

Feedback on ALTO lightning talk at PANRG – IETF100

“The IETF ALTO protocol and its extensions”

Helping end hosts and applications to have a topology-aware insight into
end to end paths

Sabine Randriamasy

IETF ALTO interim meeting – December 18th 2017

PANRG research and objectives

- Goal of PANRG
 - extend path awareness from the control plane to the edge with mechanisms allowing:
 - Endpoint discovery of paths,
 - Explicit association of properties to paths by endpoints,
 - Explicit endpoint selection of paths.
- The identified underlying hard problems are mainly:
 - Conflict of interest between hosts and networks,
 - Path property exposition by the network to third parties considering
 - network confidentiality and information aggregation level,
 - diversity of timescale for configuration, discovery, selection and dissemination.
- The PANRG would like to leverage its work on existing IETF protocols and their potential extensions.

Lightning talk: ALTO extensions relevant to PANRG

- ALTO presented as helping to address following PANRG questions
 - How are path properties defined and represented?
 - How do endpoints get access to trustworthy path properties?
 - How can endpoints select paths to use for traffic in a way that can be trusted by the network?
- See <https://datatracker.ietf.org/meeting/100/materials/slides-100-panrg-04-the-ietf-alto-protocol-and-its-extensions/>

Relevance of ALTO identified at PANRG

- Brian Trammel one of PANRG chairs presented a draft on open questions to be addressed at PANRG
 - See <https://datatracker.ietf.org/meeting/100/materials/slides-100-panrg-06-open-questions/>
- Among questions
 - *"how do endpoints discover trustworthy path properties?"*
- Presentation mentions that
 - *SCION (a path-aware Internet architecture presented at IETF99) and ALTO represent points in this design space.*

ALTO-PANRG discussions desirable

- To see how ALTO can serve PANRG goals.
- Given that PANRG focuses on the mechanisms leveraging on protocols rather than on protocols themselves,
- To see what use cases PANRG finds relevant to ALTO
- What particular ALTO feature would be existing or needed
- What existing and future protocol extensions may help for those use cases

Back-up slides

- ALTO lighting talk at PANRG session – IETF 100

Open questions in PANRG & challenges


- See https://tools.ietf.org/html/draft-trammell-panrg-questions-01_11/07
- How are path properties defined and represented?
- How do endpoints get access to trustworthy path properties?
- How can endpoints select paths to use for traffic in a way that can be trusted by the network?
- How can interfaces to the transport and application layers support the use of path awareness?
- How can a path aware network in a path aware internetwork be effectively operated, given control inputs from the network administrator as well as from the endpoints?
- Some challenges gathered from list discussions
 - Different priorities for hosts and networks
 - Privacy concerns on both sides
 - “pseudo” instead of “real” path information more appropriate for end points and more realistic from networks’ controller
 - What is the appropriate level of exposure?

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, ...


IETF ALTO WG extensions relating to PANRG

- Base protocol: [RFC 7285] = **WHERE to connect**
- The ALTO WG specifies protocol extensions for deeper insight in paths
- Relation to PANRG: **HOW to connect** given $N \geq 1$ paths / destination
 - **ALTO Path Vector Cost Mode** = multi-switch path [id-pvect]
 - Exposes abstraction of some intermediate steps of available paths
 - **ALTO Contextual Cost Values**: several costs per metric for one dest [pd-acont]
 - Exposes costs given qualitative parameters such as « access type » or others
 - **Multi-Cost ALTO** [RFC 8189]
 - Exposes costs IF path is feasible : w.r.t. constraints on cost values – path filtering
- **ALTO Network Performance Cost Metrics** [id-aperf]
 - Abstraction of network delay, jitter, packet loss, hop count, and bandwidth
- **WHEN to connect** = ALTO Cost Calendars [id-acal]
 - 1 or N destinations, one path for each

HOW to connect - ALTO path vector- in progress

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 »

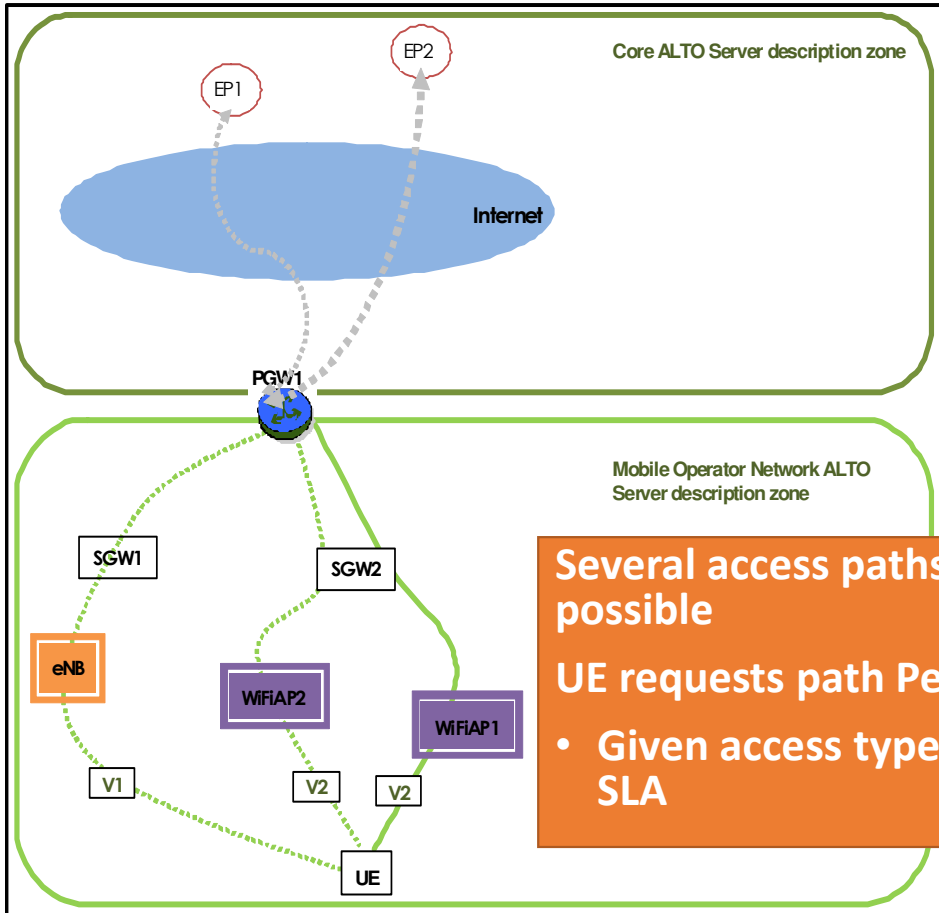


HTTP/1.1 200 OK

```
...
{
  "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 - ALTO Cost Context – in progress




Several access paths possible
UE requests path Perf
• Given access type and SLA

```
HTTP/1.1 200 OK
Content-Length: [TODO]
Content-Type: application/alto-endpointcost+json
{
  "meta" : {
    "cost-type" : {"cost-mode" : "numerical",
                  "cost-metric" : "BWscore"},
    "context-params" : [ ["cell", "wifi"], ["SLA-3"] ]
  }
  "endpoint-cost-map" : {
    "ipv4:192.0.2.2" : {
      "ipv4:192.0.2.89" : [10, 5],
      "ipv6:2000::1:2345:6789:abcd" : [4, 6]
    }
  }
}
```

Array of 2 context-based values:
[cell AND sla3, wifi AND sla3]

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 – in progress

```
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

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 Cost Mode »
 - <https://tools.ietf.org/html/draft-ietf-alto-path-vector-00>
- [pd-acont] « ALTO Contextual Cost Values »
 - <https://tools.ietf.org/id/draft-randriamasy-alto-cost-context-02.txt>
- [RFC 8189] « Multi-Cost ALTO »
 - <https://tools.ietf.org/html/rfc8189>
- [id-aperf] « ALTO Performance Cost Metrics »
 - <https://tools.ietf.org/html/draft-ietf-alto-performance-metrics-02>
- [id-acal] « ALTO Cost Calendar »
 - <https://tools.ietf.org/html/draft-ietf-alto-cost-calendar-02>