

ALTO Working Group
INTERNET-DRAFT
Intended Status: Standard Track
Expires: December 11, 2015

L. Deng
China Mobile
H. Song
Huawei
S. Kiesel
University of Stuttgart
R. Yang
Yale
Q. Wu
Huawei
June 10, 2015

Extended End Point Properties for Application-Layer Traffic Optimization
draft-deng-alto-p2p-ext-06

Abstract

The purpose of the Application-Layer Traffic Optimization (ALTO) protocol is to provide better-than-random peer selection for P2P networks. The base ALTO protocol focuses, only on providing network topological location information (i.e., network maps and cost maps). However, the peer selection method of an endpoint may also use other properties, such as geographic location. This document defines a framework and an extended set of End Point properties (EP properties) to extend the base ALTO protocol.

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at
<http://www.ietf.org/lid-abstracts.html>

The list of Internet-Draft Shadow Directories can be accessed at
<http://www.ietf.org/shadow.html>

Copyright and License Notice

Copyright (c) 2013 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	4
2. Terminology	4
3. Overview	4
3.1. Guidelines and Methodology	4
3.2. Information flow	5
3.3. Privacy considerations	6
3.3.1. Privacy-Preserving Information Mapping	6
3.2.2. Access Control	7
3.4. Relation with other properties	7
4. Endpoint Extensions	8
4.1. Location-Related Properties	8
4.1.1. Endpoint Property Type: geolocation	8
4.2. Node-related properties	9
4.2.1. End Point Property Type: participating_role	9
4.2.2. End Point Property Type: battery_limited	10
4.2.3. End Point Property Type: local_capacity	11
4.3. Network-Related Properties	11
4.3.1. End Point Property Type: network_access	11
4.3.2. End Point Property Type: forwarding_class	13
4.4. Subscription-Related Properties	14
4.4.1. End Point Property Type: volume_limited	14
4.4.2. End Point Property Type: provisioned_bandwidth	15
5. Security Considerations	16
6. IANA Considerations	16
7. References	17
7.1. Normative References	17
7.2. Informative References	17
Acknowledgements	17
Authors' Addresses	18

1. Introduction

The initial purpose for Application Layer Traffic Optimization (ALTO) protocol [RFC7285] is to provide better than random peer selection for Peer-to-Peer (P2P) networks. It is expected that ALTO can be used in serving a variety of applications and therefore it should be able to provide richer information in terms of End Point properties.

In this document, more EP property extensions are defined to provide guidance for both P2P and other applications in terms of end point selection.

2. Terminology

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].

This document makes use of the ALTO terminology defined in RFC 5693 [RFC5693].

TBA.

3. Overview

It is expected that EP properties reflecting the following list of information can be useful for an ALTO client to provide better user experience or avoid performance degradation:

- o location-related information, the information about the geographic location of the end point.

- o node-related information, the information about the end point's local features, such as software/hardware configuration and the participating role of the end point (e.g. as a end user, or a CDN server, or a P2P cache, etc.).

- o network-related information, the information about the attached network of the end point, such as the type or configuration of the access network (e.g. 2G/3G/4G, WLAN, DSL, etc.) and the information about the network topology (e.g. ASN, Rack-id, etc.).

- o subscription-related information, the information about the service provision agreement between the end point's owner (i.e. the subscriber) and the network provider.

3.1. Guidelines and Methodology

The most basic principle would be to maintain the EP property set to a minimum, which in turn implies two guidelines: non-redundancy and generality.

- o Non-redundancy, refers to the guideline that there is no complete coverage between any two properties.

- o Generality, refers to the guideline that each property should be generally applicable to a group of settings. It is not economic to define a property which is bounded to a single type of application or a single deployment scenario.

In order to make sure that the properties as defined in this document fulfill the above principle and guidelines, we intend to justify each property's definition using the following methodology:

- o Usefulness: there should be a clear motivation and application scenarios that justify the necessity and value for providing such information via EP property enquiry.

- o Non-redundancy: avoid adding a property whose value can be implied by an already defined property or any combination of them. It may be of interest to keep the discussion and suggestions on how to acquire such information via from other already defined EP properties in the document.

- o Case-independency: when designing the concrete information model for the properties, it is suggested to group application/deployment specific information into more general property definitions (with different value for different applications/scenarios) whenever possible.

3.2. Information flow

On the one hand, the same piece of information about a group of candidate endpoints may be acquired by an application in two ways: directly through one-to-many communication of application-specific message exchange with each candidate for flexibility, or indirectly via one-to-one transaction with the ALTO server for efficiency.

On the other hand, EP properties as defined in this document may as well be retrieved and aggregated into the ALTO server in two ways. One is from the endpoint itself, and the other is from the service provider which provides network service to the end point.

Note: There is currently no standardized mechanism by which a peer could publish information about itself into an ALTO server.

Therefore, it is to be decided whether or not if we should include EP properties in this document if their acquisition requires an extension to the base protocol for an endpoint to publish its information directly to the ALTO server.

An endpoint can discover the ALTO server with ALTO discovery mechanisms, and then setup a communication channel with its ALTO server. After that the endpoint property from the endpoint itself can be reported.

The ALTO server can also be configured to access the Network Management System server or other similar servers provided by the network service provider for information about end points, such as subscription related information.

3.3. Privacy considerations

Privacy considerations is a general concern for almost all EP properties, as they are by definition more stationary information regarding a specific end point.

However, each end point may have different concerns or sensitive preference over a specific EP property. For example, endpoint property regarding the service role of the endpoint, serving nodes deployed by the ISP or third party service provider, such like P2P caching server, or CDN node, may have different considerations over whether a piece of information is private or not. Therefore, it may be necessary to provide a mechanism to accommodate this type of individual customization by providing a channel for an end point to explicitly indicate this information based on its own preference.

More general, it is expected that the privacy level of a specific EP property is dependent on the nature of the information (i.e. the EP property), the type of the subscriber (i.e. the user who owns the end point in question), the type of the application (i.e. the ALTO client who is requesting the EP property) and the policy of the ISP (i.e. the owner of the ALTO server who is able to do information collection from the end points and determine how the the information is exposed to the requesting application).

Fortunately, there are generally applicable schemes to be used to address the privacy protection concerns, which may be applicable to a group of EP properties and can be configured by the ISP or the EP subscriber. In this section, several general schemes are introduced, whose application to each EP property is elaborated later in following sections.

3.3.1. Privacy-Preserving Information Mapping

On the one hand, the privacy concern is unnecessary if the specific endpoint property can also be measured/disclosed in another way. The privacy concern regarding to the accurate information of the endpoint would be alleviated if using relative numbers to rank them. For deployment considerations, it is also possible for each endpoint to make the choice whether to disclose the relative information or not, but an incentive could be used to encourage the disclosure when it is beneficial to the application.

In other words, in order to preserve the privacy of a piece of information, different data types can be defined via information mapping. In particular, in this document, each property is defined as a JSON object [RFC4627], which contains a dynamic typing attribute "content" as well as two deterministic attributes, "name" and "precision".

The "name" attribute is a string, whose value is the name of the property. The "precision" attribute is also a string, whose value comes from an attribute-dependent set. Depending on the value of the property's "precision" attribute, its "content" attribute can be a string, number, boolean or another object.

In this document, in order to define an EP property as a JSON object, we specify:

- o the string value of its "name" attribute;
- o the value set of its "precision" attribute; and
- o the definitions of its "content" attribute for each "precision".

A special string value "" for "precision" attribute is used to indicate that an EP property, which is not privacy sensitive or using information mapping, has no precision-dynamic "content" definition.

3.2.2. Access Control

On the other hand, access control to sensitive property information may also be used to mitigate the privacy concern of a defined property. Even greater flexibility can be delivered by access control at the discretion of both the network operator and the individual subscriber, which is deployment specific and out of scope for the general discussion within this document.

3.4. Relation with other properties

Endpoint information can be extremely dynamic or relatively static. Currently, this specification does not intend to provide any real-

time properties such as the available bandwidth from the endpoint [I-D.draft-wu-alto-te-metrics], whose value is subject to frequent changes and hence requires a measurement-based exposure scheme.

The basic end point properties as defined in this document, serves as a basis for the property namespace to be used to derive PID properties [I-D.draft-roome-alto-pid-properties] for the corresponding peer group, when the direct enquiry for the information per end point is not efficient or economic for the ALTO client.

4. Endpoint Extensions

This document defines new endpoint property types for the ALTO protocol [RFC 7285].

4.1. Location-Related Properties

4.1.1. Endpoint Property Type: geolocation

It is believed that the information about an individual endpoint's geo-location is of value to a variety of applications. However, it is also well accepted that geolocation of an endpoint is likely to be considered as a private piece of information to the subscriber, and therefore should be protected against undesirable privacy intrusion.

Moreover, in a data-center, the relative location of a serving node may be of interest to an ALTO client, where much finer-grained information (e.g. the hosting physical server or rack number) are relevant and can be dynamically updated by either a live migration of a serving node contained in a virtualization container or a traffic handover between active and standby instances during an HA/LB switch-off.

To this end, an EP property is defined as a JSON object, with the name "geolocation", whose "content" definition is actually dependent on the "precision" attribute, which in turn is a JSON string whose value belongs to the following JSON array:

```
geolocation_precision_set = ["countrycode", "boundingbox", "circle", "dc"]
```

If the "precision" attribute of the "geolocation" property of an endpoint is "countrycode", the following "content" attribute is defined as the ISO 3166 two-letter country codes of the region the endpoint resides in, as a JSON string.

If the "precision" attribute of the "geolocation" property of an endpoint is "boundingbox", the following "content" attribute is defined as a four-element JSON object "bounding_box":

```
bounding_box = {  
    "latul" : number,  
    "longul" : number,  
    "latbr" : number,  
    "longbr" : number  
}
```

If the "precision" attribute of the "geolocation" property of an endpoint is "circle", the following "content" attribute is defined as a three-element JSON object "circle_location":

```
circle_location = {  
    "latc" : number,  
    "longc" : number,  
    "radius" : number  
}
```

If the "precision" attribute of the "geolocation" property of an endpoint is "dc-location", the following "content" attribute is defined as a four-element JSON object "dc-location":

```
dc-location = {  
    "server-id" : number,    "rack-id" : number,  
}
```

4.2. Node-related properties

4.2.1. End Point Property Type: participating_role

Different types of end points have different roles or participating policies for a given application, which can be explored in making a better decision when choosing a serving node. For example, as described in [I-D.draft-deng-alto-p2pcache], P2P caching node can also act as p2p peers in a p2p network. If a p2p caching peer is located near the edge of the network, it will reduce the backbone traffic, as well as the uploading traffic. [RFC7069] provides one example of such caching nodes. P2P caching peers are usually expected to be given higher priority than the ordinary peers for serving a content request so as to optimize the network traffic. So it's necessary for the End Point property to support this indication.

In general, the end points which belong to different participating

parties (subscriber, ISP, or ICP) within an application's service transaction demonstrate different role/policies.

It is straightforward for an ISP to acquire the information of an end point's participation role from its local record for its subscribers, its local or third party infrastructure for a given application.

To this end, an EP property is defined as a JSON object, with the name "participating_role", whose "precision" attribute is set to "" and its "content" attribute is defined as a JSON string, whose value belongs to the following array:

```
participating_role_set=["user", "cache", "super_node"]
```

In other words, the "participating_role" property is defined as follows:

```
participating_role : {  
    "precision": "",  
    "content": ["user", "cache", "super_node"]  
}
```

4.2.2. End Point Property Type: battery_limited

Another important End Point property that will impact peer selection is what kind of power supply the peer has. It can be either the electric power or the battery supply.

And for most of the time, it is safe to bet that electric power supplied nodes would stay online longer than those battery supplied nodes, while battery powered devices are usually less willing to act as super peer, relay, etc.

And most of the nowadays intelligent equipments are aware of their power supply type. But it is necessary that the power supply of a peer can be queried through some method no matter whether or not it is limited by its battery.

To this end, an EP property is defined as a JSON object, with the name "battery_limited", whose "precision" attribute is set to "" and its "content" attribute is defined as a boolean, is either "true" or "false".

If the peer in question is actually battery-limited, the value of this property with respect to the peer is set to "true".

In other words, the "content" attribute of the "battery_limited" property is defined as a JSON boolean, "true" for a battery supplied

end point, or "false" for an electricity supplied end point or for an end point with an unknown power supply type.

```
"battery_limited": {  
    "precision": "",  
    "content": true/false  
}
```

4.2.3. End Point Property Type: local_capacity

For resource-consuming applications, it would be helpful to know the local capacity (e.g., in terms of computing, storage, and networking) of an end point before it is selected.

In other words, the "local_capacity" property is defined as a JSON object, as follows:

```
"local_capacity": {  
    "precision": "",  
    "content": {  
        "CPU": {  
            "volume": integer,  
            "meter": string  
        },  
        "memory": {  
            "volume": integer,  
            "meter": string  
        },  
        "storage": {  
            "volume": integer,  
            "meter": string  
        }  
    }  
}
```

4.3. Network-Related Properties

4.3.1. End Point Property Type: network_access

One important End Point property that will impact peer selection is the type of the node's access network.

Note: There is remaining doubt on whether or not this property is needed, since at least part of the information it reflects, for instance, the end point's provisioned bandwidth, is defined and exposed by other properties.

For instance, a mobile subscriber's access network can be cellular (2G, 3G, or 4G). Take another example of a node owned by a home subscriber, the type of its access network can be DSL, FTTB, or FTTH.

Different type of access network gives a clear indication on both the amount and the technology of the provisioned resources (e.g. the shared/guaranteed bandwidth, the interval for physical channel scheduling, etc.)

Moreover, one may prefer to specify a special access type for a node deployed in a data center too, because it is likely to be more robust, and have more network resources than either mobile or home users.

Hence application may have its own algorithm for peer selection or traffic rendering if the node access type information can be provided via an End Point property. The value for this property can be enumerated as "adsl", "ftth", "fttb", "dc", and etc.

In case that the end point has its own privacy concerns in revealing its access network type directly to potentially distrusted applications through ALTO, another indirect way of exposing the similar information can be used by "access_preference" as per ISP's judgement.

In essence, an ISP assigned "access_preference" property for the end points gives the network operator a chance to say which end point's link is "better" without having to tell what the actual criterion is.

The value for this property (defined as integer) can be set by the ISP of the ALTO server, based on its own relative preference to different network access types. A peer with the higher value is more preferable than another peer with the lower value.

For example, an ISP could use the following setting for now:

```
1 = DSL; 10 = FTTB; 12 = FTTH; 50 = DC;
```

and add "100=new_technology", when some new technology better than FTTH appears later.

To this end, an EP property is defined as a JSON object with the name "network_access", with two different values for "precision"

```
network_precision_set=["technology", "rank"]
```

In other words, the "content" of the "network_access" property is dependent on the value of its "precision" attribute.

If the value of "precision" is "technology", the following "content" attribute is defined as a JSON string, whose value belongs to the following array:

```
network_access_set = ["adsl", "ftth", "fttb", "dc", "cellular"]
```

If the value of "precision" is "rank", the following "content" attribute is defined as a JSON number, whose value indicates the relative preference over the end point in question, in terms of its access network. The end point with a higher number is more preferable to another end point with a lower number.

In summary, the "network_access" property is defined as a JSON object, as follows:

```
"network_access": {  
  "precision": "technology",  
  "content":["adsl", "ftth", "fttb", "dc", "cellular"]  
}
```

```
"network_access": {  
  "precision": "ranking",  
  "content": number  
}
```

Note: There is concern about undesirable privacy leakage via network_access properties to distrusted ALTO clients. In such cases, according to the definitions above, either the endpoint itself or the ISP who is running the ALTO server can either specify an access control policy to prevent undesirable exposure to specific ALTO clients or use a privacy preserving mapping from the raw description of access technologies to a number of abstract relative ranking information instead. Moreover, the endpoint or the ISP might choose to use another subscription related property "provisioned_bandwidth" (defined later in Section 4.4.2) instead of "network_access".

4.3.2. End Point Property Type: forwarding_class

As suggested for the NFV use-case, the End Point property "forwarding_class" is meant to indicate the type of forwarding class the end point or network supports.

Forwarding classes can be thought of as output queues. For a

classifier to assign an output queue to a packet, it must associate the packet with one of the following forwarding classes:

- o Expedited forwarding (EF), provides a low-loss, low-latency, low-jitter, assured bandwidth, end-to-end service.
- o Assured forwarding (AF), provides a group of values you can define and includes four subclasses: AF1, AF2, AF3, and AF4, each with three drop probabilities: low, medium, and high.
- o Best effort (BE), provides no service profile. For the best effort forwarding class, loss priority is typically not carried in a class-of-service (CoS) value.
- o Network control (NC), is typically high priority because it supports protocol control.

Hence, the "content" of the "forwarding_class" property is defined as a JSON string, whose value belongs to the following array:

```
forwarding_class_set = ["expedited", "assured", "network control",  
"best effort"]
```

In summary, the "forwarding_class" property is defined as a JSON object, as follows:

```
"forwarding_class": {  
  "precision": "",  
  "content": ["expedited", "assured", "network control", "best effort"]  
}
```

4.4. Subscription-Related Properties

4.4.1. End Point Property Type: volume_limited

Many wireless operators offer low-cost plans, which limit the amount of data to be transmitted within a month to some gigabytes. After that they will throttle the subscriber's bandwidth or charge extra money. Hosts with such a tariff, could be tagged by another End Point property "volume_limited" and should be avoided for peer selection to serve other peers.

The "content" value for this property (defined as a boolean) is either "true" or "false". If a peer is constrained by such a subscription plan, the value of this property with respect to the peer is set to "true".

In other words, the "volume_limited" property is defined as a JSON object with a boolean "content", "true" for an end point with such a limited data plan, or "false" for an point with unlimited or unknown data plan.

```
"volume_limited": {  
    "precision": "",  
    "content": true/false  
}
```

4.4.2. End Point Property Type: `provisioned_bandwidth`

For applications seeking for a candidate peer for uploading services, the end point's uploading bandwidth is essential for the selection.

While it is straightforward for one to expose the accurate information over an end point's bandwidth capability, the subscriber of the end point might consider it a piece of private information.

On the other hand, it is suggested that the ISP can also choose to expose its relative preference in terms of the end point's provisioned bandwidth; this ensures better load balancing within the network by avoiding undesirable hot spots caused by competition from applications for the handful most provisioned end points.

Therefore, the "provisioned_bandwidth" property is defined as a JSON object, whose "content" definition is actually dependent on the "precision" attribute, which in turn is a JSON string whose values belong to the following JSON array:

```
provisioned_bandwidth_precision_set = ["raw", "ranking"]
```

If the "precision" attribute of the "provisioned_bandwidth" property of an end point is "raw", the following "content" is filled with the accurate value of the provisioned bandwidth, as a JSON object "provisioned_bandwidth_value" with two elements:

```
provisioned_bandwidth_value = {  
    "value" : number,  
    "metric" : ["GB", "MB", "KB", "Gb", "Mb", "Kb"]  
}
```

If the "precision" attributed of the "provisioned_bandwidth" property of an end point is "ranking", the following "content" is filled with the relative ranking of the end point's provisioned bandwidth assigned by the ISP, which in turn is a JSON number where higher number indicating more preference.

In summary, the "provisioned_bandwidth" property is defined as a JSON object as follows:

```
"provisioned_bandwidth": {  
  "precision": "raw",  
  "content": {  
    "value": number,  
    "metric": ["GB", "MB", "KB", "Gb", "Mb", "Kb"]  
  }  
}
```

```
"provisioned_bandwidth": {  
  "precision": "ranking",  
  "content": number,  
}
```

5. Security Considerations

TBA.

6. IANA Considerations

This document adds the following new End Point property types to the existing registry created by ALTO protocol [RFC7285].

TBA.

7. References

7.1. Normative References

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

[RFC7285] Alimi, R., Penno, R., and Y. Yang, "ALTO Protocol", RFC7285, March 2014.

7.2. Informative References

[I-D.draft-deng-alto-p2pcache] Deng, L., Chen, W., and Q. Yi, "Considerations for ALTO with network-deployed P2P caches", draft-deng-alto-p2pcache-03 (work in progress), February 2014.

[RFC7069] Alimi, R., Rahman, A., Kutscher, D., Yang, Y., Song, H., and K. Pentikousis, "DECoupled Application Data Enroute (DECADE)", RFC 7069, November 2013.

[I-D.draft-roome-alto-pid-properties] Roome, W. and Yang, R., "PID Property Extension for ALTO Protocol", draft-roome-alto-pid-properties-01 (work in progress), February 2014.

[I-D.draft-wu-alto-te-metrics] Wu, Q., Yang, R., Lee, Y., and Randriamasy, S., "ALTO Traffic Engineering Cost Metrics", draft-wu-alto-te-metrics-03 (work-in-progress), June 2014.

Acknowledgements

The authors would like to thank, Michael Scarf, Vijay Gurbani, Reinaldo Penno, Sabine Randriamasy and Qiao Fu for their review and valuable comments.

Authors' Addresses

Lingli Deng
China Mobile
China

Email: denglingli@chinamobile.com

Haibin Song
Huawei
China

Email: haibin.song@huawei.com

Sebastian Kiesel
University of Stuttgart, Computing Center
Germany

Email: ietf-alto@skiesel.de

Richard Yang
Y. Richard Yang
Yale University

Email: yry@cs.yale.edu

Qin Wu
Huawei
China

Email: sunseawq@huawei.com