

Peer-to-peer Research Group	E. Marocco	
Internet-Draft	A. Fusco	
Intended status: Informational	Telecom Italia	
Expires: January 12, 2010	I. Rimac	
	V. Gurbani	
	Bell Labs, Alcatel-Lucent	
	July 11, 2009	

[TOC](#)

Mythbusting Peer-to-peer Traffic Localization draft-marocco-p2prg-mythbusting-01

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/1id-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 January 12, 2010.

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

Peer-to-peer traffic optimization techniques that aim at improving locality in the peer selection process have attracted great interest in

the research community and have been subject of much discussion. Some of this discussion has produced controversial myths, some rooted in reality while others remain unfounded. This document evaluates the most prominent myths attributed to P2P optimization techniques by referencing the most relevant study (or studies) that have addressed facts pertaining to the myth. Using these studies, we hope to either confirm or refute each specific myth.

Table of Contents

- [1.](#) Introduction
- [2.](#) Definitions
 - [2.1.](#) Seeder
 - [2.2.](#) Leecher
 - [2.3.](#) Swarm
 - [2.4.](#) Tit-for-tat
 - [2.5.](#) Surplus Mode
 - [2.6.](#) Transit
 - [2.7.](#) Peering
- [3.](#) Myth: Reduced Cross-domain Traffic
 - [3.1.](#) Facts
 - [3.2.](#) Discussion
 - [3.3.](#) Conclusions
- [4.](#) Myth: Increased Application Performance
 - [4.1.](#) Facts
 - [4.2.](#) Discussion
 - [4.3.](#) Conclusions
- [5.](#) Myth: Increased Uplink Bandwidth Usage
 - [5.1.](#) Facts
 - [5.2.](#) Discussion
 - [5.3.](#) Conclusions
- [6.](#) Myth: Impacts on Peering Agreements
 - [6.1.](#) Facts
 - [6.2.](#) Discussion
 - [6.3.](#) Conclusions
- [7.](#) Myth: Impacts on Transit
 - [7.1.](#) Facts
 - [7.2.](#) Discussion
 - [7.3.](#) Conclusions
- [8.](#) Myth: Swarm Weakening
 - [8.1.](#) Facts
 - [8.2.](#) Discussion
 - [8.3.](#) Conclusions
- [9.](#) Myth: Improved P2P Caching
 - [9.1.](#) Facts
 - [9.2.](#) Discussion
 - [9.3.](#) Conclusions

10.	Security Considerations
11.	Acknowledgments
12.	Informative References
S	Authors' Addresses

1. Introduction

[TOC](#)

Peer-to-peer (P2P) applications used for file-sharing, streaming and realtime communications exchange large amounts of data in connections established among the peers themselves and are responsible for an important part of the Internet traffic. Since applications have generally no knowledge of the underlying network topology, the traffic they generate is frequent cause of congestions in inter-domain links and significantly contributes to the raising of transit costs paid by network operators and Internet Service Providers (ISP).

One approach to reduce congestions and transit costs caused by P2P applications consists of enhancing the peer selection process with the introduction of proximity information. This allows the peers to identify the topologically closer resource among all the instances of the resources they are looking for. Several solutions following such an approach have recently been proposed [[Choffnes](#)] ([Choffnes, D. and F. Bustamante, "Taming the Torrent: A practical approach to reducing cross-ISP traffic in P2P systems," .](#)) [[Aggarwal](#)] ([Aggarwal, V., Akonjang, O., and A. Feldmann, "Improving User and ISP Experience through ISP-aided P2P Localitytraffic in P2P systems," .](#)) [[Xie](#)] ([Xie, H., Yang, Y., Krishnamurthy, A., Liu, Y., and A. Silberschatz, "P4P: Provider Portal for Applications," .](#)), some of which are now being considered for standardization in the IETF [[ALTO](#)] ([, "Application-Layer Traffic Optimization \(ALTO\) Working Group," .](#)).

Despite extensive research based on simulations and field trials, it is hard to predict how proposed solutions would perform in a real-world systems made of millions of peers. For this reason, possible effects and side-effects of optimization techniques based on P2P traffic localization have been a matter of frequent debate. This document describes some of the most interesting effects, referencing relevant studies which have addressed them and trying to determine whether and in what measure they are likely to happen.

Each possible effect -- or Myth -- is examined in three phases:

- *Facts: in which a list of relevant data is presented, usually collected from simulations or field trials;

- *Discussion: in which the reasons for and against the myth are discussed based on the facts previously listed;

*Conclusions: in which the authors try to express a reasonable measure of the plausibility of the myth.

This document at the current stage is little more than a strawman. With the help of the IRTF community, the authors would like to improve it, in the number of the Facts, in the quality of the Discussion and, particularly, in the trustworthiness of the Conclusions.

2. Definitions

[TOC](#)

Terminology defined in [\[I-D.ietf-alto-problem-statement\] \(Seedorf, J. and E. Burger, "Application-Layer Traffic Optimization \(ALTO\) Problem Statement," July 2009.\)](#) is reused here; other definitions should be consistent with it.

2.1. Seeder

[TOC](#)

A peer that has a complete copy of the content it is sharing, and still offers it for upload. The term "seeder" is adopted from BitTorrent terminology and is used in this document to indicate upload-only peers also in other kinds of P2P applications.

2.2. Leecher

[TOC](#)

A peer that has not yet completed the download of a specific content (but usually has already started offering for upload the part it is in possession of). The term "leecher" is adopted from BitTorrent terminology and is used in this document to indicate peers that are both uploading and downloading, also in other kinds of P2P applications.

2.3. Swarm

[TOC](#)

The group of peers that are uploading and/or downloading pieces of the same content. The term "swarm" is commonly used in BitTorrent, to indicate all seeders and leechers exchanging chunks of a particular file; however, in this document it is used more generally, for example, in the case of P2P streaming applications, to refer to all peers receiving and/or transmitting the same media stream.

2.4. Tit-for-tat

[TOC](#)

A content exchange strategy where the amount of data sent by a leecher to another leecher is roughly equal to the amount of data received from it. P2P applications, most notably BitTorrent, adopt such an approach to maximize resources shared by the users.

2.5. Surplus Mode

[TOC](#)

The status of a swarm where the number of seeders is significantly greater than the number of leechers (i.e. the total download demand is abundantly satisfied by the upload capacity).

2.6. Transit

[TOC](#)

The service through which a network can exchange IP packets with all other networks it is not directly connected to. The transit service is always regulated by a contract, according to which the customer (i.e. a network operator or an ISP) pays the transit provider per amount of data exchanged.

2.7. Peering

[TOC](#)

The direct interconnection between two separate networks for the purpose of exchanging traffic without recurring to a transit provider. Peering is usually regulated by agreements taking in account the amount of traffic generated by each party in each direction.

3. Myth: Reduced Cross-domain Traffic

[TOC](#)

The reduction in cross-domain traffic (and thus in transit costs due to it) is one of the positive effects P2P traffic localization techniques are expected to cause, and also the main reason why ISPs look at them with interest. Simulations and field tests have shown a reduction varying from 20% to 80%.

3.1. Facts

[TOC](#)

1. Various simulations and initial field trials of the [P4P solution \(Xie, H., Yang, Y., Krishnamurthy, A., Liu, Y., and A. Silberschatz, "P4P: Provider Portal for Applications," .\)](#) [Xie] on average show a 70% reduction of cross-domain traffic.
2. Data observed in [Comcast's P4P trial \(Griffiths, C., Livingood, J., and R. Woundy, "Comcast's ISP Experiences In a Recent P4P Technical Trial," October 2008.\)](#) [I-D.livingood-woundy-p4p-experiences] show a 34% reduction of the outgoing P2P traffic and an 80% reduction of the incoming.
3. Simulations of the ["oracle-based" approach \(Aggarwal, V., Akonjang, O., and A. Feldmann, "Improving User and ISP Experience through ISP-aided P2P Localitytraffic in P2P systems," .\)](#) [Aggarwal] proposed by researchers at TU Berlin show an increase in local exchanges from 10% in the unbiased case to 60%-80% in the localized case.

3.2. Discussion

[TOC](#)

Tautologically, P2P traffic localization techniques tend to localize content exchanges, and thus reduce cross-domain traffic.

3.3. Conclusions

[TOC](#)

Confirmed.

4. Myth: Increased Application Performance

[TOC](#)

Ostensibly, the increase in application performance is the main reason for the consideration of P2P traffic localization techniques in academia and industry. The expected increase depends on the specific application: file sharing applications witness an increase in the download rate, realtime communication applications observe lower delay

and jitter, and streaming applications can benefit by a high constant bitrate.

4.1. Facts

[TOC](#)

1. Various simulations and initial field trials of the [P4P solution \(Xie, H., Yang, Y., Krishnamurthy, A., Liu, Y., and A. Silberschatz, "P4P: Provider Portal for Applications," .\)](#) [Xie] show an average reduction of download completion times between 10% and 23%.
2. Data observed in [Comcast's P4P trial \(Griffiths, C., Livingood, J., and R. Woundy, "Comcast's ISP Experiences In a Recent P4P Technical Trial," October 2008.\)](#) [I-D.livingood-woundy-p4p-experiences] show and increase in download rates between 13% and 85%. Interestingly, the collected data also indicate that fine-grained localization is less effective in improving download performance compared to lower levels of localization.
3. Data collected in the [Ono experiment \(Choffnes, D. and F. Bustamante, "Taming the Torrent: A practical approach to reducing cross-ISP traffic in P2P systems," .\)](#) [Choffnes] show a 31% average download rate improvement.

*In networks where the ISP provides higher bandwidth for in-network traffic (e.g. as in the case of RDSNET, described in [\[Choffnes\] \(Choffnes, D. and F. Bustamante, "Taming the Torrent: A practical approach to reducing cross-ISP traffic in P2P systems," .\)](#)), the increase is significantly higher.

*In networks with relatively low uplink bandwidth (as the case of Easynet, described in [\[Choffnes\] \(Choffnes, D. and F. Bustamante, "Taming the Torrent: A practical approach to reducing cross-ISP traffic in P2P systems," .\)](#)), traffic localization slightly degrades application performance.
4. Simulations of the ["oracle-based" approach \(Aggarwal, V., Akonjang, O., and A. Feldmann, "Improving User and ISP Experience through ISP-aided P2P Localitytraffic in P2P systems," .\)](#) [Aggarwal] proposed by researchers at TU Berlin show a reduction in download times between 16% and 34%.
5. [Simulations by Bell Labs \(Seetharaman, S., Hilt, V., Rimac, I., and M. Ammar, "Applicability and Limitations of Locality-Awareness in BitTorrent File-Sharing," .\)](#) [Seetharaman] indicate that localization is not as effective in all scenarios

and that the user experience can suffer in certain locality-aware swarms based on the actual implementation of locality.

4.2. Discussion

[TOC](#)

It seems that traffic localization techniques often cause an improvement in application performance. However, it must be noted that such beneficial effects heavily depend on the network infrastructures. In some cases, for example in networks with relatively low uplink bandwidth, localization seems to be useless if not harmful. Also, beneficial effects depend on the swarm size; imposing locality when only a small set of local peers are available may even decrease download performance for local peers.

4.3. Conclusions

[TOC](#)

Very likely, especially for large swarms and in networks with high capacity.

5. Myth: Increased Uplink Bandwidth Usage

[TOC](#)

The increase in uplink bandwidth usage would be a negative effect, especially in environments where the access network is based on technologies providing asymmetric upstream/downstream bandwidth (e.g. DSL or DOCSIS).

5.1. Facts

[TOC](#)

1. Data observed in [Comcast's P4P trial \(Griffiths, C., Livingood, J., and R. Woundy, "Comcast's ISP Experiences In a Recent P4P Technical Trial," October 2008.\)](#)
[I-D.livingood-woundy-p4p-experiences] show no increase in the uplink traffic.
-

[TOC](#)

5.2. Discussion

Mathematically, average uplink traffic remains the same as long as the swarm is not in surplus mode. However, in some particular cases where surplus capacity is available, localization may lead to local low-bandwidth leechers connecting to each other instead of trying the external seeders. Even if such a phenomenon has not been observed in simulations and field trials, it could occur to applications that use localization as the only means for optimization when some content becomes popular in different areas at different times (as is the case of prime time TV shows distributed on BitTorrent networks minutes after getting aired in North America).

5.3. Conclusions

[TOC](#)

Unlikely.

6. Myth: Impacts on Peering Agreements

[TOC](#)

Peering agreements are usually established on a reciprocity basis, assuming that the amount of data sent and received by each party is roughly the same (or, in case of asymmetric traffic volumes, a compensation fee is paid by the party which would otherwise obtain the most gain). P2P traffic localization techniques aim at reducing cross-domain traffic and thus might also impact peering agreements.

6.1. Facts

[TOC](#)

No significant publications, simulations or trials have tried to understand how traffic localization techniques can influence factors that rule how peering agreements are established and maintained.

6.2. Discussion

[TOC](#)

This is a key topic for network operators and ISPs, and certainly deserves to be analyzed more accurately. Some random thoughts follow.

It seems reasonable to expect different effects depending on the kinds of agreements. For example:

- *ISPs with different customer bases: the ISP whose customers generate more P2P traffic can achieve a greater reduction of cross-domain traffic and thus could probably be in a position to re-negotiate the contract ruling the peering agreement;

- *ISPs with similar customer bases:

- ISPs with different access technologies: customers of the ISP which provides higher bandwidth -- and, in particular, higher uplink bandwidth -- will have more incentives for keeping their P2P traffic local. Consequently, the ISP with a better infrastructure will be able to achieve a greater reduction in cross-domain traffic and will be probably in a position to re-negotiate the peering agreement;

- ISPs with similar access technologies: both ISPs would achieve roughly the same reduction in cross-domain traffic and thus the conditions under which the peering agreement had been established would not change much.

As a consequence of the reasoning above, it seems reasonable to expect that the simple fact that one ISP starts localizing its P2P traffic will be a strong incentive for the ISPs it peers with to do that as well.

It's worth noting that traffic manipulation techniques have been reportedly used by ISPs to obtain peering agreements [\[Norton\] \(Norton, W., "The art of Peering," .\)](#) with larger ISPs. One of the most used technique involves injecting forged traffic into the target ISP's network, in order to increase its transit costs. Such a techniques aims at increasing the relevance of the source ISP in the target's transit bill and thus motivate the latter to sign a peering agreement. However, traffic injection is exclusively unidirectional and easy to detect. On the other hand, if a localization-like service were used to direct P2P requests toward the target network, the resulting traffic would appear fully legitimate and, since in popular applications that follow the tit-for-tat approach peers tend to upload to the peers they download from, in many cases also bi-directional.

6.3. Conclusions

[TOC](#)

Likely.

7. Myth: Impacts on Transit

[TOC](#)

One of the main goals of P2P traffic localization techniques is to allow ISPs to keep local a part of the traffic generated by their customers and thus save on transit costs. However, similar techniques based on de-localization rather than localization may be used by those ISP that are also transit providers to artificially increase the amount of data exchanged with networks they provide transit to (i.e. pushing the peers run by their customers to establish connections with peers in the networks that pay them for transit).

7.1. Facts

[TOC](#)

No significant publications, simulations or trials have tried to study effects of traffic localization techniques on the dynamics of transit provision economics.

7.2. Discussion

[TOC](#)

It is actually very hard to predict how the economics of transit provision would be affected by the tricks some transit providers could play on their customers making use of P2P traffic localization -- or, in this particular case, de-localization -- techniques. This is also a key topic for ISPs, definitely worth an accurate investigation. Probably, the only lesson contentions concerning transit and peering agreement have taught so far [\[CogentVsAOL\]](#) ([Washington Post, "Peering Dispute With AOL Slows Cogent Customer Access,"](#) .) [\[SprintVsCogent\]](#) ([PC World, "Sprint-Cogent Dispute Puts Small Rip in Fabric of Internet,"](#) .) is that, at the end of the day, no economic factor, no matter how much relevant it is, is able to isolate different networks from each other.

7.3. Conclusions

[TOC](#)

Likely.

[TOC](#)

8. Myth: Swarm Weakening

Peer selection techniques based on locality information are certainly beneficial in areas where the density of peers is high enough, but may cause damages otherwise. Some studies have tried to understand to what extent locality can be pushed without damaging peers in isolated parts of the network.

8.1. Facts

[TOC](#)

1. [Simulations run by researchers at INRIA \(Le Blond, S., Legout, A., and W. Dabbous, "Pushing BitTorrent Locality to the Limit," .\)](#) [LeBlond] have shown that, in BitTorrent, even when peer selection is heavily based on locality, swarms do not get damaged.
 2. [Simulations by Bell Labs \(Seetharaman, S., Hilt, V., Rimal, I., and M. Ammar, "Applicability and Limitations of Locality-Awareness in BitTorrent File-Sharing," .\)](#) [Seetharaman] indicate that the user experience can suffer in certain locality-aware swarms based on the actual implementation of locality.
-

8.2. Discussion

[TOC](#)

It seems reasonable to expect that excessive traffic localization will cause some degree of deterioration in P2P swarms based on the tit-for-tat approach, and the damages of such deterioration will likely affect most users in networks with lower density of peers. However, as shown in [\[LeBlond\] \(Le Blond, S., Legout, A., and W. Dabbous, "Pushing BitTorrent Locality to the Limit," .\)](#), the right balance of randomness and locality depends on the P2P algorithm.

On the other hand, P2P systems not adopting the tit-for-tat approach (e.g. the eDonkey network) should not be damaged by locality-based optimizations.

8.3. Conclusions

[TOC](#)

Plausible, in some circumstances.

9. Myth: Improved P2P Caching

[TOC](#)

P2P caching has been proposed as a possible solution to reduce cross-domain as well as uplink P2P traffic. Since the cache servers ultimately act as seeders, a cache-aware localization service would allow a seamless integration of a caching infrastructure with P2P applications [\[I-D.weaver-alto-edge-caches\]](#) (Weaver, N., "Peer to Peer Localization Services and Edge Caches," March 2009.).

9.1. Facts

[TOC](#)

1. [A traffic analysis performed in a major Israeli ISP \(Leibowitz, N., Bergman, A., Ben-Shaul, R., and A. Shavit, "Are file swapping networks cacheable? Characterizing p2p traffic," .\)](#) [Leibowitz] has shown that P2P traffic has a theoretical caching potential of 67% byte-hit-rate.
-

9.2. Discussion

[TOC](#)

P2P Caching may be beneficial for both ISPs and network users, and locality-based optimizations may help the ISP to direct the peers towards caches.

9.3. Conclusions

[TOC](#)

Plausible.

10. Security Considerations

[TOC](#)

No considerations at this time.

[TOC](#)

11. Acknowledgments

This documents tries to summarize discussions happened in live meetings and on several mailing lists: all those who are reading this have probably contributed more ideas and more material than the authors themselves.

12. Informative References

[TOC](#)

[ALTO]	"Application-Layer Traffic Optimization (ALTO) Working Group."
[Aggarwal]	Aggarwal, V., Akonjang, O., and A. Feldmann, "Improving User and ISP Experience through ISP-aided P2P Localitytraffic in P2P systems."
[Choffnes]	Choffnes, D. and F. Bustamante, "Taming the Torrent: A practical approach to reducing cross-ISP traffic in P2P systems."
[CogentVsAOL]	Washington Post, "Peering Dispute With AOL Slows Cogent Customer Access."
[I-D.ietf-alto-problem-statement]	Seedorf, J. and E. Burger, "Application-Layer Traffic Optimization (ALTO) Problem Statement," draft-ietf-alto-problem-statement-02 (work in progress), July 2009 (TXT).
[I-D.livingood-woundy-p4p-experiences]	Griffiths, C., Livingood, J., and R. Woundy, "Comcast's ISP Experiences In a Recent P4P Technical Trial," draft-livingood-woundy-p4p-experiences-02 (work in progress), October 2008 (TXT).
[I-D.weaver-alto-edge-caches]	Weaver, N., "Peer to Peer Localization Services and Edge Caches," draft-weaver-alto-edge-caches-00 (work in progress), March 2009 (TXT).
[LeBlond]	Le Blond, S., Legout, A., and W. Dabbous, "Pushing BitTorrent Locality to the Limit."
[Leibowitz]	Leibowitz, N., Bergman, A., Ben-Shaul, R., and A. Shavit, "Are file swapping networks cacheable? Characterizing p2p traffic."
[Norton]	Norton, W., "The art of Peering."
[Seetharaman]	Seetharaman, S., Hilt, V., Rimac, I., and M. Ammar, "Applicability and Limitations of Locality-Awareness in BitTorrent File-Sharing."
[SprintVsCogent]	PC World, "Sprint-Cogent Dispute Puts Small Rip in Fabric of Internet."
[Xie]	

Xie, H., Yang, Y., Krishnamurthy, A., Liu, Y., and A. Silberschatz, "P4P: Provider Portal for Applications."

Authors' Addresses

[TOC](#)

	Enrico Marocco
	Telecom Italia
Email:	enrico.marocco@telecomitalia.it
	Antonio Fusco
	Telecom Italia
Email:	antonio2.fusco@telecomitalia.it
	Ivica Rimac
	Bell Labs, Alcatel-Lucent
Email:	rimac@bell-labs.com
	Vijay K. Gurbani
	Bell Labs, Alcatel-Lucent
Email:	vkg@alcatel-lucent.com