

Application Layer Traffic	N. Weaver	
Optimization (ALTO) Working Group	International Computer Science	
Internet-Draft	Institute	
Intended status: Informational	March 04, 2009	
Expires: September 5, 2009		

[TOC](#)

## **Peer to Peer Localization Services and Edge Caches draft-weaver-alto-edge-caches-00**

### **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/ietf/lid-abstracts.txt>.

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

This Internet-Draft will expire on September 5, 2009.

### **Copyright Notice**

Copyright (c) 2009 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 in effect on the date of publication of this document (<http://trustee.ietf.org/license-info>). Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

### **Abstract**

Without caches in the infrastructure, peer to peer content delivery's primary effect is cost shifting rather than cost savings. Even with perfect localization, depending on the relative cost of last-mile

uplink bandwidth verses transport bandwidth, P2P may substantially increase aggregate cost. Yet the addition of edge caches, caches located in the ISPs near the customers, radically change the economics of P2P content delivery. Edge caches interact very strongly with localization services for P2P content delivery, and any localization service must be tightly integrated into edge-cache operation.

---

## Table of Contents

<a href="#">1.</a>	Introduction
<a href="#">2.</a>	The Design of Edge Caches
<a href="#">2.1.</a>	Safe Incentives for Edge Caches
<a href="#">3.</a>	An Economic Model for Delivery Costs
<a href="#">3.1.</a>	The Limits of Localization
<a href="#">4.</a>	Edge-Cache Interactions with Localization
<a href="#">5.</a>	Conclusions
<a href="#">6.</a>	Acknowledgements
<a href="#">7.</a>	IANA Considerations
<a href="#">8.</a>	Security Considerations
<a href="#">9.</a>	References
<a href="#">9.1.</a>	Normative References
<a href="#">9.2.</a>	Informative References
<a href="#">§</a>	Author's Address

---

## 1. Introduction

[TOC](#)

When compared with conventional content delivery, peer to peer content delivery of bulk data is significant at shifting costs from the content provider to the ISPs, but can often significantly magnify the aggregate cost of delivery. Depending on the particular costs to an ISP, even perfect localization (restriction of P2P activity to within the ISP's network) may still result in significantly higher aggregate costs over conventional content delivery, although localization does reduce transit costs.

However, if edge-caches are introduced into the architecture, the economics can change radically. Rather than increasing transport costs, P2P with ISP-provided edge caches reduce transport costs for all parties, achieving costs reductions for the ISP analogous to those seen with edge-based HTTP servers such as [Akamai \(Akamai Inc, "The Akamai CDN," 2008.\)](#) [akamai]. Yet unlike edge-based web servers, edge-caches for P2P are failure-transparent: when they fail, or do not have the right data, the failure does not impact correct operation of the P2P system.

It is critical that ALTO or other localization services for bulk-data P2P be both edge-cache aware and assist edge-caches in their operation, for localization without edge-caches may not produce significant cost savings to the ISPs or performance benefits to the customers, but edge-caches need localization services both to ease client discovery and to provide necessary topological information for edge-cache operation. This document begins with a brief discussion of [edge caches for P2P \(The Design of Edge Caches\)](#), then outlines a [simple cost model of content delivery \(An Economic Model for Delivery Costs\)](#), which argues why both localization and edge-caches are necessary for cost-effective content delivery. It then discusses [how localization and edge-caches should interact \(Edge-Cache Interactions with Localization\)](#), before a brief [conclusions section \(Conclusions\)](#). 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 \(Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels," March 1997.\)](#) [RFC2119].

---

## 2. The Design of Edge Caches

[TOC](#)

An edge-cache is simply a special P2P node which lives in the ISP's network close to, but not at, the final recipients. Thus it incurs no transit cost in communicating with ISP-local peers, and is close in latency and has a high-bandwidth connection into the ISP's internal network.

The role of an edge cache is to coordinate transfers between local peers and the rest of the Internet, as well as to cache data for subsequent use, within the existing or modified P2P protocol. For example, a BitTorrent edge cache can participate in a swarm, offering up data only to ISP-local peers once it has a complete file, and refusing to seed or leech (but only tit-for-tat) with peers outside the ISP before it has obtained the entire file.

One feature of an edge-cache is that it can be unreliable. Since, from the point of view of the other peers, it is simply another P2P participant, if the edge-cache fails to include a block, a file, or fails altogether, the P2P system will still work properly. This is in sharp contrast to edge-based HTTP caches or CDNs, where a failure in the node may result in failures to the user.

A side consequence of unreliability is that an edge-cache can therefore be inexpensive. For example, a 1U server (based on a Mini-ITX motherboard) capable of holding 4 SATA disks might cost less than \$800. With a price of \$130 for a 1.5TB drive, an edge cache costing less than \$1400 could cache over 5 TB of data. Such a low-cost system might suffer significantly higher transient failure rates than a higher-quality server, necessitating a reboot, reimage, and disabling of bad

disks, but as failures are low-consequence, such caches can be cheap to deploy.

Finally, a P2P edge-cache doesn't require changing existing P2P protocols. As long as local peers will find the edge cache, or the edge-cache can find the local peers, edge-caches can be introduced into existing protocols without change. In particular, BitTorrent is highly amenable to edge-caches without requiring client changes.

---

## 2.1. Safe Incentives for Edge Caches

[TOC](#)

The biggest impediment to building edge-caches is not technical but legal. Given a P2P swarm, a single edge cache or collection of caches should be able to monitor the swarm and find participants. But an edge cache needs to be notified both about a particular P2P swarm and that it is acceptable to cache the swarm.

It is outside the scope of this document for a detailed discussion, but there exist many possibilities, such as P2P content providers (such as Linux ISO images) registering their content, users of the ISP asserting that a swarm is legitimate (and consenting to be identified if a copyright holder objects), and agreements with third party data providers (such as Amazon S3) which support BitTorrent and other P2P content distribution.

---

## 3. An Economic Model for Delivery Costs

[TOC](#)

For purposes of this discussion, we assume that different portions of the network have different costs to transmit or receive one unit of data. Although costs really vary by time of day and network conditions (for example, the cost to an ISP of traffic on an uncongested uplink on the last mile is effectively 0, but can be huge if there is congestion, or peering arrangements may make the cost of uplink transit negative), for simplicity we will ignore these effects for now.

CP: This is the cost for the content provider to send one unit of data

CDN: This is the cost for the content provider to send one unit of data through a third party, edge-based CDN

CT: This is the cost for the ISP to receive one unit of data from the general internet

CTU: This is the cost for the ISP to send one unit of data to the general internet

CL: This is the cost for the ISP to send one unit of data to the end customer across the last mile

CLU: This is the cost for the ISP to receive one unit of data from an end customer across the last mile.

With such a basic cost model, it becomes possible to estimate the costs for different content delivery mechanisms.

Central (conventional) HTTP traffic: For such traffic, the content provider pays  $N \cdot CP$ , while the ISP pays  $N \cdot (CT + CL)$ . The costs increase linearly with the number of requests.

Edge-located HTTP content delivery networks (such as Akamai): For such traffic, the content provider pays  $N \cdot CDN$ , while the ISP pays  $N \cdot CL$ . This is obviously the best case for the ISP, but the cost of the CDN may not be favorable to the content provider.

Conventional P2P without localization: If we assume the P2P system is highly efficient, the content provider pays only  $CP$  regardless of the number of users. The ISP will need to pay  $N \cdot (CL + CLU)$  for all users on the last mile, and some value less than  $N \cdot (CT + CTU)$  for transit.

Conventional P2P with perfect localization: If the P2P system is perfect, including localizing the traffic completely within the ISP, the content provider pays only  $CP$ , while the ISP will need to pay  $N \cdot (CL + CLU)$  but only  $(CT + CTU)$  for transit.

Conventional P2P with perfect localization and perfect edge caches: Adding in edge-caches changes the situation. Now the content provider pays only  $CP$ , while the ISP pays  $N \cdot CL + CT + CTU$ .

---

### 3.1. The Limits of Localization

[TOC](#)

Such a simple cost model illustrates the major limitation of localization. If  $CLU$ , the cost of the last mile uplink, is more than  $CT$ , the cost of the transit downlink, P2P can significantly increase the costs to the ISP over conventional HTTP delivery, even with perfect localization and perfect operation. For some networks, such as DOCSIS cable modems, this is often the case, as increasing network capacity on the shared last mile may require new infrastructure or repurposing bandwidth otherwise used for higher-value services such as television channels.

Yet it shows that if edge-caches are added into the system, everybody sees a cost savings: both the content provider and the ISP benefit from lower cost, but without the reliability concerns present in edge-based HTTP CDNs. Thus edge-caches represent the best of both worlds: for a content provider, edge-caches in the P2P system have the same low cost as a conventional P2P system, but for the ISP, the edge-caches have the same low cost as an edge-located CDN.

---

## 4. Edge-Cache Interactions with Localization

[TOC](#)

Since edge-caches are critical to realize the true potential of P2P to create an aggregate cost savings, they need to be considered when

developing other portions of a common P2P infrastructure. In particular, edge-caches both interact with and benefit from localization services, and thus it is critical that both localization and edge-caching be codesigned to interoperate. Thus some edge-cache concerns which directly relate to localization.

Edge-cache discovery: Any localization service which supports the discovery of "preferable" nodes should give preference to any relevant edge-caches in the system. Thus the localization service will drive traffic towards the relevant edge caches, resulting in greater performance and lower cost-of-delivery.

Edge-cache content notification: Any localization service should also act as content notification, notifying the edge-cache about a user's desire to fetch a particular piece of content. The edge-cache may use this information, along with other constraints and heuristics, to determine whether it should participate in this distribution system. For example, a particular ISP's edge-cache for BitTorrent could be configured to cache torrents requested from Amazon S3 or other sources based on a contractual relationship, but reject torrents hosted elsewhere.

Peer-access control: The edge-cache, when contacted by a peer, needs to know whether the peer is local to its network. Thus the localization service should support queries from the edge cache as to whether a peer would be considered local to the ISP.

Support for file descriptors: In order for both the localization service and the edge-cache to track files as they are requested, ALTO requests from peers should include both a per-file unique ID and a variable length field containing the protocol's representation of the file requested (eg, for BitTorrent, the .torrent file). This has some minor privacy implications, but greatly enhances both the ability of localization to know which peers are involved in a particular transfer and the ability of edge-caches to determine which data to fetch.

---

## 5. Conclusions

[TOC](#)

Edge-caches are critical if P2P is to achieve the promised aggregate cost savings. Without an edge-cache, localization's benefits are limited, as even perfect localization is unable to reduce the transfers over the last-mile uplink. Yet edge-caches also need to rely on localization, both to drive traffic to the edge cache, to discover new content, and to determine which peers are allowed to access the edge-cache. Thus localization protocols should include edge-caches in their focus, and edge-caches will need to use localization protocols.

---

[TOC](#)

## 6. Acknowledgements

Grant info here. All opinions are those of the author, not the funding institution.

Feedback on the general concept and economic models for P2P edge caches from Richard Woundy, Jason Livingood, Vern Paxson, Christian Kreibich, and others.

---

## 7. IANA Considerations

[TOC](#)

None

---

## 8. Security Considerations

[TOC](#)

The privacy concerns of edge-caches and localization are only mild to moderate. It is already possible for P2P nodes to observe what other nodes are downloading or making available, and an edge-cache simply represents another such node in the system. Any P2P system which wishes to avoid this problem will not want to use localization (because of the impacts on traffic analysis), and ISPs will not want to cache such data (because most of the data will represent illegal content).

This is also why localization services such as ALTO should have a query interface that doesn't just give a list of IP addressees to rank, but also has query modes which present ALTO with a UUID and a content identifier, so a localization system can keep track of other systems which have already requested the same content.

---

## 9. References

[TOC](#)

### 9.1. Normative References

[TOC](#)

[RFC2119]	<a href="#">Bradner, S.</a> , " <a href="#">Key words for use in RFCs to Indicate Requirement Levels</a> ," BCP 14, RFC 2119, March 1997 ( <a href="#">TXT</a> , <a href="#">HTML</a> , <a href="#">XML</a> ).
-----------	--

---

## 9.2. Informative References

[TOC](#)

[akamai]	Akamai Inc, " <a href="#">The Akamai CDN</a> ," 2008.
----------	---

---

## Author's Address

[TOC](#)

	Nicholas Weaver
	International Computer Science Institute
	1947 Center Street suite 600
	Berkeley, CA 94704
	USA
Phone:	+1 510 666 2903
Email:	<a href="mailto:nweaver@icsi.berkeley.edu">nweaver@icsi.berkeley.edu</a>