Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 10, 2015

S. Randriamasy, Ed.
Alcatel-Lucent Bell Labs
R. Yang
Yale University
Q. Wu
Huawei
L. Deng
China Mobile
N. Schwan
Thales Deutschland

March 9, 2015

ALTO Cost Calendar
draft-randriamasy-alto-cost-calendar-03

Abstract

The goal of Application-Layer Traffic Optimization (ALTO) is to bridge the gap between network and applications by provisioning network related information in order to allow applications to make informed decisions. The present draft proposes to extend the cost information provided by the ALTO protocol. The purpose is to broaden the decision possibilities of applications to not only decide 'where' to connect to, but also 'when'. This is useful to applications that have a degree of freedom on when to schedule data transfers, such as non- instantaneous data replication between data centers or service provisioning to end systems with irregular connectivity. ALTO guidance to schedule application traffic can also efficiently help for load balancing and resources efficiency.

The draft proposes a new cost Mode called "Calendar" Mode, that is applicable to time-sensitive ALTO metrics and allows Applications to carefully schedule their connections or data transfers. In the Calendar Mode, an ALTO Server exposes ALTO Cost Values in JSON arrays where each value corresponds to a given time interval. The time intervals as well as other Calendar attributes are specified in the IRD. Besides the functional time-shift enhancement the ALTO Cost Calendar also allows to schedule the ALTO requests themselves and thus save a number of ALTO transactions.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

IETF is currently standardizing the ALTO protocol which aims for providing guidance to overlay applications, that need to select one or several hosts from a set of candidates that are able to provide a desired resource. This guidance is based on parameters that affect performance and efficiency of the data transmission between the hosts, e.g., the topological distance. The goal of ALTO is to improve the Quality of Experience (QoE) in the application while simultaneously optimizing resource usage in the underlying network infrastructure.

The ALTO protocol therefore [RFC7285] specifies a Network Map, which defines groupings of endpoints in a network region (called a PID) as seen by the ALTO server. The Endpoint Cost Service and the Endpoint (EP) Ranking Service then provide rankings for connections between the specified network regions and thus incentives for application clients to connect to ISP preferred endpoints, e.g. to reduce costs imposed to the network provider. Thereby ALTO intentionally avoids the provisioning of realtime information as explained in the ALTO Problem Statement [RFC5693] and ALTO Requirements [RFC5693]) drafts that write "Such information is better suited to be transferred through an in-band technique at the transport layer instead". Thus the current Cost Map and Endpoint Cost Service are providing, for a given Cost Type, exactly one rating per link between two PIDs or to an Endpoint. Applications are expected to query one of these two services in order to retrieve the currently valid cost values. They therefore need to plan their ALTO information requests according to the estimated frequency of cost value change.

Since network costs can fluctuate, due to diurnal patterns of traffic demand andor network maintenance, an ALTO client should interpret the returned costs as those at the query moment. Providing network costs for only the current time, however, may not be sufficient, in particular, for applications that can schedule their traffic in a span of time, for example, by deferring backup to night during traffic trough.

In case these value changes are predicable over a certain period of time and the application does not require immediate data transfer, it would save time to get the whole set of cost values over the period in one ALTO response and using these values to schedule data transfers would allow to optimise the network resources usage and QoE.

In this draft we introduce use cases that describe applications that have a degree of freedom on scheduling data transfers over a period of time, thus they do not need to start a transfer instantaneously on a retrieved request. For this kind of applications we propose to extend the Cost Map and Endpoint Cost Services by adding a calendar on the cost values, allowing applications to time-shift data transfers.

This document extends RFC7285 to allow an ALTO server to provide network costs for a given duration of time. A sequence of network costs across a time span for a given pair of network locations is referred to as an ALTO cost calendar for the pair of network locations. In addition to this functional ALTO enhancement, we expect to further gain by gathering multiple Cost Values for one cost type as firstly one Cost Map reporting on N Cost Values is less bulky than N Cost Maps containing one Cost value each and secondly, this reduces N ALTO transactions to a single one. This is valuable for both the storage of these ALTO maps and their transfer. Similar gains can be obtained for the ALTO Endpoint Cost Service.

In this draft an "ALTO Calendar" is presented as a Cost Mode that is applicable to time-sensitive ALTO metrics and allows applications using such metrics to carefully schedule their connections or data transfers. In the Calendar Mode, an ALTO Server exposes ALTO Cost Values in JSON arrays where each value corresponds to a given time interval. The time intervals as well as other Calendar attributes (the ones suggested by Richard) are specified in the IRD and allow the ALTO Client to interpret the received ALTO values. This draft proposes a set of Calendar attributes to be added to the IRD, for discussion in the ALTO WG. In order to support the calendaring of Cost values represented in Modes such as 'string' this draft also proposes one option, which is to extend the Cost Mode to a combination of several indicators, such as 'string' and 'calendar'.

The remainder of this draft first provides a variety of use cases that motivate the need for a 'calendar' cost mode. It then specifies the needed extensions to the ALTO protocol and details some example messages.

2. Motivating Use cases for ALTO Cost Schedule

This section introduces use cases showing the benefits of providing ALTO Cost values in 'calendar' mode. Most likely, the ALTO Cost Calendar would be used for the Endpoint Cost Service, assuming that a limited set of feasible Endpoints for a non-real time application is already identified, that they do not need to be accessed immediately and that their access can be scheduled within a given time period. The Cost Map service, filtered or not, is also applicable as long as the size of the Map is manageable.

Last, the ALTO Cost calendar is beneficial to optimizing ALTO transactions themselves. Indeed, let us assume that an Application Client is located in an end sytem with limited resources and/or has an access to the network that is either intermittent or provides an acceptable QoE in limited but predictable time periods. In that case, it needs to both schedule its resources demanding networking activities and its ALTO requests. Instead of having to figure out when the cost values may change and having to carefully schedule multiple ALTO requests, it could aviod this by relying on Cost Shedule attributes that indicate the time granularity, the validity and time scope of the cost information, together with the time related cost values themselves.

2.1. Bulk Data Transfer scheduling

Large Internet Content Providers (ICPs) like Facebook or YouTube, as well as CDNs rely on data replication across multiple sites to offload the core site and increase user experience through shorter latency from a local site. Typically the usage pattern of these data centers or caches follows a location dependent diurnal pattern.

In the examples above, data needs to be replicated across the various locations of a Internet Content Provider (ICP), leading to bulk data transfers between datacenters on a diurnal pattern.

In the mean time, there is a degree of freedom on when the content is transmitted from the origin server to the caching node, or from the core site to a local site. However, scheduling these data transfers is a non-trivial task as the transfer should not infer with the user peak demand to avoid degradation of user experience and to decrease billing costs for the datacenter operator by leveraging off-peak hours for the transfer. This peak demand typically follows a diurnal pattern according to the geographic region of the datacenter.

As a result, it would be very helpful to let these ICPs to have a good knowledge about the link utilization patterns between the different datacenters from the networks before making a more intelligent scheduling decision. While this usage data today already is gathered and also used for the scheduling of data transfer, provisioning this data gets increasingly complex with the number of CDN nodes and in particular the number of datacenter operators that are involved. For example, privacy concerns prevent that this kind of data is shared across administrative domains. The ALTO Cost Calendar specified later in this document avoids this problem by presenting an abstracted view of time sensitive utilization maps through a dedicated ALTO service to allow ICPs a coherent scheduling of such data transfers across administrative domains.

2.2. Endsystems with limited connectivity or access to datacenters

Another use case that benefits from the availability of multi-timeframe cost information is based on applications that are limited by their connectivity either in time or resources or both. For example applications running on devices in remote locations or in developing countries that need to synchronize their state with a data center periodically, in particular if sometimes there is no connection at all. Example applications is enterprise database update, remote learning, remote computation distributed on several data center endpoints.

Wireless connections have a variable quality and may even be intermittent. On the other hand, the wireless network conditions are often predicable and have a rapid impact on applications. Non real time applications and time-insensitive data transfers such as client patching, archive syncing, etc. can benefit from careful scheduling. It is thus desirable to provide ALTO clients with routing costs to connection nodes (i.e. Application Endpoints) over different time periods. This would allow end systems using ALTO aware application clients to schedule their connections to application endpoints.

Another challenge arises with end systems using resources located in datacenters and trading content and resources scattered around the world. For non-real time applications, the interaction with Endpoints can be scheduled at the time slots corresponding to the best possible network conditions in order to improve the QoE. For instance, resource Ra downloaded from Endpoint EPa at time t1, Resource Rb uploaded to EPb at time t2, some batch computation involving Ra and Rb done on EPc at time t3 and results R(A,B) downloaded to EPd and EPe at time t4. Example applications are similar to the ones cited in the previous paragraph.

```
 +-----+                                          +-----+
 | EPa |                                          | EPb | <----- Rb
 +-----+                                          +-----+   (t2=50)
    |                    +-------+                   |
    Ra -------------->  | EPc   |                   |
    (time t1=10)         |       |                   |
                         |t3=100 |  <----------------- Rb
                         +-------+
                             | \
                             |  \
                          R(Ra,Rb)
                         (t4=200)
                             |     \
                             |      -------------------.
                             V                          V
                          +-----+                    +-----+
                          | EPd |                    | EPe |
                          +-----+                    +-----+
```

2.3. SDN Controller guided access to application endpoints

The Software Defined Networking (SDN), see [sdnrg], is a model that attempts to manage and reconfigure networks in a more flexible way in order to better cope with the traffic challenges posed by nowadays resources greedy applications. To this end, one option is "moving the control plane out of the network elements into "controllers", see [SDN charter, http://www.1-4-5.net/~dmm/sdnrg/sdnrg.html], that implements the network control and management. The SDN Controllers are deemed to gather the network state information and provide it in an abstracted form to SDN aware applications while gathering their requirements in QoE and exchanging other application "management" information and commands.

The relevance of ALTO to perform a number of SDN functions has been recently highlighted. An ALTO Server can assist an SDN Controller by hosting abstracted network information that can be provided to SDN aware applications via an ALTO Client. It can also assist other SDN Control operations using information in and outside the ALTO scope.

The SDN primitive "Get network resources" provides applications with informations allowing them to evaluate the expected QoE. QoE related information includes delay and bandwidth at the application endpoints as well as on the network paths. Such information may be provided via the ALTO Service by proposed extensions of the ALTO protocol that define new ALTO Cost Types allowing to abstract and report QoE to applications.

One key objective of an SDN controller is the ability to balance the application traffic whenever possible. For non real time applications, data and resources transfer can be time shifted, resources availability may often be predicable and last, strong incentives for applications to time shift their traffic may be given by network operators appropriately setting routing cost values at different time values, according to their policy to cope with network occupation over time.

To achieve this objective, the SDN controller can:

1. get the network state history from its controlled network elements through its southbound API

2. possibly derive an estimation or a prediction of these values over given time frames

3. compute estimates and/or network provider preferences on end to end paths and store their abstraction in an ALTO Server in the form of ALTO Cost Calendar values defined for different time periods

4. deliver these values to the SDN applications via the ALTO Endpoint Cost Service, as estimations covering the past and/or the future and/or preferences.

This way:

- On one hand, the applications get the best possible QoE, as they can pick the best time for them to access one or more Endpoints,

- One the other hand the SDN controller achieves load balancing as it may guide the application traffic so as to better distribute the traffic over time, and thus optimize its resources usage.

2.4. Large flow scheduling on extended ALTO topologies

[draft-yang-alto-topology-00] presents initial thinking on extending ALTO for topology exposure services, that would provide flexible abstractions based on the raw network topology. Among other features, an ALTO topology may expose several paths between a source (src) and destination (dst), or topology details may be provided on restricted parts. This work was presented to the ALTO WG at IETF88.

The presentation slides [slides-88-alto-5-topology] on [draft-yang-alto-topology-00] expose a use case entitled "Large Flow Scheduling". This case includes a "daylife example" where a Google Map service proposes multiple routes between 2 points A and B, each calculated w.r.t. length and estimated time. For each of these selected paths, the map service exposes a time-sensitive qualitative value taking 4 values between Slow and Fast. A user of this application may thus organize its transfer w.r.t. metrics, paths and time, provided s/he does not have to commute immediately.

The use case on Large flow scheduling on extended ALTO topologies in the present section illustrates one modality of ALTO topology service, that would expose several paths between end to end (src, dst) pairs, computed w.r.t. one of more metrics, possibly under given constraints. On top of this enriched topology service, non real-time applications may also choose the time of data/resources transfer, taking thus advantage of a richer set of decision variables.

The use case "Large Flow Scheduling" of presentation [slides-88-alto-5-topology]can thus be adapted as follows:

- Step1 - obtain the set T transfer tasks {(src, dest, data)}

- Step2 - identify one or more paths for each (src, dst): several information sources exist among which:

  * (a) ALTO CostMap with a "path" metric, // not specified here

  * (b) an ALTO Topology Service providing a path computation hint (e.g. w.r.t. routingcost and/or other metrics)

- Step 3 - while T not empty:

  * 1 - query for example values for some metric 'available bandwidth' on paths:

    + to this end, query the values in the ALTO 'calendar' Mode: on the selected (src, dst) for a set of time intervals. With this mode, the ALTO client will receive an array of values, each applicable to a time slot .

  * 2 - schedule data transfer at the time slots corresponding to the preferred value.

2.5. Time-sensitve TE metrics Calendaring

Draft [draft-wu-alto-te-metrics] , proposes to extend the set of ALTO metrics with 11 ALTO traffic engineering (TE) metrics to reflect measurement on network delay, jitter, packet loss, hop count, and bandwidth. ALTO TE metrics that are time-sensitive, either by nature such as bandwidth and delay related metrics, or due to "normally" changing network conditions or both.

The values of ALTO TE metrics are typically collected from routing protocols and provided in a non-real time manner. In "normally" changing network conditions, TE metric values remain uniformly distributed over given time intervals and can be aggregated over bigger time intervals of periodic patterns. For example, an ALTO Server may collect values for e.g. delay from a routing protocol produced by measurements done every second over a measurement period of 30 seconds. The ALTO Server may then aggregate these values over two measurement periods (i.e. 60 seconds) and repeat the operation as it wishes. Then every hour, the ALTO Server provides these delay values in 'calendar' mode, encoded as an array of 60 values, assumed to estimate network performance statistics on each minute of this hour.

Another example is Bandwidth Calendaring. Bandwidth Calendaring allows network operators to reserve resources in advance according to agreements with their customers, enabling them to transmit data with specified starting time and duration, for example, for a scheduled bulk data replication between data centers. Traditionally, this can be supported by a Network Management System operation such as path pre-establishment and activation on the agreed starting time. However, this does not provide efficient network usage since the established paths exclude the possibility of being used by other services even when they are not used for undertaking any service.

A Cost calendar provided by an ALTO server can support the scheduled bulk data replication application with better efficiency since it can alleviate the burden of processing on network elements. This requires the ALTO server to maintain the calendared TE cost metrics on the end to end paths associated to data transfer.

To support cost calendaring for these time-sensitive ALTO TE metrics, the network topology and the dynamicity of the traffic need to be considered. For example, a small topology with low density and low capacity that carries inpredictable, heavy and bursty traffic has few chances to exhibit stationary TE metric value patterns over large periods and would benefit to use the ALTO Calendar over smaller time slots. Some ALTO TE metric values, even aggregated over time may need to be updated at a frequency that would require doing ALTO request at a pace that would be overload both the ALTO Client and the Server.

3. Design considerations for an ALTO calendar

This section enumerates a set of challenges in designing the calendaring specifications, and will be updated upon discussions in the ALTO WG.

An ALTO Cost calendar provided by the ALTO Server is an array of values for a given metric, where each value corresponds to a time interval which length is specified for this metric in the IRD, together with other attributes describing the time scope of the calendar. Most likely, the ALTO Cost Calendar would be used for the Endpoint Cost Service, assuming that a limited set of feasible Endpoints for a non-real time application is already identified, that they do not need to be accessed immediately and that their access can be scheduled within a given time period. The Cost Map service, filtered or not, is also applicable as long as the size of the Map is manageable.

Given a cost calendar (i.e., a sequence of cost values such as [1, 2, 3, 4]), an ALTO client needs meta information to interpret the values (Q1) Which network metric (e.g., routingcost, or latency) do the numbers represent (Q2) Are the values absolute values or relative ranking order or strings ? (i.e., numerical or ordinal or strings) and (Q3) What is the time associated with each value ? RFC7285 defines CostType, which consists of two fields named cost-metric and cost-mode, to provide information for Q1 and Q2. A design issue in providing ALTO cost calendar is to provide information for Q3.

3.1. Purpose of an ALTO calendar

A calendar is used to schedule transfers of application data or services and has several characteristics:

- o the Calendar values are assumed to be stationary on each time interval,
- o the ALTO Server may provide values on past time periods that can be interpreted as historical experience and used to anticipate future cost values,
- o the ALTO Server may provide stationary values on present or future time periods that can be interpreted as predictions on cost values,

- o the ALTO Server may provide stationary values on time intervals covering the past, and/or present and/or future.

- o for metrics provided with units and claiming to be aggregated from network measurements, the values can be interpreted as estimations.

- o For abstracted metrics provided with no units such as the 'routingcost' defined in the base ALTO protocol or abstracted unitless scores on network performances such as some potential 'bandwidth score' or 'unreliability cost', the values can be interpreted as network provider preferences.

Note that we distinguish between "estimates" that we see as value aggregations represented with units such as bytes, seconds, percentage and "preferences" that we see as abstracted costs or scores w.r.t. a metric or state such as 'routingcost', 'bandwidthscore', 'link quality'.

The method used to generate the estimation and aggregation of measured values is currently outside the scope of this draft and expected to be documented in the applicable metric definition document.

3.2. Design requirements for an ALTO calendar

TO BE COMPLETED IN FURTHER DRAFT VERSIONS

An ALTO Calendar can be seen as a cyclic value array pattern that is valid for a certain time period with specified beginning date, duration and number of time intervals.

- o needs to convey cyclic network provider preferences expressed w.r.t. given ALTO metric values (e.g., hourly, daily, weekly measurement/prediction)

- o needs to convey cyclic network status if the ALTO Server claims to provide aggregated information on network status (e.g., hourly, daily, weekly measurement/prediction)

- o needs to be able to convey the result of a particular instance of time (e.g., to convey predicted network status during a maintenance outage on July 4, 2014 from 5-7pm)

- o needs at least the following attributes to report on cyclic patterns:

    * generic time zone,

   * applicable time interval for each calendar value (measurement estimation with units or unitless preference value) : combining <nb-int-unit> and <interval-unit> to reflect for example: 1hour, 2minutes, 1week, 1month

   * date range of the Calendar, e.g. number of intervals allowing to derive the calendar time range in terms of: year, month, week, day, hour, min, secs

- needs to expose validity period of the calendar: indicating when the next ALTO Calendar for this date range should be fetched if needed,

- needs to provide time stamps:

   * last-update-time: specifying when the metric values were last computed ,

   * next-update-time: specifying when the calendar values will be re-computed, indicating thus when an ALTO client should fetch an update if it uses a Calendar.

   * calendar-start-date: specifying when the current already computed calendar starts,

   * next-calendar-start-date: specifying when the already computed calendar will have different values, indicating thus that the ALTO client should fetch the next pre-computed calendar

It may be useful to keep a cyclic network status with date, in case of exceptional predicted events such as New Year evening on a Tuesday or any worldwide event generating a lot of traffic.Traffic calendars may be particularly useful in such cases.

4. ALTO extensions for a Cost Calendar

The usage of a time-related ALTO Cost Calendar is rather proactive in that it can be used like a "time table" to figure out the best time to schedule data transfer and also anticipate predictable events including predictable flash crowds. An ALTO Cost Calendar should be viewed as a synthetic abstraction of real measurements that can be historic or be a prediction for upcoming time periods.

Specifications on the cost "calendar" attributes are proposed here and will be completed in further versions of this draft, upon discussion with the ALTO WG.

The format of ALTO requests and responses will be specified in further versions of this draft, as in particular it may be necessary that the ALTO response indicates the computation and validity dates of the provided ALTO Calendar.

4.1. ALTO Cost-Mode: Calendar

This draft introduces a new ALTO Cost Mode called "calendar". This mode applies preferably to Costs that can be expressed in a single-valued Cost Mode. In that sense, when the "numerical" mode is available for a Cost-Type, the cost expressed in the "calendar" mode is an extension of its expression from one value in the "numerical" mode to an array of several values varying over time.

Types of Cost values such as JSONBool can also be expressed in the "calendar" mode, as states may be "true" or "false" depending on given time periods. They may be expressed as a single value which is either "true" or "false" following a decision rule outside the ALTO protocol.

DISCUSSION: The current draft focuses on Calendars representing values encoded in the ’numerical’ mode. However, Calendars should also be able to represent time sensitive values represented by strings, such as "medium", "high", "low". To support this, one option is to extend the Cost Mode to be a combination (attribute) indicators, for example cost-mode : "ordinal;calendar"

cost-mode : "ordinal;calendar"

We chose ";" as separator to be consistent with existing formats such as the HTTP list of multiple options (e.g., Accept).

4.2. ALTO Calendar attributes in the IRD

To ensure that the application client understands the provided information in the cost calendar in an unambiguous way, we specify the Calendar attributes in the ALTO IRD "meta" information, that defines the time scope of the "calendared" cost values. The Calendar attributes in the IRD are meant to carry constant dateless values.

- o time-interval-size:

  - * expresses the unit in which the duration of an ALTO calendar time interval duration is expressed appended to the number of these units. The time unit, ranges from "second" to "year". The number is encoded with an integer. Example values are: "5 minute" , "2 hour". These vales mean that each calendar value is provided on a time interval that lasts respectively 5 minutes and 2 hours.

o numb-intervals:

* the integer number of values of the cost calendar array, at least equal to 1.

o repeat:

* an integer number representing the number of times that the Calendar pattern repeats. That is : if for example a given daily pattern is represented Calendar is made of "num-intervals" = 24 cost values each applicable to a time slot lasting "time-interval-size" = 1 hour and the value of 'repeat' is 5, the client can interpret that 5 daily patterns starting from the date indicated in the Server response have identical values.

- Attributes 'time-interval-size' and 'numb-intervals', when mutlipled, reflect the duration of the provided calendar. For example an ALTO Server may provide a calendar for ALTO values changing every 'time-interval-size' equal to 5 minutes. If 'numb-intervals' has the value 12, then the duration of the provided calendar is "1 hour". Note also that in this example, a 5 minutes interval may cover the aggregation of real TE measurements done every 30 seconds, but this latter aspect is outside the scope of this draft as it is to be specified in the definition of the ALTO metric.

- Attribute 'repeat' reflects the frequancy at which the patterns represented by a Calendar may change. This information is completed by the 'start' attribute provided in Server responses to calendar requests.

NOTE that: To cope with existing representation fomats, further versions will re-name these attributes. The current proposal for renaming in further versions is to replace "time-interval-size" by "interval" and "num-interval" by "count" :

4.3. Example of calendared information resources in the IRD

This section describes an example IRD and related ALTO calendar transaction in a scenario where an ALTO Server offers the Calendar Mode for several Cost Types that are either specified in the base ALTO protocol or proposed in other drafts see [draft-wu-alto-te-metrics] or suggested here as examples, like a cost metric reporting on measured packet loss and called 'TEpktloss. The provided example transactions are based on the use cases of section 2.

These examples describe situations where a client has the choice of trading content or resources with several Endpoints and needs to decide with which Endpoint it will trade and at what time. For instance, one may assume that the Endpoints are spread over different time-zones, or have intermittent access. The ALTO Calendar mode specified below allows these clients to retrieve Endpoint Cost Maps valid for a certain timeframe (e.g. 24 hours), and get a set of values, each applicable on a specified time interval (e.g. 1 hour). Thus the application can optimize the needed data transfer according to this information.

In the example IRD of the present draft, the available Endpoint Costs metrics are: "routingcost", "AShopcount", 'TEpktloss' and 'Availbandwidth'. "routingcost" and "AShopcount" are available in the "numerical" Cost Mode. 'TEpktloss' , 'Availbandwidth' and "routingcost" as well are available in the "calendar" Cost Mode.

We suppose that the ALTO Client GETs the IRD on Tuesday July 1st 2014 at 13:00

- o The Calendar for 'TEpktloss'': is an hourly pattern that consists of 12 values provided each on a time interval of 5 minutes, and the values are the same for the next 2 hours.

- o The Calendar for 'Availbandwidth': is a daily pattern that consists of 12 values provided each on time intervals of 2 hours, with the first interval starting at 0h00. This information is typically used to enable applications to see which time intervals in a day are the most favorable to operate, and which "busy " time intervals should be avoided. The pattern is the same for the next "repeat" = 7days.

- o The Calendar for 'routingcost': is a daily pattern that consists of an array of 24 time intervals lasting each 1 hour. The routingcost calendar covers a 1 day period, starting at midnight. This may be applicable for networks with poor or intermittent connectivity where the operator may integrate monetary as well as network performance metrics in the provided 'routingcost' values.The pattern is the same for the nex "repeat" = 4 days.

4.3.1. Example IRD with ALTO cost Calendars

The example IRD given in this section includes 2 particular URIs:

o "http://alto.example.com/endpointcost/lookup", in which the ALTO Server offers the numerical mode for metrics "routingcost" and "AShopcount".

o "http://alto.example.com/endpointcost/calendar/lookup", in which the ALTO Server provides "calendar" mode for metrics 'TEpktloss' and 'Availbandwidth' and 'routingcost'.

For Cost Type 'calendar-routing', this example assumes that the ALTO Server has defined 3 different daily patterns each represented by a Calendar, to cover the week of Monday June 30th at 00:00 to Sunday July 6th 23:59:

- C1 for Monday, Tuesday, Wednesday, Thursday, (week days)

- C2 for Saturday, Sunday, (week end)

- C3 for Friday (maintenance outage on July 4, 2014 from 02:00:00 GMT to 04:00:00 GMT, or big holiday such as New Year evening)

The example ALTO response shown in a further section also illustrates how specific calendar attributes allow an ALTO client to fetch 3 Calendars instead of 7 and thus to reduce the volume of on-the-wire data exchange. For Cost Type 'calendar-routing' , the IRD provides a value for attribute 'num-calendars' which is equal to 3.

```
GET /directory HTTP/1.1
   Host: alto.example.com
   Accept: application/alto-directory+json,application/alto-error+json


HTTP/1.1 200 OK
   Content-Length: [TODO]
   Content-Type: application/alto-directory+json

   {
     "meta" : {
        "cost-types": {
           "num-routingcost": {
              "cost-mode" : "numerical",
              "cost-metric" : "routingcost"
              },
           "num-AShopcount": {
              "cost-mode" : "numerical",
              "cost-metric" : "hopcount"
              },
           "calendar-TEpktloss": {
              "cost-mode"  : "calendar",
```

```
              "cost-metric": "TEpktloss",
              "description": {
                  "time-interval-size" : "5 minute",
                  "numb-intervals" : 12,
                  "repeat" : 2
              }
           },
           "calendar-bw": {
              "cost-mode"  : "calendar",
              "cost-metric": "Availbandwidth",
              "description": {
                  "time-interval-size" : "2 hour",
                  "numb-intervals" : 12,
                  "repeat" : 7
              }
          },
          "calendar-routing": {
             "cost-mode"  : "calendar",
             "cost-metric": "routingcost",
             "description": {
                 "time-interval-size" : "1 hour",
                 "numb-intervals" : 24,
                 "repeat: 4"

         }
         ... other meta ...
       },

    "resources" : {

       ... usual ALTO resources such as Network Map, Cost Maps ...

       "endpoint-cost" : {
         "uri" : "http://alto.example.com/endpointcost/lookup",
         "media-types" : [ "application/alto-endpointcost+json" ],
         "accepts" : [ "application/alto-endpointcostparams+json" ],
         "capabilities" : {
           "cost-constraints" : true,
           "cost-type-names" : [ "num-routingcost", "num-AShopcount" ]
           }
         },
       "endpoint-cost-calendar-map" : {
         "uri" : "http://alto.example.com/endpointcost/calendar/lookup",
         "media-types" : [ "application/alto-endpointcost+json" ],
         "accepts" : [ "application/alto-endpointcostparams+json" ],
         "capabilities" : {
           "cost-constraints" : true,
           "cost-type-names" : [ "calendar-routingcost",
```

```
                                    "calendar-TEpktloss",
                                    "calendar-bw"]
           }
         }
     }
  }
```

4.4. ALTO Calendar information in ALTO responses

ALTO responses convey additional attributes with usually non constant values that inform the ALTO Client about the next date at which the calendar values stored in the ALTO Server will change and at which time updates calendar values will be uploaded in the ALTO Server. A number of Calendar attributes in ALTO responses are dates. The reference time zone for the provided values is GMT. Indeed, the option chosen to express the time format is the HTTP header fields formats such as:

Date: Tue, 15 Nov 1994 08:12:31 GMT

- calendar-start-time:

  * the date corresponding to the first value in the calendar values array

- time-interval-size: as defined for the IRD

- numb-intervals: as defined for the IRD

- repeat: as defined for the IRD

- Attribute 'calendar-start-time' indicates when the calendar provided to the ALTO client starts. If the 'calendar-start-time' date is past, the application can also use the information to compute statistics on values provided by ALTO over time to guide applications. Besides estimating some customized prediction the ALTO Client may use these values to assess their reliability w.r.t. some real measures of QoE.

Discussion: like for the attributes of the IRD, "calendar-start-time" can be renames "start" to comply with existing formats.

4.4.1. Example transaction for a routingcost Calendar to face intermittent connectivity

Let us assume an Application Client located in an end sytem with limited resources and having an access to the network that is either intermittent or provides an acceptable quality in limited but possibly predictable time periods. Therefore, it needs to both schedule its resources demanding networking activities and minimize its ALTO transactions.

The Application Client has the choice to trade content or resources with a set of Endpoints of moderate 'routingcost', and needs to decide with which Endpoint it will trade at what time. For instance, one may assume that the Endpoints are spread on different time-zones, or have intermittent access. In this example, the 'routingcost' is assumed to be the time sentitive decision metric, with values provided in the ALTO Calendar Mode.

The ALTO Client embedded in the Application Client queries an ALTO Calendar on 'routingcost' and will get the Calendar covering the 24 hours time period "containing" the date and time of the ALTO client request. We suppose in this example that the ALTO Client sends its request on Tuesday July 1st 2014 at 13:15

The present example also illustrates how attributes "repeat" and "calendar-start-time" allow an ALTO client to fetch 3 Calendars instead of 7 and thus to reduce the volume of on-the-wire data exchange, because the ALTO Server has defined 3 different daily patterns each represented by a Calendar, to cover the week of Monday June 30th at 00:00 to Sunday July 6th 23:59:

- C1 for Monday, Tuesday, Wednesday, Thursday, (week days)

- C2 for Saturday, Sunday, (week end)

- C3 for Friday (maintenance outage on July 4, 2014 from 5-7pm, or holiday such as New Year evening)

```
POST endpointcost/calendar/lookup HTTP/1.1
  Host: alto.example.com
  Content-Length: [TODO]
  Content-Type: application/alto-endpointcostparams+json
  Accept: application/alto-endpointcost+json,application/alto-error+json

  {
    "cost-type" : {"cost-mode" : "calendar", "cost-metric" : "routingcost"},
    "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-endpointcost+json

  {
    "meta" : {
      "calendar-start-time" : Tue, 1 Jul 2014 00:00:00 GMT,
      "time-interval-size" : "1 hour",
      "numb-intervals" : 24,
      "repeat": 4
    },
    "cost-type" : {"cost-mode" : "calendar", "cost-metric" : "routingcost"},
    "endpoint-cost-calendar-map" : {
        "ipv4:192.0.2.2": {
          "ipv4:192.0.2.89"    : [7, ... 24 values],
          "ipv4:198.51.100.34" : [4, ... 24 values],
          "ipv4:203.0.113.45"  : [2, ... 24 values]
        }
    }
  }
```

NOTE also: that the Calendar could enumerate 24*7 hourly values to represent the pattern for such a week.

4.4.2. Example transaction for a bandwidth calendar

An example of non-real time information that can be provisioned in a 'calendar' is the expected path bandwidth. While the transmission rate can be measured in real time by end systems, the operator of a data center is in the position of formulating preferences for given paths, at given time periods for example to avoid traffic peaks due to diurnal usage patterns. In this example, we assume that an ALTO Client requests a bandwidth calendar as specified in the IRD to shedule its bulk data transfers as described in the use cases of sections 2.1 and 2.5.

We suppose in this example that the ALTO Client sends its request on Tuesday July 1st 2014 at 13:15

```
POST endpointcost/calendar/lookup HTTP/1.1
  Host: alto.example.com
  Content-Length: [TODO]
  Content-Type: application/alto-endpointcostparams+json
  Accept: application/alto-endpointcost+json,application/alto-error+json

  {
    "cost-type" : {"cost-mode" : "calendar", "cost-metric" : "Availbandwidth"},
    "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-endpointcost+json

  {
    "meta" : {
      "calendar-start-time" : Tue, 1 Jul 2014 00:00:00 GMT,
      "time-interval-size" : "2 hour",
      "numb-intervals" : 12,
      "repeat" : 7
    },
    "cost-type" : {"cost-mode" : "calendar", "cost-metric" : "Availbandwidth"},
    "endpoint-cost-calendar-map" : {
        "ipv4:192.0.2.2": {
          "ipv4:192.0.2.89"    : [7, ... 12 values],
          "ipv4:198.51.100.34" : [4, ... 12 values],
          "ipv4:203.0.113.45"  : [2, ... 12 values]
        }
    }
  }
```

5. Use cases for ALTO Cost Schedule

This section introduces use cases showing the benefits of providing ALTO Cost values in 'calendar' mode. Most likely, the ALTO Cost Calendar would be used for the Endpoint Cost Service, assuming that a limited set of feasible Endpoints for a non-real time application is already identified, that they do not need to be accessed immediately and that their access can be scheduled within a given time period. The Cost Map service, filtered or not, is also applicable as long as the size of the Map is manageable.

Last, the ALTO Cost calendar is beneficial to optimizing ALTO transactions themselves. Indeed, let us assume that an Application Client is located in an end sytem with limited resources and/or has an access to the network that is either intermittent or provides an acceptable QoE in limited but predictable time periods. In that case, it needs to both schedule its resources demanding networking activities and its ALTO requests. Instead of having to figure out when the cost values may change and having to carefully schedule multiple ALTO requests, it could aviod this by relying on Cost Shedule attributes that indicate the time granularity, the validity and time scope of the cost information, together with the time related cost values themselves.

5.1. Bulk Data Transfer scheduling

Large Internet Content Providers (ICPs) like Facebook or YouTube, as well as CDNs rely on data replication across multiple sites to offload the core site and increase user experience through shorter latency from a local site. Typically the usage pattern of these data centers or caches follows a location dependent diurnal pattern.

In the examples above, data needs to be replicated across the various locations of a Internet Content Provider (ICP), leading to bulk data transfers between datacenters on a diurnal pattern.

In the mean time, there is a degree of freedom on when the content is transmitted from the origin server to the caching node, or from the core site to a local site. However, scheduling these data transfers is a non-trivial task as the transfer should not infer with the user peak demand to avoid degradation of user experience and to decrease billing costs for the datacenter operator by leveraging off-peak hours for the transfer. This peak demand typically follows a diurnal pattern according to the geographic region of the datacenter.

As a result, it would be very helpful to let these ICPs to have a good knowledge about the link utilization patterns between the different datacenters from the networks before making a more intelligent scheduling decision. While this usage data today already is gathered and also used for the scheduling of data transfer, provisioning this data gets increasingly complex with the number of CDN nodes and in particular the number of datacenter operators that are involved. For example, privacy concerns prevent that this kind of data is shared across administrative domains. The ALTO Cost

Calendar specified later in this document avoids this problem by presenting an abstracted view of time sensitive utilization maps through a dedicated ALTO service to allow ICPs a coherent scheduling of such data transfers across administrative domains.

5.2. Endsystems with limited connectivity or access to datacenters

Another use case that benefits from the availability of multi-timeframe cost information is based on applications that are limited by their connectivity either in time or resources or both. For example applications running on devices in remote locations or in developing countries that need to synchronize their state with a data center periodically, in particular if sometimes there is no connection at all. Example applications is enterprise database update, remote learning, remote computation distributed on several data center endpoints.

Wireless connections have a variable quality and may even be intermittent. On the other hand, the wireless network conditions are often predicable and have a rapid impact on applications. Non real time applications and time-insensitive data transfers such as client patching, archive syncing, etc. can benefit from careful scheduling. It is thus desirable to provide ALTO clients with routing costs to connection nodes (i.e. Application Endpoints) over different time periods. This would allow end systems using ALTO aware application clients to schedule their connections to application endpoints.

Another challenge arises with end systems using resources located in datacenters and trading content and resources scattered around the world. For non-real time applications, the interaction with Endpoints can be scheduled at the time slots corresponding to the best possible network conditions in order to improve the QoE. For instance, resource Ra downloaded from Endpoint EPa at time t1, Resource Rb uploaded to EPb at time t2, some batch computation involving Ra and Rb done on EPc at time t3 and results R(A,B) downloaded to EPd and EPe at time t4. Example applications are similar to the ones cited in the previous paragraph.

```
     +-----+                                           +-----+
     | EPa |                                           | EPb | <----- Rb
     +-----+                                           +-----+   (t2=50)
        |                   +-------+                     |
        Ra -------------->  | EPc   |                     |
        (time t1=10)        |       |                     |
                            |t3=100 |  <----------------- Rb
                            +-------+
                                | \
                                |  \
                              R(Ra,Rb)
                             (t4=200)
                                |     \
                                |      -------------------.
                                V                         V
                             +-----+                   +-----+
                             | EPd |                   | EPe |
                             +-----+                   +-----+
```

5.3. SDN Controller guided access to application endpoints

The Software Defined Networking (SDN), see [sdnrg], is a model that attempts to manage and reconfigure networks in a more flexible way in order to better cope with the traffic challenges posed by nowadays resources greedy applications. To this end, one option is "moving the control plane out of the network elements into "controllers", see [SDN charter, http://www.1-4-5.net/˜dmm/sdnrg/sdnrg.html], that implements the network control and management. The SDN Controllers are deemed to gather the network state information and provide it in an abstracted form to SDN aware applications while gathering their requirements in QoE and exchanging other application "management" information and commands.

The relevance of ALTO to perform a number of SDN functions has been recently highlighted. An ALTO Server can assist an SDN Controller by hosting abstracted network information that can be provided to SDN aware applications via an ALTO Client. It can also assist other SDN Control operations using information in and outside the ALTO scope.

The SDN primitive "Get network resources" provides applications with informations allowing them to evaluate the expected QoE. QoE related information includes delay and bandwidth at the application endpoints as well as on the network paths. Such information may be provided via the ALTO Service by proposed extensions of the ALTO protocol that define new ALTO Cost Types allowing to abstract and report QoE to applications.

One key objective of an SDN controller is the ability to balance the application traffic whenever possible. For non real time applications, data and resources transfer can be time shifted, resources availability may often be predicable and last, strong incentives for applications to time shift their traffic may be given by network operators appropriately setting routing cost values at different time values, according to their policy to cope with network occupation over time.

To achieve this objective, the SDN controller can:

1. get the network state history from its controlled network elements through its southbound API

2. possibly derive an estimation or a prediction of these values over given time frames

3. compute estimates and/or network provider preferences on end to end paths and store their abstraction in an ALTO Server in the form of ALTO Cost Calendar values defined for different time periods

4. deliver these values to the SDN applications via the ALTO Endpoint Cost Service, as estimations covering the past and/or the future and/or preferences.

This way:

- On one hand, the applications get the best possible QoE, as they can pick the best time for them to access one or more Endpoints,

- One the other hand the SDN controller achieves load balancing as it may guide the application traffic so as to better distribute the traffic over time, and thus optimize its resources usage.

5.4. Large flow scheduling on extended ALTO topologies

[draft-yang-alto-topology-00] presents initial thinking on extending ALTO for topology exposure services, that would provide flexible abstractions based on the raw network topology. Among other features, an ALTO topology may expose several paths between a source (src) and destination (dst), or topology details may be provided on restricted parts. This work was presented to the ALTO WG at IETF88.

The presentation slides [slides-88-alto-5-topology] on [draft-yang-alto-topology-00] expose a use case entitled "Large Flow Scheduling". This case includes a "daylife example" where a Google Map service proposes multiple routes between 2 points A and B, each calculated w.r.t. length and estimated time. For each of these selected paths, the map service exposes a time-sensitive qualitative value taking 4 values between Slow and Fast. A user of this application may thus organize its transfer w.r.t. metrics, paths and time, provided s/he does not have to commute immediately.

The use case on Large flow scheduling on extended ALTO topologies in the present section illustrates one modality of ALTO topology service, that would expose several paths between end to end (src, dst) pairs, computed w.r.t. one of more metrics, possibly under given constraints. On top of this enriched topology service, non real-time applications may also choose the time of data/resources transfer, taking thus advantage of a richer set of decision variables.

The use case "Large Flow Scheduling" of presentation [slides-88-alto-5-topology]can thus be adapted as follows:

- Step1 - obtain the set T transfer tasks {(src, dest, data)}

- Step2 - identify one or more paths for each (src, dst): several information sources exist among which:

  * (a) ALTO CostMap with a "path" metric, // not specified here

  * (b) an ALTO Topology Service providing a path computation hint (e.g. w.r.t. routingcost and/or other metrics)

- Step 3 - while T not empty:

  * 1 - query for example values for some metric ’available bandwidth’ on paths:

    + to this end, query the values in the ALTO ’calendar’ Mode: on the selected (src, dst) for a set of time intervals. With this mode, the ALTO client will receive an array of values, each applicable to a time slot .

  * 2 - schedule data transfer at the time slots corresponding to the preferred value.

5.5. Time-sensitve TE metrics Calendaring

Draft [draft-wu-alto-te-metrics] , proposes to extend the set of ALTO metrics with 11 ALTO traffic engineering (TE) metrics to reflect measurement on network delay, jitter, packet loss, hop count, and bandwidth. ALTO TE metrics that are time-sensitive, either by nature such as bandwidth and delay related metrics, or due to "normally" changing network conditions or both.

The values of ALTO TE metrics are typically collected from routing protocols and provided in a non-real time manner. In "normally" changing network conditions, TE metric values remain uniformly distributed over given time intervals and can be aggregated over bigger time intervals of periodic patterns. For example, an ALTO Server may collect values for e.g. delay from a routing protocol produced by measurements done every second over a measurement period of 30 seconds. The ALTO Server may then aggregate these values over two measurement periods (i.e. 60 seconds) and repeat the operation as it wishes. Then every hour, the ALTO Server provides these delay values in 'calendar' mode, encoded as an array of 60 values, assumed to estimate network performance statistics on each minute of this hour.

Another example is Bandwidth Calendaring. Bandwidth Calendaring allows network operators to reserve resources in advance according to agreements with their customers, enabling them to transmit data with specified starting time and duration, for example, for a scheduled bulk data replication between data centers. Traditionally, this can be supported by a Network Management System operation such as path pre-establishment and activation on the agreed starting time. However, this does not provide efficient network usage since the established paths exclude the possibility of being used by other services even when they are not used for undertaking any service.

A Cost calendar provided by an ALTO server can support the scheduled bulk data replication application with better efficiency since it can alleviate the burden of processing on network elements. This requires the ALTO server to maintain the calendared TE cost metrics on the end to end paths associated to data transfer.

To support cost calendaring for these time-sensitive ALTO TE metrics, the network topology and the dynamicity of the traffic need to be considered. For example, a small topology with low density and low capacity that carries inpredictable, heavy and bursty traffic has few chances to exhibit stationary TE metric value patterns over large periods and would benefit to use the ALTO Calendar over smaller time slots. Some ALTO TE metric values, even aggregated over time may need to be updated at a frequency that would require doing ALTO request at a pace that would be overload both the ALTO Client and the Server.

6. IANA Considerations

Information for the ALTO Endpoint property registry maintained by the IANA and related to the new Endpoints supported by the acting ALTO server. These definitions will be formulated according to the syntax defined in Section on "ALTO Endpoint Property Registry" of [ID-alto-protocol],

Information for the ALTO Cost Type Registry maintained by the IANA and related to the new Cost Types supported by the acting ALTO server. These definitions will be formulated according to the syntax defined in Section on "ALTO Cost Type Registry" of [RFC7285],

6.1. Information for IANA on proposed Cost Types

When a new ALTO Cost Type is defined, accepted by the ALTO working group and requests for IANA registration MUST include the following information, detailed in Section 11.2: Identifier, Intended Semantics, Security Considerations.

6.2. Information for IANA on proposed Endpoint Propeeries

Likewise, an ALTO Endpoint Property Registry could serve the same purposes as the ALTO Cost Type registry. Application to IANA registration for Endpoint Properties would follow a similar process.

7. Acknowledgements

Thank you to Diego Lopez, He Peng and Haibin Song and the ALTO WG for fruitful discussions.

8. References

8.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC5693] Seedorf, J. and E. Burger, "Application-Layer Traffic Optimization (ALTO) Problem Statement", RFC 5693, October 2009.

8.2. Informative References

[ID-alto-protocol]
R.Alimi, R. Penno, Y. Yang, Eds., "ALTO Protocol, RFC 7285", September 2014.

[RFC7285] R. Alimi, R. Yang, R. Penno, Eds., "ALTO Protocol", September 2014.

[article-gslh-alto-sdn]
V. Gurbani, M. Scharf, T.Lakshman, and V. Hilt, , "Abstracting network state in Software Defined Networks (SDN) for rendezvous services, IEEE International Conference on Communications (ICC) Workshop on Software Defined Networks (SDN)", June 2012.

[draft-jenkins-alto-cdn-use-cases-01]
B. Niven-Jenkins (Ed.), G. Watson, N. Bitar, J. Medved, S. Previdi, , "Use Cases for ALTO within CDNs, draft-jenkins-alto-cdn-use-cases-01", June 2011.

[draft-randriamasy-multi-cost-alto]
S. Randriamasy, Ed., W. Roome, N. Schwan, , "Multi-Cost ALTO (work in progress), draft-randriamasy-alto-multi-cost-07", October 2012.

[draft-wu-alto-te-metrics]
Q. Wu, Y. Yang, Y. Lee, D. Dhody, S. Randriamasy, , "ALTO Traffic Engineering Cost Metrics (work in progress)", October 2014.

[draft-xie-alto-sdn]
H. Xie, T. Tsou, D. Lopez, H. Yin, , "Use Cases for ALTO with Software Defined Networks (work in progress), draft-xie-alto-sdn-extension-use-cases-01", January 2013.

[draft-yang-alto-topology-00]
Y. Yang, , "ALTO Topology Considerations (work in progress)", July 2013.

[sdnrg] "Software Defined Network Research Group, http://trac.tools.ietf.org/group/irtf/trac/wiki/sdnrg", .

[slides-88-alto-5-topology]
G. Bernstein, Y. Lee, Y. Yang, , , "ALTO Topology Service: Use Cases, Requirements and Framework (presentation slides IETF88 ALTO WG session), http://tools.ietf.org/agenda/88/slides/slides-88-alto-5.pdf", November 2013.

Authors' Addresses

Sabine Randriamasy (editor)
Alcatel-Lucent Bell Labs
Route de Villejust
NOZAY 91460
FRANCE

Email: Sabine.Randriamasy@alcatel-lucent.com

Richard Yang
Yale University
51 Prospect st
New Haven, CT 06520
USA

Email: yry@cs.yale.edu

Qin Wu
Huawei
101 Software Avenue, Yuhua District
Nanjing, Jiangsu 210012
China

Email: sunseawq@huawei.com

Lingli Deng
China Mobile
China

Email: denglingli@chinamobile.com

Nico Schwan
Thales Deutschland

Email: nico.schwan@thalesgroup.com