

ALTO  
Internet-Draft  
Intended status: Informational  
Expires: August 17, 2014

L. Deng  
W. Chen  
Q. Yi  
China Mobile  
Y. Zhang  
Chinese Academy of Sciences  
February 13, 2014

Considerations for ALTO with network-deployed P2P caches  
draft-deng-alto-p2pcache-03.txt

## Abstract

Uploading from peers located in a public WLAN hotspot has been reported to severely impact other current users' experiences, and raised caution from the operator's side who are willing to increasingly participate in building public WLAN facilities to offload the explosive mobile data traffic from cellular networks. Cooperation between the network operator and the P2P service providers in form of intra-domain P2P caches is expected to be an effective mechanism to solve the problem. This draft introduces considerations on ALTO deployment in terms of P2P caches and discusses potential extensions to the ALTO protocol to standardize this mutual cooperation.

## 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 August 17, 2014.

## Copyright Notice

Copyright (c) 2014 IETF Trust and the persons identified as the document authors. All rights reserved.

## Internet-Draft Considerations for ALTO with network-deployed February 2014

This document is subject to [BCP 78](http://trustee.ietf.org/license-info) 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

<a href="#">1.</a>	Introduction . . . . .	<a href="#">2</a>
<a href="#">2.</a>	Terminology . . . . .	<a href="#">3</a>
<a href="#">3.</a>	Motivation . . . . .	<a href="#">4</a>
<a href="#">4.</a>	Architecture . . . . .	<a href="#">5</a>
<a href="#">4.1.</a>	Forwarding Cache for Wired subscribers . . . . .	<a href="#">5</a>
<a href="#">4.2.</a>	Bidirectional Cache for WLAN subscribers . . . . .	<a href="#">6</a>
<a href="#">4.3.</a>	Generalized cache architecture for intra-ISP networks . . . . .	<a href="#">7</a>
<a href="#">5.</a>	Considerations for ALTO deployment with P2P Caches . . . . .	<a href="#">8</a>
<a href="#">5.1.</a>	Forwarding Cache: vertical separator from outsiders . . . . .	<a href="#">9</a>
5.2.	Bidirectional Cache: horizontal division within insiders . . . . .	9
<a href="#">6.</a>	Open issues . . . . .	<a href="#">9</a>
<a href="#">6.1.</a>	Question: selective caching . . . . .	<a href="#">9</a>
<a href="#">6.2.</a>	Discussion: cooperative caching and ALTO extension . . . . .	<a href="#">10</a>
<a href="#">7.</a>	Security Considerations . . . . .	<a href="#">10</a>
<a href="#">8.</a>	IANA Considerations . . . . .	<a href="#">10</a>
<a href="#">9.</a>	References . . . . .	<a href="#">10</a>
<a href="#">9.1.</a>	Normative References . . . . .	<a href="#">10</a>
<a href="#">9.2.</a>	Informative References . . . . .	<a href="#">11</a>
	Authors' Addresses . . . . .	<a href="#">11</a>

[1.](#) Introduction

P2P applications like file sharing and multimedia streaming are so popular that lots of P2P technologies have been increasingly utilized throughout the world. The goal of Application-Layer Traffic Optimization (ALTO) [[I-D.ietf-alto-protocol](#)] is to provide guidance to these applications, which have to select one or several hosts from a set of candidates that are able to provide a desired resource.

Meanwhile, since wireless accesses to Internet have become pervasive with widely deployed WLANs, more and more people access Internet

services via WLAN and the amount of P2P traffic in WLAN is explosively growing. In addition to a huge number of individually setup WLANs at homes, there has been an increasing trend for the government, organizations, and even traditional network operators to set up publicly accessible WLAN facilities. Even though the service

Internet-Draft Considerations for ALTO with network-deployed February 2014

may be free of charge, to use the resources effectively in a fair way, and to avoid congestion for the purpose of service availability are vital for these public WLANs.

However, recent statistics reveal that P2P traffic accounts for 80% in part of China Mobile's WLANs, and traffic congestion at APs (access points) frequently occurs because of P2P applications. P2P traffic in WLANs not only causes problems on their own delivery quality, but also degrades the performance of other Internet applications in WLANs.

## [2.](#) Terminology

ALTO: application layer traffic optimization. For ALTO protocol, please refer to [[I-D.ietf-alto-protocol](#)].

AP: a wireless access point (or WAP) is a device that allows wireless devices to connect to a wired network using WLAN. The AP usually connects to a AC as a standalone device, but it can also be an integral component of the router itself.

AC: a wireless access controller (or WAC) is the network entity that provides wire access via APs to the network infrastructure in the data plane, control plane, management plane, or a combination therein.

DCF: Distributed coordination function (or DCF) is the fundamental MAC technique of the IEEE 802.11 based WLAN standard. DCF employs a CSMA/CA with binary exponential backoff algorithm.

Forwarding Cache: is a traditional content cache, which caches content flows from outside its coverage and serves subsequent requestors under its coverage for the content.

Reverse Cache: is a special content cache proposed for WLAN peers, which caches content flows from inside its coverage and serves

subsequent requestor from outside its coverage for the content on behalf of the peers within its coverage.

**Bidirectional Cache:** is a combination of a forwarding cache and a reverse cache.

**Transparent Cache:** is a content cache deployed by network operators for third party SPs (e.g. P2P streaming services), which participates in the overlay's service provision implicitly via request redirection.

**Cooperative Cache:** is a content cache deployed by network operators in cooperation with specific content delivering SPs (e.g. P2P streaming services), which participates in the overlay's service provision explicitly.

### 3. Motivation

On one hand, it is well accepted that compared to fixed networks, mobile networks have some special characters, including small link bandwidth, high cost, limited radio frequency resource, and terminal battery. Therefore, it is recommended by [[I-D.ietf-alto-deployments](#)] that in mobile network, the usage of wireless link should be decreased as far as possible and be high-efficient. For example, in the case of a P2P service, the clients in the fixed network should decrease the data transport from the clients in the mobile networks, as well as the clients in the mobile networks should prefer the data transmission from the clients in the fixed networks.

On the other hand, efforts have been put on using forwarding caches to optimize traffic pattern in such scenarios, which demonstrates great improvement in user experience and considerable traffic reduction at interworking points. What's more, owing to the characteristics of the DCF model in 802.11, there is an constant unfairness between uplink and downlink traffic in competing wireless channel resources, which leads to downlink congestion in WLAN resultant from P2P traffic (which constitutes the major part of uploading traffic). In particular, There is basic indication under the current DCF model, that in order to achieve fairness among mobile stations, in the long run, each stations has a equal chance of

getting the wireless slot given they all have a sustainable long traffic to send. In other words, there is an implicit unfairness between uplink and downlink traffic under the BSS model, where the AP holds the throat of all the downlink traffic and competing with the uplink traffic from potentially a large group of other user mobile devices. However, traditional P2P cache (as a forward cache) cannot help here, since it does nothing to stop a WLAN peer from uploading. Hence, bidirectional caches are proposed, which contains a reverse cache as well as a forward cache can be deployed at the AC (Access Controller). The reverse cache can provide uploading service instead of the WLAN peers under the AC's coverage. As a result, the uplink bandwidth consumption at each AP can be reduced and the uplink congestion can be alleviated effectively. Simulation results in file-sharing scenarios show that, employing cache instead of WLAN peers accelerates file transfer by 42% while improving the throughput of other Internet applications under the same AP by 28%.

With deployed P2P caches on the internal network, especially at a position as low as the AC-level, it could be sub-optimal to simply

Internet-Draft Considerations for ALTO with network-deployed February 2014

use the accessing network type as the divider for different PIDs and assign sufficient high cost within the wireless PID to prefer accessing remote peers over local peers blindly.

To this end, this draft discusses the optimal ALTO deployment recommendations for a P2P application in terms of a wireless accessing network with network-deployed application-agnostic caches.

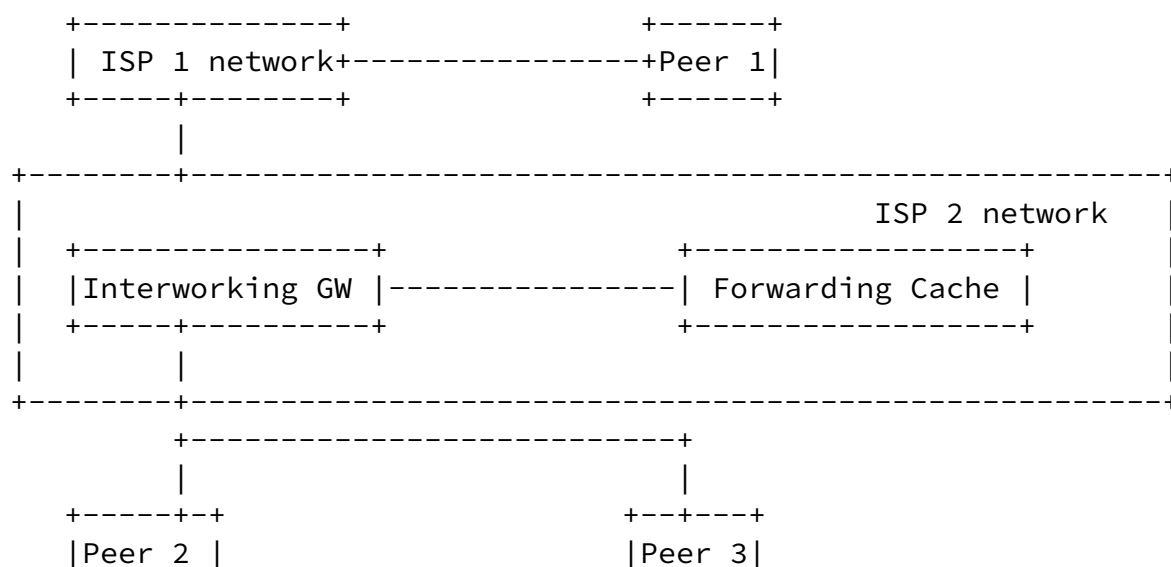
In summary, the goal here is to illuminate applications through ALTO about these existing network capabilities to make full use of them in achieving application performance improvement and network cost reduction.

## [4.](#) Architecture

### [4.1.](#) Forwarding Cache for Wired subscribers

Fig. 1 illustrates the proposed architecture of a traditional unidirectional P2P cache (or Forwarding Cache for short) system for wired subscribers, deployed mainly for the purpose of reducing interworking P2P traffic. Forwarding Caches are assumed to be deployed at the interworking gateways to maximize their coverage for

local subscribers. In transparent mode, they buffer downloading content from outside ISP networks, intercept the upcoming outgoing P2P requests from local subscribers and serve them with cached content instead. In cooperative mode, they register to the tracker as a super peer and are under the regulation of the tracker's private protocol.




+-----+

+-----+

Figure 1: Architecture of Forwarding Cache at interworking gateway

#### 4.2. Bidirectional Cache for WLAN subscribers

Fig. 2 illustrates the proposed architecture of a bidirectional cache system in WLAN. In a WLAN, all AP will connect to a device named AC, and the AC can be seen as the gateway to Internet of the WLAN. For most settings, both the traffic flowing into the WLAN and the traffic flowing out of the WLAN pass through the AC, hence Bidirectional Caches are assumed to be deployed at AC to exploit the traffic locality. Besides the normal functions of a Forwarding Cache, A Bidirectional Cache in transparent mode buffers uploading content from inside the WLAN network, intercepts the upcoming outgoing P2P responses from local WLAN subscribers and serve the correspondent requester (be it another local WLAN subscriber or an outsider) with cached content instead. In cooperative mode, they register to the tracker as a super peer and are under the regulation of the tracker's private protocol.



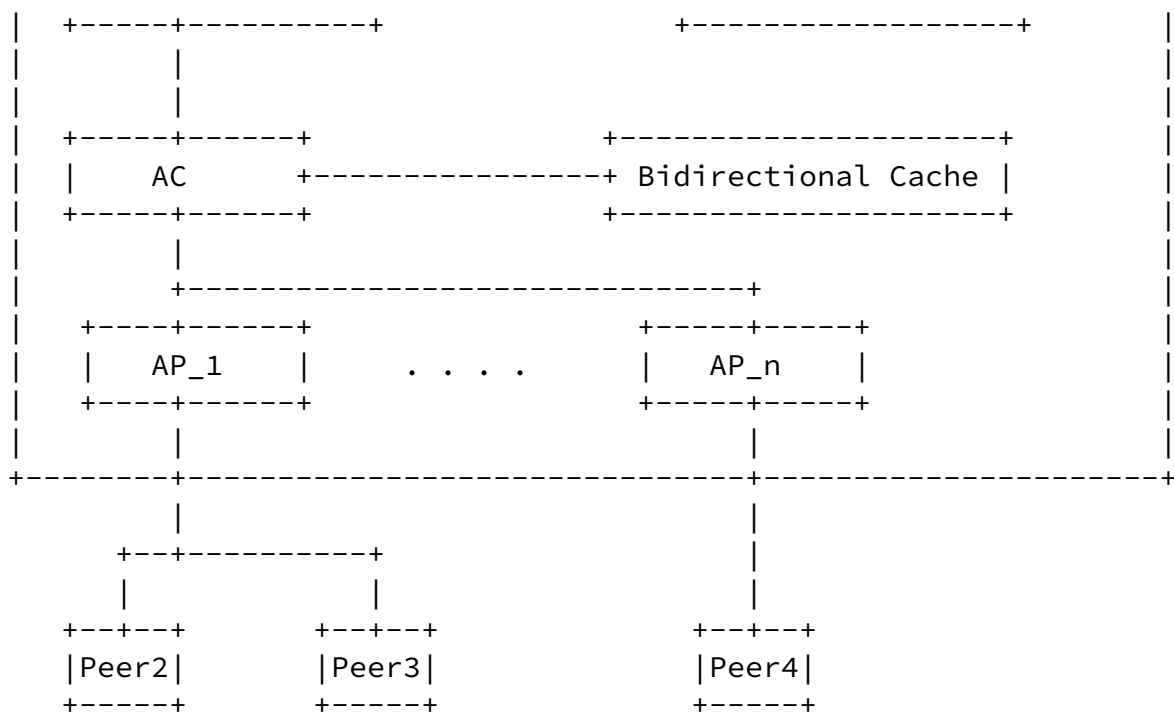


Figure 2: Architecture of Bidirectional Cache in WLAN

#### 4.3. Generalized cache architecture for intra-ISP networks

Fig. 3 generalized the overall architecture of the potential P2P cache deployments inside an ISP with various access network types. As it shows, P2P caches may be deployed at various levels, including: the interworking gateway linking with other ISPs, internal access network gateways linking with different types of accessing networks (e.g. WLAN, cellular and wired), and even within an accessing network at the entries of individual WLAN sub-networks. Moreover, depending on the network context and the operator's policy, each cache can be a Forwarding Cache or a Bidirectional Cache.



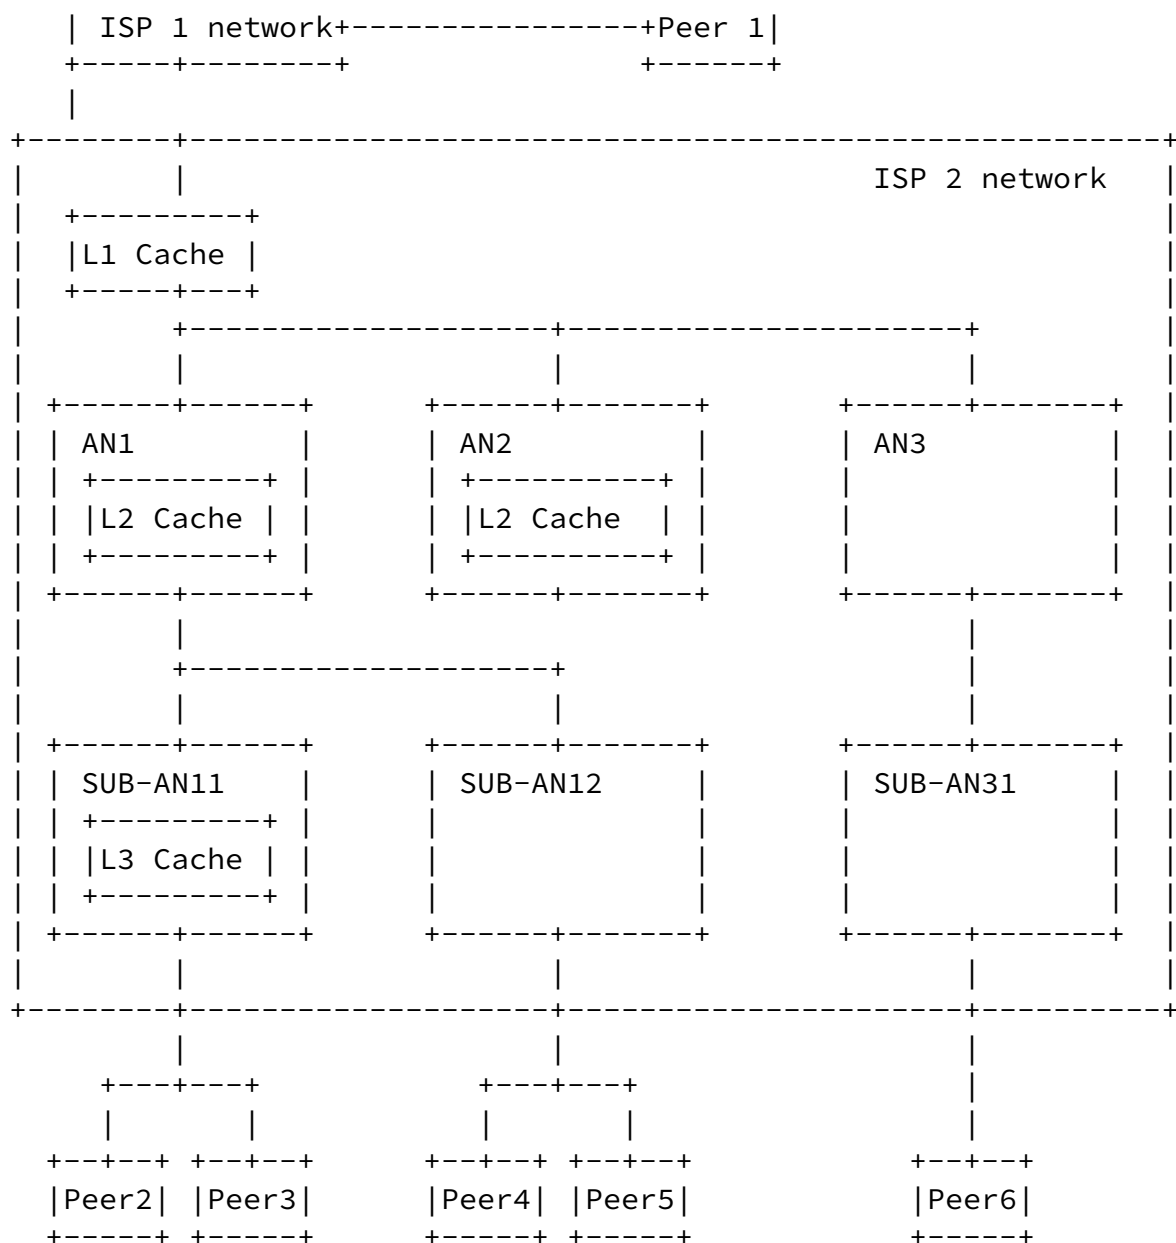


Figure 3: Generalized Architecture of intra-ISP Caches

## 5. Considerations for ALT0 deployment with P2P Caches

For wired network operator, forwarding caching effectively localizes the downloading P2P traffic within the sub-net under its coverage resulting in reduction of network cost for cross-boundary peer selection, whereas reverse caching blocks the uploading traffic outside a wireless sub-net leading to elimination of network cost for wireless uploading peer selection. In other words, caching between pairs of endpoints changes the traffic cost along the way.

Therefore, it is expected that by cooperation between the network operator and the P2P SP in building up various caching system and sharing information through ALTO protocol about these facilities brings benefits to both party.

In order to do that, it is proposed to use locations of caches as dividers of different PIDs to guide intra-ISP network abstraction and mark costs among them according to the location and type of relevant caches. Otherwise, if we stick to using network types as dividers for PIDs, there is no chance that a cost map would ever grasp this kind of information about node pairs within the same PID.

#### [5.1.](#) Forwarding Cache: vertical separator from outsiders

It is reasonable to use Forwarding Caches as separators for different PIDs, since it accelerates P2P traffic in a particular direction, indicating varied costs among these adjacent partitions. For instance as shown in Fig.3, assuming the L2 Cache in AN1 of ISP 2 network is a Forwarding Cache, the downloading traffic from other peers outside to AN1 but within the same ISP2 can be buffered once and served AN1 network subsequently. The cost from AN2 or AN3 to AN1 is reduced as result, but not vice versa. In other words, the ISP 2 network should be sub-divided into {AN1} and {AN2, AN3}, and the incoming P2P cost for {AN1} is reduced, for the sake of the L2 Forwarding Cache located at the entry of AN1.

#### [5.2.](#) Bidirectional Cache: horizontal division within insiders

Since Bidirectional Cache are deployed in wireless accessing networks to further reduce the outgoing and local-in-local-out traffic costs in both directions, it seems straightforward to join the adjacent partitions together and modify the cost between insiders to zero. However, there is hidden layering within the Bidirectional Cache coverage, as the blocking of uploading traffic only works for traffic traverse pass the Bidirectional Cache. If the cache is located too high as to be outside a local routing subnet, the local traffic flows within the subnet cannot benefit from the Bidirectional Cache.

### [6.](#) Open issues

#### [6.1.](#) Question: selective caching

As there is both CAPEX and OPEX expenditures for dedicated P2P Cache devices, it may be cost-efficient for caches to make buffering/serving decisions based on the popularity of the specific content. How to expose this application-relevant information to ALTO under

such context is an open issue.

Internet-Draft Considerations for ALTO with network-deployed February 2014

## [6.2.](#) Discussion: cooperative caching and ALTO extension

Luckily, in the cooperative-mode, a cache is playing as a normal peer under the tracker, and the latter can make the "right" decision in choosing in favor of the former under the guidance of the ALTO response while the tracker itself would take care of the content availability problem. If the cache doesn't have the content in question, it would no appear in the peer list handed in to ALTO server by the tracker.

In this case, the ALTO server can collect the information about caching sub-system in the network, identify those "caching" peers in the peer list of an cost request from an ALTO client, and arrange the returned rank list accordingly. For example, a simple candidate-ranking policy for a cost query to a WLAN peer, could be caching peers at the beginning, then inside wired peers, and lastly outside wired peers.

Moreover, the P2P SP and WLAN network operator may benefit even more by group popular files according to peers' geographic location or access types, and adapt its internal caching scheduling decisions about which files to be cached on which spot. In order to do that, it would be helpful that the ALTO server could provide to the client with the requesting peer's subscription types (i.e. wired/WLAN/cellular/...) as well as geographic locations. Specific definitions are specified separately in [I-D.[draft-deng-alto-p2p-ext](#)]

## [7.](#) Security Considerations

TBA

## [8.](#) IANA Considerations

None.

## [9.](#) References

### [9.1.](#) Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC2234] Crocker, D., Ed. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", [RFC 2234](#), November 1997.

Deng, et al.

Expires August 17, 2014

[Page 10]

---

Internet-Draft Considerations for ALTO with network-deployed February 2014

## [9.2](#). Informative References

- [I-D.[draft-deng-alto-p2p-ext](#)]  
Deng, L. and H. Song, "End Point Properties for Peer Selection", [draft-deng-alto-p2p-ext-00](#) (work in progress), February 2014.
- [I-D.[ietf-alto-deployments](#)]  
Stimerling, M., Kiesel, S., Previdi, S., and M. Scharf, "ALTO Deployment Considerations", [draft-ietf-alto-deployments-08](#) (work in progress), October 2013.
- [I-D.[ietf-alto-protocol](#)]  
Alimi, R., Penno, R., and Y. Yang, "ALTO Protocol", [draft-ietf-alto-protocol-25](#) (work in progress), January 2014.

## Authors' Addresses

Lingli Deng  
China Mobile

Email: [denglingli@chinamobile.com](mailto:denglingli@chinamobile.com)

Wei Chen  
China Mobile

Email: [chenwei@chinamobile.com](mailto:chenwei@chinamobile.com)

Qiuchao Yi  
China Mobile

Email: yiqiuchao@chinamobile.com

Yan Zhang  
Chinese Academy of Sciences

Email: zhangy@hpn1.ac.cn

Deng, et al.

Expires August 17, 2014

[Page 11]