

ALTO WG
Internet-Draft
Intended status: Standards Track
Expires: April 11, 2015

W. Roome
Alcatel-Lucent
October 8, 2014

Resource Attributes for ALTO Protocol
draft-roome-alto-resource-attr-00

Abstract

This document extends the Application-Layer Traffic Optimization (ALTO) Protocol [RFC7285] by defining additional descriptive attributes for the resources offered by an ALTO Server.

Status of this Memo

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

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on April 11, 2015.

Copyright Notice

Copyright (c) 2014 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	3
1.1.	Attributes Versus Capabilities	3
2.	Changes To ALTO Protocol	4
2.1.	IRD Entries	4
3.	Resource Attributes	5
3.1.	Attribute Names	5
3.2.	Attribute Values	6
3.3.	Proposed Resource Attributes	6
3.4.	Authoritative Endpoint Sets	7
4.	Use Cases	8
5.	Alternatives And Discussion (Delete if draft is published)	8
6.	Security Considerations	9
7.	IANA Considerations	9
8.	References	9
	Author's Address	9

1. Introduction

In the ALTO Protocol [RFC7285], the Information Resource Directory (IRD) defines the resources, or services, offered by an ALTO server. An IRD contains one entry for each resource. Each entry provides the information a client needs to use that resource: the URI for the resource, the type of data the server returns, the type of data the client sends (if any), the identifiers of any resources on which this resource depends, and optional capabilities for the resource.

However, resources have many additional attributes. For example, a Network Map resource maps PID (Provider-defined Identifier) names to network address prefixes. A small Network Map resource might define 20 PIDs and a total of 100 prefixes, while a large Network Map might have 5,000 PIDs and 500,000 prefixes. A client might like to know the size of the Network Map before retrieving it, but the IRD gives no hint.

Also, an ALTO server provided by an Internet Service Provider (ISP) might have detailed information for network addresses managed by that ISP, but only coarse information (or no information at all) for other network addresses. If an ALTO Client needs cost information for a particular network address, and the client knows about several different ALTO Servers, the client would prefer to use the ALTO Server with the best data for that address. But the IRD entries give no hint as to the best server; the client would have to try each server and guess as to which was the most accurate.

This document defines a framework for declaring such attributes in IRD entries, and defines an initial set of attributes.

1.1. Attributes Versus Capabilities

The ALTO Protocol already defines a "capabilities" section for IRD entries, so one alternative is simply to define additional capabilities. We prefer not to do that for several reasons:

- o Capabilities are always defined; the ALTO protocol either requires the IRD entry to specify a value, or else defines a default. Attributes are optional, and do not have defaults.
- o Capability names are defined by the ALTO protocol; adding a new capability requires a new RFC. Attribute names are managed by IANA, so new attributes can be added without extending the protocol.
- o A client can deduce attributes by using the resource. Capabilities cannot be deduced.

- o Attribute values are approximate, and/or may change frequently. Capabilities are accurate, and should not change unless the resource is redefined. For example, if a Cost Map attribute says the map has 500 cost points, that means that was the size when the client retrieved the IRD. If the client retrieves the map a few minutes later, the size might be different. However, if a capability says a Cost Maps returns a "routingcost" cost metric, it is an error if the resource returns a "hopcount" instead.

2. Changes To ALTO Protocol

2.1. IRD Entries

Resources attributes are defined by a new field, named "attributes", at the same level as the existing "uri" and "media-type" fields in IRD entries. That is, in [RFC7285] Sec. 9.2.2, IRDResourceEntry is revised as:

```
    object {
      JSONString    uri;
      JSONString    media-type;
      [JSONString   accepts;]
      [Capabilities capabilities;]
      [ResourceID   uses<0..*>;]
      [Attributes   attributes;]      // Added
    } IRDResourceEntry;

    object {
      ...
    } Attributes;      // Added
```

Here is an example of an IRD with attributes:

```
...
"resources" : {
  "my-default-network-map" : {
    "uri" : "http://alto.example.com/networkmap",
    "media-type" : "application/alto-networkmap+json",
    "attributes" : {
      "pid-count-current" : 50,
      "prefix-count-current" : 150,
      "authoritative-prefixes" : {
        "ipv4": [ "1.2.0.0/16", "3.4.0.0/16" ]
      },
      "authoritative-exclusions" : {
        "ipv4": [ "1.2.1.0/25" ]
      }
    }
  },
  "numerical-routing-cost-map" : {
    "uri" : "http://alto.example.com/costmap/num/routingcost",
    "media-type" : "application/alto-costmap+json",
    "capabilities" : {
      "cost-type-names" : [ "num-routing" ]
    },
    "uses" : [ "my-default-network-map" ],
    "attributes" : {
      "cost-count-current" : 500,
      "cost-count-range" : [ 250, 750 ]
    }
  },
  ...
}
```

3. Resource Attributes

3.1. Attribute Names

An attribute name is encoded as a string. The string MUST be no more than 32 characters, and it MUST NOT contain characters other than US-ASCII alphanumeric characters (U+0030-U+0039, U+0041-U+005A, and U+0061-U+007A), the hyphen ('-', U+002D), the colon (':', U+003A), the low line ('_', U+005F), or the '.' separator (U+002E). The '.' separator is reserved for future use and MUST NOT be used unless specifically indicated by a companion or extension document.

Identifiers prefixed with "priv:" are reserved for Private Use [RFC5226] without a need to register with IANA. All other identifiers MUST be registered in the "ALTO Resource Attribute Registry" (see Section 7). For an identifier with the "priv:"

prefix, an additional string (e.g., company identifier or random string) MUST follow (i.e., "priv:" only is not a valid identifier) to reduce potential collisions.

Section 3.3 defines an initial set of Resource Attributes.

3.2. Attribute Values

The type of an attribute value depends on the attribute. When registering an attribute with IANA, the JSON value type MUST be specified.

3.3. Proposed Resource Attributes

We propose the following Resource Attributes:

pid-count-current:

Resource type: Network Map

Value type: JSON Number

Semantics: The number of PIDs in the current Network Map.

prefix-count-current:

Resource type: Network Map

Value type: JSON Number

Semantics: The total number of prefixes in the current Network Map.

pid-count-range:

Resource type: Network Map

Value type: JSON Array with two JSON Numbers

Semantics: Low and high values for the number of PIDs in the Network Map. This range SHOULD cover the expected size of the map for the foreseeable future.

prefix-count-range:

Resource type: Network Map

Value type: JSON Array with two JSON Numbers

Semantics: Low and high values for the total number of prefixes in the Network Map. This range SHOULD cover the expected size of the map for the foreseeable future.

authoritative-prefixes:

Resource types: Network Map and Endpoint Cost Map

Value type: A JSON Object of type NetworkMapData (Section 11.2.1.6 of [RFC7285])

This resource is authoritative (see Section 3.4) for all endpoints which are covered by a prefix in this set, and which are not covered by a prefix in the "authoritative-exclusions"

set.

authoritative-exclusions:

Resource types: Network Map and Endpoint Cost Map
Value type: A JSON Object of type NetworkMapData (Section 11.2.1.6 of [RFC7285])
See the "authoritative-prefixes" resource, above.

cost-count-current:

Resource type: Cost Map
Value type: JSON Number
Semantics: The number of cost points in the current Cost Map.
Note that although a Cost Map is an NxN matrix, where N is the number of PIDs, a Cost Map is not required to define a cost for every {source,destination} pair.

cost-count-range

Resource type: Cost Map
Value type: JSON Array with two JSON Numbers
Semantics: Low and high values for the total number of cost points in the Cost Map. This range SHOULD cover the expected size of the map for the foreseeable future.

3.4. Authoritative Endpoint Sets

The set of endpoints defined by the "authoritative-prefixes" and "authoritative-exclusions" attributes are the endpoints for which a resource provides authoritative data. For an ALTO Server provided by an ISP, this will usually be the endpoints of the ISP's customers.

While we do not rigorously define "authoritative", in general if a resource is authoritative for a set of endpoints, then:

- o No other ALTO Server has more accurate data for those endpoints.
- o The ALTO Server SHOULD fully specify the costs for those endpoints. That is, for every endpoint EA in the authoritative set and every other endpoint EX, the associated cost map service(s) SHOULD define the costs from EA to EX and from EX to EA.
- o PIDs covering endpoints in the authoritative set SHOULD be more detailed (smaller, finer-grained, etc.) than PIDs for other endpoints.

We define the authoritative set with two sets of prefixes, one inclusive, the other exclusive, for the following reason. The most likely case is that an authoritative set represents the addresses of

an ISP's customers. In this case, the inclusive set would be the prefixes that have been allocated to the ISP by the appropriate authority. If customers have migrated to other ISPs and taken their addresses, those prefixes form the exclusionary set. While the exclusionary set is not necessary, it provides an efficient way to represent a large block of addresses with a few exceptions. This two-level model seems to be a good balance between efficiency and simplicity.

4. Use Cases

Here are brief descriptions of several ways in which clients can take advantage of Resource Attributes:

- o Advance knowledge of the size of a network or cost map allows an ALTO client to allocate sufficient space to hold the map, or to select the appropriate way to store it. For example, a Cost Map is an $N \times N$ matrix, where N is the number of PIDs. But it may be sparse. If the number of costs in the Cost Map is considerably less than the square of the number of PIDs in the associated Network Map, then the client knows in advance that the Cost Map is likely to be sparse, and can use a sparse matrix object. On the other hand, if the number of costs is close to the square of the number of PIDs, the client knows that there is no advantage in using sparse matrix techniques.
- o A client such as a P2P tracker can use the authoritative sets to select the appropriate ALTO Server. For example, suppose the tracker has obtained the IRDs for a number of ISP-maintained ALTO Servers. When a peer at address E asks the tracker for a set of peers, the tracker would locate the Network Map (or Endpoint Cost Service) whose authoritative set includes E . This will mostly be the ALTO Server for the requesting peer's ISP. The tracker uses that ALTO resource to evaluate costs between E and the other peers, and returns the peers with the lowest costs.
- o If an ALTO Server offers several different Network Maps, a client can use the network map size and/or authoritative sets to select the appropriate Network Map.

5. Alternatives And Discussion (Delete if draft is published)

Here a few points for which I welcome discussion:

- o Can anyone suggest a better name than "attributes"? I considered "properties", but that conflicts with Endpoint Properties.

Another possibility is "meta", but we have overloaded that name.

- o I don't like the term "authoritive" to describe the endpoints on which a Network Map is centered. However, I can't think of anything better, and it does seem to accurately reflect the intent. Can anyone suggest an alternative?

6. Security Considerations

Because attributes should not provide any information which cannot be deduced from the data already provided by an ALTO Server, this extension does not introduce any security considerations not already present in the ALTO Protocol.

7. IANA Considerations

This document defines a registry for ALTO Resource Attributes. When a new ALTO Resource Attribute is defined and accepted by the ALTO working group, requests for IANA registration MUST include the identifier, applicable resource type(s), JSON type and intended semantics.

Section 3.3 defines the initial set of Resource Attributes.

8. References

- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", RFC 5226, BCP 26, May 2008.
- [RFC7285] Almi, R., Penno, R., Yang, Y., Kiesel, S., Previdi, S., Roome, W., Shalunov, S., and R. Woundy, "Application-Layer Traffic Optimization (ALTO) Protocol", RFC 7285, September 2014.

Author's Address

Wendy Roome
Alcatel-Lucent/Bell Labs
600 Mountain Ave, Rm 3B-324
Murray Hill, NJ 07974
USA

Phone: +1-908-582-7974
Email: w.roome@alcatel-lucent.com