draft-ietf-alto-multi-cost-01.txt

Updates since IETF93

S. Randriamasy

W. Roome

N. Schwan

Multi-Cost ALTO in a nutshell

- Returns ***array*** of costs instead of *scalar* cost
- Defines 'OR‘ constraints,
  - Supports **decision trade-offs** such as:
  - *"give me costs among {those PIDs/Endpoints} with either moderate ‘routingcost’ or ‘hopcount’ equal to 0*
    - For example: 'hopcount' = 0 **OR** routingcost in [5, 10]"
- Proposes additional abstract cost metrics
- Applicable service information resources:
  - Filtered Cost Map (FCM),
    - **For full Multi-Cost Map: use empty SRC & DEST**
  - Endpoint Cost Service (ECS)
- Does not introduce new media types
- Backwards compatible with legacy ALTO Clients

WG feedback on v0

- Proposal for Introduction
- Request for clarification
  - Why full Multi-Cost maps only provided as Filtered Cost Maps
  - Difference between « testable-cost-types » and « multi-cost-types » in capabilities and constraints
- Nits and wording

Updates in v01- 1/2

- Section 3.5 Full Cost Map Resources
  - Augmented § 1: explain how a legacy client
    - would not understand Server response having « meta » with array of cost-types and
    - thus would not understand the mapping of cost values in array with cost-types.
- Section 4.1.1 Accept input parameters (to FCM)
  - `testable-cost-types`: appended text to explain how this features supports
    - value requests for cost-type T1 with constraints on cost-type T2 while client does not want values on T2.
    - Servers providing values on T1 and T2 with constraints on T2 only
  - `constraints`:
    - corrected nits and errors

Updates in v01 – 2/2

- § 9.2 Informative References
  - Removed references not used in draft
- Proposal for introduction
  - Text on motivation present in Introduction
  - Text on design choices already present in Section 3 Overview of approach
    - 3.2: Compatibility with legacy clients
    - 3.3: Filtered Multi Cost Map resources
    - 3.5: Full Cost Map resources

  ➔Need to consider newToC or add condensed text on design choices in Introduction

Next steps

- Consider updates of Section 1.Introduction
- Thank you to Richard Yang and Wang Xin
- Get their feedback on updates
- Get last WG feedback and Prepare for WGLC

Thank you

Back-up follows

Example § 5.1: Filtered multi-cost map resource in IRD

```
"filtered-multicost-map" : {
  "uri" : "http://alto.example.com/multi/costmap/filtered",
   "media-types" : ["application/alto-costmap+json" ],
   "accepts" : ["application/alto-costmapfilter+json" ],
   "uses" : [ "my-default-network-map" ],
   "capabilities" : {
      "cost-constraints" : true,
      "max-cost-types" : 2,
      "cost-type-names" : [ "num-routingcost",
                            "num-hopcount"],
      "testable-cost-type-names" : [ "num-routingcost",
                                     "num-hopcount"]
    }
},
```

Indicates that this service is MC compatible

Multi-Cost ALTO clients « see » also fields in slanted blue

Base ALTO clients « see » fields in black and ignore others

Example § 5.2: full MC Map - with testable cost types-1

```
POST multi/costmap/filtered HTTP/1.1
Host: alto.example.com
Content-Type: application/alto-costmapfilter+json
Accept: application/alto-costmap+json,application/alto-error+json
{
     "multi-cost-types" : [
       {"cost-mode": "numerical", "cost-metric": "routingcost"},
       {"cost-mode": "numerical", "cost-metric": "hopcount"}
    ],
     "testable-cost-types" : [
       {"cost-mode": "numerical", "cost-metric": "routingcost"},
       {"cost-mode": "numerical", "cost-metric": "hopcount"}
    ],
     "or-constraints": [
           ["[0] le 10", "[1] le 2"],
           ["[0] le 3",  "[1] le 6"]
     ],
    "pids" : {
       "srcs" : [ ],
       "dsts" : [ ]
    }
}
```

Motivation – use cases

- Use multiple selection metrics for endpoints and e2e paths
  - To jointly meet application needs while keeping network awareness
    - E.g. by ***jointly*** getting 'routingcost' meeting NP interests and 'bandwidth score' meeting app interests
- Save time and bandwidth on ALTO requests
  - 1 Multi-Cost transaction on N metrics rather than N on 1 metric
  - 1 Multi-Cost Map is smaller than N Cost Maps
- Consistency of metric values
  - Different cost-types may change at different paces
  - For multi-variate optimization
- Enrich filtering constraints to represent compromises, e.g.
  - *select paths with moderate 'routingcost' OR null 'hopcount'*

Multi-Cost transactions

- Multi-Cost Requests and responses convey an *Array* of costs
  - Array may contain any Cost Mode combination
    - Requested Cost-types array

      `["num-routingcost", "ord-hopcount", "string-status"]`
    - Taking values:

      `[23, 6, "medium"]`
  - **RULE:** cost values for each Source/Destination pair MUST be provided in the same order as in the array of Multi-Cost Types

Design

- Suggested new properties and costs
  - Aggregate values with or without units
    - EP-Nominal Memory, EP-Nominal Bandwidth
    - EP Occupied memory, EP Occupied bandwidth,
    - Path Occupation Cost, // or Bandwidth Score,

- Multi-Cost filtering constraints
  - Combine AND and OR operators
  - Are applied to cost-types present in value request
    - NOTE: [draft-lee-alto-app-net-info-exchange] proposes to use constraints on metrics not present in value request