANE property map »




HTTP/1.1 200 OK

```
...
{
  "meta": {
    "dependent-vtags": [...],
    "multi-cost-types": [
      {"cost-mode": "array", "cost-metric": "ane-path"},
      {"cost-mode": "numerical", "cost-metric": "BWcapa"}
    ]
    "vtag": { //information to get ANE properties}
  },
  "cost-map": {
    "PID1": {
      "PID2": [{"ane": "L15", "ane": "L56", "ane": "L67", "ane": "L72"}, 100]},
    "PID3": {
      "PID4": [{"ane": "L35", "ane": "L57", "ane": "L74"}, 100]   }}}}
```

The application thus knows whether flows share bottleneck and how much total capacity they get

HOW to connect – Multi-Cost ALTO with constraints



Client can request « routingcost » + « BW capacity »
On paths with EITHER lower cost and bandwidth
OR higher cost and bandwidth

ALTO response provided ONLY on paths meeting the constraints

POST /costmap/filtered HTTP/1.1

```
...
{
  "multi-cost-types" : [
    "cost-mode": "numerical", "cost-metric": "routingcost",
    "cost-mode": "numerical", "cost-metric": "BWcapa"
  ],
  "or-constraints" : [ ["[0] le 30", "[1] le 70" ],
                      ["[0] le 50", "[1] gt 80"] ],
  "pids" : { "srcs" : [ "PID1" ],
            "dsts" : [ "PID2", "PID3" ] }
}
```

```
HTTP/1.1 200 OK
Content-Type: application/alto-costmap+json
{
  "meta" : {
    "dependent-vtags" : [...],
    "multi-cost-types" : [ ... ]
  }
  "cost-map" : {
    "PID1" : {
      "PID2" : [40, 90]
    }
  }
}
```

WHEN to connect – ALTO Cost Calendars

```
HTTP/1.1 200 OK
Content-Type: application/alto-costmap+json
Content-Length: ###
{
"meta" : {...
  "cost-type" : {"cost-mode": "numerical", "cost-metric": « BWcapa"},

  "calendar-response-attributes" : {
    "calendar-start-time" : Wed, 18 Oct 2017 00:00:00 GMT,
    "time-interval-size" : "4 hour",
    "numb-intervals" : 6 }
  } // end meta

"cost-map" : {
  "PID1" : {
    "PID2" : [90, 70, 80, 90],
    "PID4" : [90, 80, 60, 80]
  } }
}
```

- Array of time-dependent cost values
- Attributes specifying how to understand them