

Network Working Group
Internet-Draft
Intended status: Informational
Expires: September 3, 2012

M. Defeche
University of Liege
E. Vyncke, Ed.
Cisco Systems
March 2, 2012

Measuring IPv6 Traffic in BitTorrent Networks
draft-vyncke-ipv6-traffic-in-p2p-networks-01

Abstract

This document is a follow-up of a University thesis which aims to measure the evolution over time of IPv6 traffic and to analyze the geographical distribution of IPv6 nodes. The first measurements were done during the Summer 2009 using a specific-purpose program which connects to the BitTorrent peer-to-peer network and this document adds measurements done with the same program but in October 2011 and February 2012.

The study was made in Peer-to-Peer (P2P) networks because they are responsible for a big part of Internet traffic and because their structure and functioning permit a rapid discovery of a large number of nodes from all over the world. In addition, the P2P users are more likely to be interested by IPv6 as IPv6 does not have the same NAT problems as IPv4.

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 September 3, 2012.

Copyright Notice

Copyright (c) 2012 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.

Table of Contents

1.	Introduction	3
2.	Some explanations about BitTorrent	3
2.1.	Peer Wire Protocol	4
2.2.	Tracker	4
2.3.	Peer Exchange	5
2.4.	Distributed Hash Table	5
2.5.	Local Service Delivery	6
3.	Tools used for Measurement	6
4.	What was measured	6
4.1.	IPv6 Addresses	7
4.1.1.	Native IPv6	8
4.1.2.	Teredo	8
4.1.3.	6to4	9
4.1.4.	ISATAP	9
4.1.5.	Other Addresses	9
4.2.	Traffic Measurements	10
5.	IANA Considerations	10
6.	Security Considerations	10
7.	Acknowledgements	11
8.	References	11
8.1.	Normative References	11
8.2.	Informative References	11
	Authors' Addresses	12

1. Introduction

An IPv6 vs. IPv4 measurement was made in Peer-to-Peer (P2P) networks because they are responsible for most Internet traffic [[TFE3](#)] in 2009 and because their structure and functioning permit a rapid discovery of a large number of nodes from all over the world. In addition, the P2P users are more likely to be interested by IPv6 as IPv6 does not have the same NAT problems as IPv4.

The measurements were made in October 2011 and February 2012 re-using the application developed in October 2009 and presented in [[I-D.defeche-ipv6-traffic-in-p2p-networks](#)]. This was part of the Master Thesis [[THESIS](#)].

Measurements include: number of IPv4 vs. IPv6 nodes, which kind of IPv6 connectivity and geographical distribution. This measurement is intended to be run multiple times per year and the results are displayed on-line [[online](#)].

2. Some explanations about BitTorrent

Let's start with some explanations about BitTorrent.

BitTorrent is based on an hybrid decentralized network which is particularly well suited to the transfer of large files. BitTorrent generates the largest amount of traffic of all P2P networks and was installed on 28.20% of PCs in September 2007, and this number is certainly higher at present. BitTorrent also includes different protocols to discover peers, namely DHT, PEX and LSD which will be discussed later. Thanks to these mechanisms BitTorrent can be completely decentralized. The different clients are all compatible with the core protocol but some divergences concerning PEX, DHT and LSD appear between Azureus and the mainline, which represents at least the BitTorrent and uTorrent clients. Swarming is one of the basis of the protocol and IPv6 specification exists although it is not implemented by all clients.

Since BitTorrent is the only protocol that offers in theory a good support to IPv6, our choice was limited in 2009. But there are other reasons why BitTorrent is the network protocol that best matches our needs.

- o The number of different copies of the same file(s) is far smaller than in other networks. This leads to a larger number of peers sharing the same file(s). Therefore swarming can be more efficient thanks to a larger number of sources.

- o As BitTorrent is generally used to share large files, peers stay connected longer, giving us more time to discover them.
- o BitTorrent is responsible for the largest part of the P2P traffic throughout the world.
- o BitTorrent is widely used in most regions of the world.
- o Thanks to many different extensions like PEX ([Section 2.3](#)), DHT ([Section 2.4](#)) and LSD ([Section 2.5](#)), the discovery of peers is improved greatly.

2.1. Peer Wire Protocol

The Peer wire Protocol [[TFE5](#)], or PWP, specifies the way peers communicate in an asynchronous fashion with each other to exchange data and signalling messages. It is based on TCP connections that function in Full-Duplex and Pipelining mode to get better performances. This protocol does not define how to choose pieces to request, nor how to select peers to download from and to upload to. Certain algorithms, which are explained below, give some solutions to attain a good propagation of pieces in the swarm and to make peers happy with their download rate compared to their upload rate.

2.2. Tracker

The trackers [[TFE5](#)] act like servers but do not deal with the transfer of files ; their only purposes are to manage the swarm and to respond to periodic client requests for information about peers sharing the same torrent. Since the transfer of files is completely supported by peers, the bandwidth requirement is very low and thus a single tracker can handle many swarms, each one containing a large number of peers. This protocol commonly called THP is used by clients to communicate over HTTP with trackers. As a matter of fact, the trackers run an HTTP server. Peers contact trackers that are present in the metadata file for the following purposes :

- o to enter a swarm
- o to leave a swarm
- o to inform the tracker that the download is complete
- o to periodically give information on the download state and retrieve information about a random peer set. This time interval is defined by the tracker. If a peer misses a periodical request it can be considered as disconnected by the tracker and thus not present in the swarm any more

This communication permits trackers to keep track of peers that are in the swarm and to avoid referencing disconnected peers.

Tracker responses do not support IPv6 peers without this extension [[TFE12](#)] which means that they do not include IPv6 peers in their responses. This extension adds two new parameters for the tracker requests:

- o `ipv6` : the client IPv6 address. It can be added when the client contacts the tracker over IPv4;
- o `ipv4` : the client IPv4 address. Conversely, when communication is made over IPv6.

It permits clients having a dual stack to advertize both its addresses in the swarm. The port of the additional address is the same as the primary one.

2.3. Peer Exchange

The Peer Exchange or PEX is a means to discover new peers through peers that the client is already connected to. As a matter of fact, peers trade information concerning the peers they are connected to. Only few initial peers are needed to rapidly find a large number of peers. This mechanism permits a reduction of the tracker load and an improvement in the robustness as the tracker dependency is decreased.

2.4. Distributed Hash Table

The DHT technology [[TFE13](#)] is a way to store a hash table over a network, thus in a distributed way, each peers contains a part of it. Files and nodes are identified by a same length key, which is 20 bytes in BitTorrent. Each peer also maintains a list of different peers to efficiently route its searches. DHT in BitTorrent is an implementation of Kademlia which is based on the XOR metric that is the distance between two nodes or between a node and a file can be determined by a XOR of their keys.

This technology is used to decentralize the tracking mechanism to once more decrease the dependence on trackers. Even if the trackers are down or if no trackers are specified in the metadata file, the DHT technology permits the discovery of peers sharing the same torrent thanks to the info key hash as identifier. Conversely to PEX, no initial peers are needed.

2.5. Local Service Delivery

The Local Service Discovery or LSD permits the discovery of peers that are downloading the same torrent in the same local network. The transfer rate is much higher between two peers in the same local network than between two peers in different networks since the bandwidth limitation is that of the local network and not the one of the Internet connection which is far smaller, especially the upload stream. Briefly, LSD works as follows : the hash of the info key is broadcasted in the local network to find out peers sharing the same torrent.

3. Tools used for Measurement

As millions of pieces of information can be reached, the choice of a MySQL database came naturally for effectiveness reasons and because concurrent access is managed. Each program requests, inserts, and/or updates information in this database.

The computer that holds the database and executes the programs has native IPv6 and IPv4 connectivity, which will permit a better evaluation of the two versions than if we use Teredo, for instance.

Our specialized BitTorrent client joins different swarms and periodically changes to other swarms. While in a swarm we try to get connected to as many peers as possible thanks to all protocols supported by BitTorrent. In this way we are able to easily, automatically and efficiently discover peers. Ideally we should choose swarms with large numbers of peers in order to effectively retrieve information. Concerning the legality issue, we can use a trick to avoid downloading and uploading any files. In fact, we claim that we are not interested in any pieces so we will not download anything and we will not upload either since we have nothing to upload. So we are present in the swarm but without taking part in the sharing of files.

4. What was measured

The measurements were done in two periods:

- o May 2009 to July 2009: 5,000,000 peers were discovered but we were only able to establish a BitTorrent connection with 1,500,000 peers;
- o 1 day in October 2009: 100,000 peers were discovered;

- o 1 day in October 2011: 180,000 peers were discovered;
- o 2 days in February 2012: 664,685 peers were discovered.

4.1. IPv6 Addresses

We will now analyze the distinct IPv6 addresses found via the different protocols and classify them into different categories. The next section will explain the utilization of these addresses over the world in greater detail.

	Native	Teredo	6to4	ISATAP	Tunnel Brokers
Number in 2009	1,216	99,634	41,425	24	102
Percentage in 2009	0.85	69.72	28.99	0.02	0.08
Number in 2011	258	1,280	636	3	n/a
Percentage in 2011	11.85	58.80	29.22	0.14	n/a
Number in 2012	1,466	4,483	3,582	8	n/a
Percentage in 2012	15.37	47.00	37.55	0.08	n/a

Table 1: Different Types of IPv6 Addresses

The next table describes which weird and plain wrong IPv6 address our probe has received. It can be expected that a normal BitTorrent client will try to contact those addresses and therefore will generate packets whose destination IPv6 address is illegal.

	6bone	Site Local	IPv4 Compatible	IPv4 Mapped	Bogon
Number in 2009	436	24	1	94	74
Percentage in 2009	0.31	0.02	0.00	0.07	0.05
Number in 2011	0	0	0	0	61
Percentage in 2011	0	0	0	0	100
Number in February 2012	0	0	0	2	108

Table 2: Different Types of IPv6 Addresses (Cont.)

4.1.1. Native IPv6

In the 2011 study, only 197 distinct native IPv6 addresses were discovered in 25 different countries during our study. More than 40 % of these addresses came from the French ISP Free.

In the February 2012 study, many native IPv6 addresses have been discovered 1,466 (15% of all IPv6 addresses). France has still the highest ratio of native IPv6 addresses vs any addresses with 2.10% of the French BitTorrent peers using IPv6 (again Free being the largest followed by SFR). The next countries are:

- o Chine with 0.65% (only CERNET)
- o Japan with 0.59% (mainly KDDI)
- o United-States of America: 0.51% (mainly AT&T but also Comcast, Hurricane Electric and Verizon)

4.1.2. Teredo

Teredo with 47 % of IPv6 addresses found is clearly the most utilized way to get IPv6 connectivity. This ratio is also unusually high compared to other Internet measurements. It is probably linked to:

- o users: BitTorrent users are mainly residential and not professional, therefore Teredo is allowed to traverse the firewall while most enterprises block all outbound UDP traffic (and blocking Teredo in the same shot) and the most common Teredo implementation (Microsoft) also disables Teredo when used in an Active Directory network;
- o application: Teredo is only used by Windows to connect to an IPv6 which does not have any IPv4 address (in other word [[RFC3484](#)] severelly limits the normal use of Teredo), but, with BitTorrent the peers appear always as single stack (i.e. it appears like having only an IPv6 address even if it actually have both version of the IP protocol).

When we analyze the servers employed to configure these addresses we notice that only four servers represent 99.6 % of the used ones. Their addresses are as follows (February 2012 but mostly unchanged compared to 2011):

1. 65.55.158.118 : is deployed by Microsoft (58.80 %);
2. 94.245.115.184 : is deployed by Microsoft (1.3 %);

3. 94.245.121.251 : is deployed by Microsoft (17.20 %);
4. 94.245.121.253 : is deployed by Microsoft (22.30 %);

while in October 2009, three different servers represented 99 % of the traffic:

1. 213.199.162.214 : is deployed by Microsoft (46.12%);
2. 65.55.158.80 : is also set up by Microsoft (41.33%);
3. 207.46.48.150 : is once more owned by Microsoft (12.41%).

Thus 100 % of the peers that were using Teredo were likely to be using one of the Microsoft Operating Systems.

4.1.3. 6to4

6to4 is also still broadly used to provide IPv6 connectivity with 37.55% in February 2012, 29.21 % in 2011 and was 28.99 % in 2009 were created by this transition mechanism. Nevertheless, it is still far behind Teredo. It must be noted that the proportion of 6to4 nodes has not decreased as indicated by other measurements, but, the other measurements count only the web access to dual-stack servers where happy eye-balls [[I-D.ietf-v6ops-happy-eyeballs](#)] and [[RFC3484](#)] can influence the browser behavior in order to avoid all tunneling (this includes Teredo and 6to4).

4.1.4. ISATAP

We only discovered 8 different peers that used ISATAP ([[RFC5214](#)]) to get IPv6 connectivity in a non IPv6-capable infrastructure. This mechanism is less common than Teredo ([[RFC4380](#)]) and 6to4 ([[RFC3056](#)]). However, this bad result can be moderated by the fact that ISATAP is mostly destined to enterprises which often have a strict control on applications used by their employees. Thus installing a BitTorrent client is not often possible, and even if it is the firewall can filter its traffic.

4.1.5. Other Addresses

Although IPv6 addresses with the 6bone prefix should not exist anymore, in October 2009 we found up to 436 addresses with this prefix. In fact, all these addresses have the old Teredo prefix 3ffe:831f::/32 which was used in the 6bone network. The 6bone addresses are no more visible in 2011 and 2012.

In October 2009, we found 94 IPv4-mapped addresses that were all

discovered by PEX.

The good news is that in February 2012, we found not a single 6bone address and only 2 IPv4-mapped addresses in our peers.

Bogons are IP blocks that are reserved for private or special uses plus those that are not yet allocated by the IANA. Thus a bogon is an illegal IP address that must not appear on the Internet since they are theoretically unroutable. At present, the IPv6 unicast space is limited to the 2000::/3 prefix. During our study we discovered up to 1108 bogons via PEX.

4.2. Traffic Measurements

Most of the European countries have at least 1 % of their peers that can use IPv6 with better results for Latvia (3.88%), France (3.02%), Finland (2.90%), Romania (2.83%). The first non-European country is Chile with 2.78%. Another surprising fact is that China has only 1.19 % of its peers using IPv6 with 0.65% being native IPv6. This result may be negatively affected by the filter set up in this country.

As explained in the analysis of IPv6 addresses section, we only found few native IPv6 addresses. This analysis confirms what we discussed previously, which is that Japan and France have the highest rate but they are now joined by 'newcomers' such as China, USA but also Czech Republic and Slovenia.

5. IANA Considerations

There are no extra IANA consideration for this document.

6. Security Considerations

This I-D does not describe any specific protocol, so, the security section is mostly irrelevant. The measures were done with the specific security issues in mind:

- o copyrighted content: only a first fragment is downloaded and never stored on long term storage, so, it is assumed that even if copyrighted content is actually received it will never be user or propagated further to other peers;
- o denial of services: timers and rate limiters are used when getting the list of torrents from our directory, when contacting other peers, and so on

7. Acknowledgements

Many thanks to Professor Guy Leduc of the University of Liege and his research assistants, Cyril Soldani and Sylvain Martin. Martin is also grateful to Arvid Norberg who implemented the LibTorrent Rasterbar library [[TFE23](#)]. We also exchanged ideas with Nathan Ward.

8. References

8.1. Normative References

- [TFE12] Hazel, G., "IPv6 Tracker Extension", May 2008, <http://www.bittorrent.org/beps/bep_0007.html>.
- [TFE13] Maymounkov, P. and D. Mazieres, "Kademlia : A Peer-to-peer Information System Based on the XOR Metric", 2002, <<http://pdos.csail.mit.edu/~petar/papers/maymounkov-kademlia-lncs.pdf>>.
- [TFE5] Cohen, B., "The BitTorrent Protocol Specification", 2008, <http://www.bittorrent.org/beps/bep_0003.html>.

8.2. Informative References

- [I-D.defeche-ipv6-traffic-in-p2p-networks]
Defeche, M. and E. Vyncke, "Measuring IPv6 Traffic in BitTorrent Networks", [draft-defeche-ipv6-traffic-in-p2p-networks-00](#) (work in progress), October 2009.
- [I-D.ietf-v6ops-happy-eyeballs]
Wing, D. and A. Yourtchenko, "Happy Eyeballs: Success with Dual-Stack Hosts", [draft-ietf-v6ops-happy-eyeballs-07](#) (work in progress), December 2011.
- [RFC3056] Carpenter, B. and K. Moore, "Connection of IPv6 Domains via IPv4 Clouds", [RFC 3056](#), February 2001.
- [RFC3484] Draves, R., "Default Address Selection for Internet Protocol version 6 (IPv6)", [RFC 3484](#), February 2003.
- [RFC4380] Huitema, C., "Teredo: Tunneling IPv6 over UDP through Network Address Translations (NATs)", [RFC 4380](#), February 2006.
- [RFC5214] Templin, F., Gleeson, T., and D. Thaler, "Intra-Site Automatic Tunnel Addressing Protocol (ISATAP)", [RFC 5214](#),

March 2008.

- [TFE23] Norberg, A., "LibTorrent Rasterbar Library",
<<http://www.rasterbar.com/products/libtorrent/>>.
- [TFE3] Schulze, H., "Internet Study 2008/2009", 2009.
- [THESIS] Defeche, M., "Measure and Analysis of IPv6 Traffic in
Peer-to-Peer Networks. Master Thesis", August 2009.
- [geoip] Maxmind, "GeoLite Country", 2012,
<<http://www.maxmind.com/app/geolitecountry>>.
- [online] Vyncke, E., "IPv6-enabled BitTorrent Peers",
<<http://www.vyncke.org/ipv6status/p2p.php>>.

Authors' Addresses

Martin Defeche
University of Liege

Email: martin.defeche@gmail.com

Eric Vyncke (editor)
Cisco Systems
De Kleetlaan 6a
Diegem 1831
Belgium

Email: evyncke@cisco.com

