

ALTO Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 30, 2015

U. Rauschenbach
Nokia Networks
October 27, 2014

ALTO in wireless access networks
draft-rauschenbach-alto-wireless-access-00

Abstract

The Application-Layer Traffic Optimization (ALTO) specification defines the concept of Provider-defined Identifiers (PIDs) as the aggregation of network endpoints for network nodes. Each PID can be associated with a set of endpoint addresses, e.g. aggregating endpoints that use the same network access point.

This document focuses on mobile networks and introduces proposes the toidea of using use cells of cellular access networks as an aggregation points. This allows applications to make decisions based on the path cost of using the current cell as a network attachment point, or to even choose which network access points network attachment point to select. Use cases are described which can benefit from this. The draft elaborates on possible ALTO modifications enabling such use cases. The intent of the draft is to start discussion on the topic.

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 [RFC2119].

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 30, 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
2. Problem statement
3. Use cases
 - 3.1. Cell as aggregation point
 - 3.2. Passive reaction to handovers
 - 3.2.1. Cost calendar to extend battery life for background tasks
 - 3.2.2. Cost calendar to optimize application sessions
 - 3.3. Connection management and traffic offload
4. Requirements and solution ideas
5. References
- Author's Address

1. Introduction

The Application-Layer Traffic Optimization (ALTO) [RFC7285] specification allows providing information about a network to support applications in making decisions regarding the most optimal use of the network. Originally geared towards supporting peer-to-peer applications in IP networks, the ALTO WG has recently broadened its scope.

ALTO defines the concept of Provider-defined Identifiers (PIDs) to aggregate network endpoints. Each PID can be associated with a set of endpoint addresses. The concept of PIDs is generic and potentially allows many types of endpoint addresses - the ALTO specification mentions IP addresses, overlay IDs and MAC addresses. However, RFC 7285 only defines address types for IPv4 and IPv6 addresses and leaves it to additional documents to define additional types.

This document introduces a number of use cases occurring in wireless networks which benefit from associating a cell or access point with a PID. In such scenarios, the cell or access point through which network attachment occurs or may occur in the future needs to be identifiable by the ALTO client. Wireless access points provide information such as cell ID to identify them. However, currently ALTO provides only IP addresses for endpoints, and arbitrary strings as PID values. Allowing an ALTO map to provide information about an actual cell or access point can benefit use cases when the terminal needs metrics describing the path cost of using a particular cell or access point before the terminal attaches to the network via that cell or access point, as well as for using a cell ID to aggregate information about all endpoints attached to that cell or access point.

2. Problem statement

ALTO defines the concept of PIDs which can be used to aggregate network nodes. A PID can aggregate e.g. cover all nodes connected to the network via a particular POP. It is assumed that in fixed networks, the actual POP via which an endpoint is connected to the network rarely changes. In cellular wireless networks, the situation is different. Network access happens through cells. Nodes frequently change cells when they move. Traffic conditions and other costs may vary widely from cell to cell. Therefore, it makes sense to include information about actual cells through which network attachment happens in the ALTO maps. However, currently this is not possible in the ALTO data entities. Similarly, access to WiFi networks happens through wireless access points which are currently also not considered in ALTO.

3. Use cases

3.1. Cell as aggregation point

Wireless terminals connect to a cellular network through different cells, and to Wi-Fi networks through different wireless access points. The quality of network access may vary greatly from connection point to connection point, e.g. due to congestion or the differing capacity of backhaul (connection point is used to refer to both cells and access points from now on). In particular in cellular networks, the connection point actually used frequently changes.

If the ALTO information model would be able to identify the connection point with a PID, then different costs could be provided for different map paths to the connection points. Also, there may potentially be a large number of cellular subscribers connected to a typical cell. It may not matter which endpoint addresses they use, but only via which connection point they are attached to the network. Exploiting this may greatly reduce map size as numerous endpoints under a connection point can be aggregated into a PID, and also reduce the number of map updates as endpoints move so that they can be attached via different connection points due to endpoint movement or handovers.

3.2. Passive reaction to handovers

Applications can benefit from knowledge or assumption about the cells a terminal is connected to or will with some probability be connected to in the near future.

Applications may know which cells a terminal typically connects throughout a day e.g. from observing daily commute patterns. They may also be provided candidate cells with favorable conditions by including in the ALTO response a set of non-congested cells in the vicinity of a queried cell. Once the terminal connects to such a cell while moving, the application can execute tasks which require e.g. higher throughput.

Also, if non-congested and congested cells overlap (such as e.g. a non-congested small cell or WiFi hotspot and a congested macro cell) and the terminal is in motion, an application can make assumptions about future congestion and prepare / react accordingly.

Finally, applications may obtain candidate cells by accessing information about neighboring cells through the APIs provided by modern smartphones.

Based on such knowledge, the following two classes of use cases can be distinguished.

3.2.1. Cost calendar to extend battery life for background tasks

Large fractions of Internet traffic today occur over wireless systems. Not all traffic is real-time and many tasks can be performed in the background. In particular, clients that download large volumes of data (e.g., mail applications or social network clients) can choose the time when they will actually download large items such as attachments or video clips.

If a cost calendar does include information that a client can use to compute the achievable throughput of a service via different cells (e.g. based on information about available bandwidth or congestion information per cell), it can initiate or schedule large downloads to take place when it can connect to a cell that currently allows high throughput.

If a high throughput cell is used for background downloads, the battery life can be extended as the terminal spends less time in

battery-consuming states with the processor busy and the radio on and instead spends more time in battery-conserving idle modes.

3.2.2. Cost calendar to optimize application sessions

Cost and other metrics usually fluctuate over time e.g. based on the day of week. Such knowledge can be used to provide hints to applications how to adapt proactively. For instance, during (or prior to) user mobility (e.g. a commute), an application may make use of such heuristic information to prefetch farther in advance (beyond just in time) more items within an ongoing streaming session at times of low congestion, prior to it entering time intervals or cells with higher congestion etc. Such knowledge can also contribute to picking a suitable video rate in progressive streaming. In a video telephony session, applications may proactively switch to audio-only calls in congested areas. This use case is somewhat related to the previous one; however, it focuses on optimizing the user experience of foreground tasks, rather than on scheduling background tasks.

3.3. Connection management and traffic offload

The previous use cases have focussed on the application reacting on a change in network access point. However, often various connection points of different wireless networks are available to which an endpoint may actively attach. As mobile endpoints get more intelligent, functions for connection management and traffic offload can consider not only the quality of the radio connection when doing access selection and traffic offload, but also other properties such as cost parameters and additional metrics provided by ALTO. Also, when multiple different wireless access options are available, traffic may be distributed between these options based on ALTO cost maps so the network can be used more optimally. However, this requires that the wireless terminal can associate ALTO map entries with actual cell identifiers in the wireless networks.

4. Requirements and solution ideas

This section elaborates solution ideas as a starting point for the technical discussion.

Aggregation in ALTO is modeled by a PID which represents a group of nodes. A PID is identified by a provider-defined string with no general meaning, and an endpoint is identified by its IP address.

Cellular networks use a cell identifier (cell ID) to identify the cells through which network attachment happens. In Public Land Mobile Networks (PLMNs) based on E-UTRAN [TS36.311], the Cell ID is composed of the PLMN ID and the ECI, which in turn are composite identifiers. The PLMN ID (Public Land Mobile Network Identifier) identifies a wireless communications system. It consists of the 3 digit Mobile Country Code (MCC) and the 3 digit Mobile Network Code (MNC). ECI is the E-UTRAN Cell Identifier which identifies a Cell within a PLMN. An ECI is composed of eNB ID and Cell ID. The eNB ID (20 bits) is the eNodeB Identifier which identifies an eNB within a PLMN. The Cell ID (8 bits) identifies a (sub)cell within a particular eNodeB.

WiFi-based wireless networks provide data such as SSID, access point MAC address and other values that allow to identify the network access point. Note that the Hotspot 2.0 specification [HS2.0] allows a terminal querying a number of parameters prior to connecting to an access point (such as Venue Name, Roaming Consortium, IP Address Type Availability, 3GPP Cellular Network, Domain Name, Hotspot Query List, Hotspot Capability List and others). The terminal searching for a particular network access can retrieve the information elements through the IEEE 802.11 ANQP from the access point prior to connection establishment. However, once connected to an access

point, the terminal has to retrieve the information of other access points in the area by virtually detaching from the access point and running access network discovery over the air to the other access points. It would be beneficial when the terminal would be able to retrieve the information of adjacent access points and access networks by way of e.g. an ALTO query over the existing connection.

In order to enable the use cases defined above in ALTO, the IDs which identify the wireless connection point would have to be added to the ALTO data model.

To support the use cases above, two things are required:

1. To enable querying properties of a certain cell such as bandwidth or degree of congestion. Such queries could be supported through the mechanism of endpoint property queries (RFC7285 section 11.4). Possible solutions are:
 - * Solution A: The cellular access point would need to be defined as an additional endpoint inside the group aggregated by a PID, i.e. a new address type that represents a cell ID would be required.
 - * Solution B: A property query would be needed for the PID, plus a property that represents the cell identifier.
2. To enable cells as identifiable aggregation points, the notion of a PID (which is now just a string) would need to be extended by including cell identification. Assigning different PIDs to different cells would then allow to create different cost map sections for paths of access through different cells (e.g., costs would be higher for access through congested cells than for access through non-congested cells).
 - * Solution: A PID would need to be extended by a cell identifier (related to Solution B above)

5. References

- [HS2.0] The WiFi Alliance, "Hotspot 2.0 (Release 2) Technical Specification, Version 1.0.0", August 2014.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC7285] Alimi, 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.
- [TS36.311] Third Generation Partnership Project, "3GPP TS 36.331 V12: Technical Specification Group Radio Access Network; Evolved Universal Terrestrial Radio Access (E-UTRA); Radio Resource Control (RRC); Protocol specification (Release 12)", September 2014.

Author's Address

Uwe Rauschenbach
Nokia Networks

Email: uwe.rauschenbach@nsn.com