

IPv6 Operations
Internet-Draft
Intended status: Informational
Expires: April 19, 2010

M. Defeche
University of Liege
E. Vyncke
Cisco Systems
October 16, 2009

Measuring IPv6 Traffic in BitTorrent Networks
<[draft-defeche-ipv6-traffic-in-p2p-networks-00.txt](#)>

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of [BCP 78](#) and [BCP 79](#). This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

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 April 19, 2010.

Copyright Notice

Copyright (c) 2009 IETF Trust and the persons identified as the document authors. All rights reserved.

Internet-Draft

IPv6 Traffic in BitTorrent

October 2009

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

This document summarizes a University thesis which aims to measure the evolution over time of IPv6 traffic, to analyze the geographical distribution of IPv6 nodes and to compare the two Internet Protocol versions on many different criteria (RTT, latency, MTU). The measurements were done during the Summer 2009 using a specific-purpose program which connects to the BitTorrent peer-to-peer network.

The study was made in Peer-to-Peer (P2P) networks because they are responsible for most 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.

Internet-Draft

IPv6 Traffic in BitTorrent

October 2009

Table of Contents

1.	Introduction	4
2.	Some explanations about BitTorrent	4
2.1.	Peer Wire Protocol	5
2.2.	Tracker	5
2.3.	Peer Exchange	6
2.4.	Distributed Hash Table	6
2.5.	Local Service Delivery	6
3.	Tools used for Measurement	7
4.	What was measured	8
4.1.	IPv6 BitTorrent Clients	8
4.2.	IPv6 Addresses	8
4.2.1.	Native IPv6	9
4.2.2.	Teredo	9
4.2.3.	6to4	10
4.2.4.	Tunnel Brokers	10
4.2.5.	ISATAP	11
4.2.6.	Other Addresses	11
4.2.7.	Summary about IPv6 Addresses	12
4.3.	Traffic Measurements	12
4.4.	Comparaison IPv4 vs. IPv6	12
4.5.	Round-Trip Time	13
4.6.	Comparaison of Hops	14
4.7.	MTU Comparaison	14
4.8.	Monthly Evolution	15
4.9.	Protocols Used to Discover Peers	15
5.	IANA Considerations	16
6.	Security Considerations	16
7.	Acknowledgements	17
8.	References	17
8.1.	Normative References	17
8.2.	Informative References	17
	Authors' Addresses	18

1. Introduction

An IPv6 vs. IPv4 measurement was made in Peer-to-Peer (P2P) networks because they are responsible for most Internet traffic [\[TFE3\]](#) 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 May, June, July and October 2009 to get an idea of the evolution within a couple of months. The intend is to run the same measurements every month for a couple of years.

Measurements include: number of IPv4 vs. IPv6 nodes, which kind of IPv6 connectivity, comparing the MTU, RTT between the two versions of IP and geographical distribution.

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 [micro]Torrent 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 is limited. 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 and because this information is retrieved and handled by five different programs, 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 runs the operating system Debian GNU/Linux 5.0 and has native IPv6 and IPv4 connectivity, which will permit a better evaluation of the two versions than if we use Teredo, for instance.

All our programs evolved constantly to better match our needs and to improve their performance. Most of them were rethought several times to achieve our goals. A long time was spent to debug the LibTorrent Rasterbar library to get a full IPv6-capable BitTorrent client and to identify problems that The Pirate Bay had with IPv6.

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.

As we are interested in the number of routers on the path between us and the remote peer we must discover the value of the Time-To-Live or Hop Limit depending on the IP version. Since the TTL or Hop Limit is initially set depending on the operating system, for instance 64 for Linux, and because the number of routers on the path should not exceed 32 we can easily calculate the number of times the field was decremented.

Ping and ping6 commands are used to measure the latency to all

discovered peers. Tracepath and tracepath6 commands are used to

measure the TTL and the MTU between peers. Of course, when the IPv6 connectivity of a peer is through a transition mechanism such as 6to4 or Teredo, then it is easy to extract the IPv4 address of the peer from its IPv6 address and comparison can be done.

[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.

The number of peers to which we managed to establish a BitTorrent connection may seem low in comparison with the number of discovered peers but is not. In fact, certain peers are already disconnected when we obtain information about them. The study made by Tribler on PEX clearly indicated that information retrieved by PEX is mostly garbage. Furthermore, when a BitTorrent client has reached its maximum number of connections, it rejects connection attempts made by remote peers and stops initiating new connections. These are the two principal reasons that explain the difference between the number of discovered peers and peers we were connected to.

[4.1.](#) IPv6 BitTorrent Clients

Among the 621,625 peers whose client is known to us, 542,537 of them utilized one which supports IPv6, that is 87.28%. We could adjust the percentage of peers having IPv6 connectivity with this information. However, the IPv6 results will not be adjusted with this value since it has no influence on the different proportion analyzed. And among the 542,537 peers that used a BitTorrent client that supports IPv6, 53,828 were using IPv6, that is 9.92%.

[4.2.](#) IPv6 Addresses

We will now analyze the 142,904 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	1,216	99,634	41,425	24	102
Percentage	0.85	69.72	28.99	0.02	0.08

Table 1: Different Types of IPv6 Addresses

	6bone	Site Local	IPv4 Compatible	IPv4 Mapped	Bogon
Number	436	24	1	94	74
Percentage	0.31	0.02	0.00	0.07	0.05

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

[4.2.1.](#) Native IPv6

Only 1,216 distinct native IPv6 addresses were discovered in 28 different countries during our study. More than 73% of these addresses came from the French ISP Free. Even if we downloaded torrents whose names contained either FRENCH or VOSTFR we were not connected to more French peers than Italian peers, for instance. As the high number of native IPv6 addresses from France is not due to the high number of peers from France, this phenomenon can be explained principally because Free seems to be the only large ISP that proposes IPv6 connectivity. We also noticed that many Universities and institutes around the world offer IPv6 connectivity especially in Portugal where there are nine of them, which is more than any other country.

[4.2.2.](#) Teredo

Teredo with 69.72% 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);

- o application: Teredo is only used by Windows to connect to an IPv6 which does not have any IPv4 address (in other word [[RFC3484](#)])

severely limits the normal use of Teredo), but, with BitTorrent the peers appear always as single stack.

When we analyze the servers employed to configure these addresses we notice that only three servers represent more than 99% of the utilized ones. Their addresses are as follows :

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 more than 99% of the peers that were using Teredo were likely to be using one of the Microsoft Operating Systems. A test on my personal computer showed that Miredo, which is a type of Teredo IPv6 tunneling software for Linux, does not use one of the Microsoft Teredo servers.

[4.2.3.](#) 6to4

6to4 is also broadly used to provide IPv6 connectivity as 28.99% of the addresses found were created by this transition mechanism. Nevertheless, it is still far behind Teredo. As 6to4 is designed to provide IPv6 connectivity to many hosts with a single public IPv4 address, we checked if many 6to4 hosts were using the same 6to4 router. The result is not very satisfactory since 99.47% of the 6to4 hosts employ a distinct 6to4 router. However, we noticed that 9 hosts were using the same 6to4 router which is quite high. It seems that the 6to4 mechanism is mostly used for small LAN and not by large sites.

6to4 is more commonly used than native IPv6. We found 6to4 users in almost every country in Europe. The USA has far more 6to4 users than other countries.

[4.2.4.](#) Tunnel Brokers

We can determine, via the list of Tunnel Brokers of the SixXS service

[[[TFE28](#)]] and IPv6 Task Force [[[TFE29](#)]] websites, whether an ISP is a tunnel broker. Unfortunately, these lists are not exhaustive so we may have missed a peer using a Tunnel Broker. However, as we are aware of the largest Tunnel Broker, those we missed should be small and thus cannot really influence the result. No dynamic identification of the Tunnel Broker was implemented due to the difficulty of identifying them all. The Tunnel Brokers in Table 3 that we discovered represent only 0.08% of the IPv6 addresses found.

Tunnel Broker	Region	Number	Percentage
Hurricane Electric	Worldwide	25	24.51
Internode	Australia	2	1.96
Hexago	Canada	22	21.57
SixXs	Worldwide	49	48.04
XS4ALL	Netherlands	4	3.92

Table 3: List of Tunnel Brokers

Only 5 different Tunnel Brokers were found. The most widely used one is SixXS with 48.04% of the total Tunnel Broker utilization. We noticed that 61% of the peers that were using this Tunnel Broker came from the Czech Republic. Hexago is rather popular for a Tunnel Broker that operates only in Canada.

[4.2.5](#). ISATAP

We only discovered 24 different peers that used ISATAP ([[RFC4212](#)]) to get IPv6 connectivity in a non IPv6-capable infrastructure. Half of these addresses have a global unicast prefix and the other half have a 6to4 prefix. 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.

[4.2.6](#). Other Addresses

Although IPv6 addresses with the 6bone prefix should not exist anymore 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.

We found 94 IPv4-mapped addresses that were all discovered by PEX.

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 74 bogons via PEX, most of which had the b524::/16 prefix.

We found 24 site-local addresses via PEX. These addresses came from peers that were connected to other peers in their site that they

found via LSD. These addresses are obviously useless to us since we are not in the same site.

[4.2.7.](#) Summary about IPv6 Addresses

The most impressive results are those of Teredo and 6to4 that represent together 98.71% of IPv6 addresses discovered. The native IPv6 addresses form only a very small fraction of the total.

Another important conclusion is that PEX is not ideally implemented in many BitTorrent clients since they exchange data that cannot be used by the remote peers like bogons or addresses reserved for the private-use networks. Filtering the addresses to send via PEX avoids wasting processing time to initiate connections that cannot be established and limits the number of addresses to send which saves bandwidth.

[4.3.](#) Traffic Measurements

Most of the European countries have at least 3% of their peers using IPv6 with better results for Sweden, Portugal and Latvia. Another surprising fact is that China has only 2.64% of its peers using IPv6 which seems strange since China is one of the leading countries in IPv6 deployment. This result may be negatively affected by the filter set up in this country. Even if they have many free IPv4

block left, the USA still has more than 4.8% of their peers using IPv6 which is more than their current policy supposes.

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 with 1.09% and 0.65% respectively. China with 0.65%, has a good percentage in comparison with other countries. This is explained by its interest to widely deploy native IPv6. Thus China has particularly poor results concerning 6to4 and Teredo traffic.

[4.4.](#) Comparaison IPv4 vs. IPv6

In this section we will compare the two IP versions with different IT statistics to find out whether the new version is really better. Of course, we must take into account the fact that IPv6 transition mechanisms have a structure that can considerably influence certain statistics. We will thus treat the IPv6 peers differently according to the mechanism they use, if they use one at all.

[4.5.](#) Round-Trip Time

First of all, even if we tried to obtain RTT statistics about each peer we did not get it all because some were disconnected when we did the test and others were not reachable via ICMP as many routers do not forward this type of ICMP message. We obtained this information for approximately 36% of the IPv4 peers and 18% of the IPv6 ones.

As we tested it for IPv4 addresses embedded in IPv6 ones we will compare the two results. We have the RTT information of 32% of the embedded IPv4 addresses, most of which come from 6to4. As a matter of fact, many NAT devices block the ICMP traffic which explains why we have so little RTT information about IPv4 addresses embedded in Teredo ones.

Table 4 indicates at first sight that IPv4 is two times faster than IPv6. This global result is not representative of the new IP version speed as it is greatly influenced by the transition mechanisms that

are normally much slower than native connectivity. The standard deviation average concerns the average of the standard deviation of all peers made thanks to the three different attempts.

Global	IPv4	IPv6
Average	324.8	812.8
Standard Deviation	403.6	890.8

Table 4: Global RTT Information in msec

Table 5 shows the results of the native IPv6 peers in comparison with those of the IPv6 (all kind of connectivity) and IPv4 peers. We can immediately see that the IPv6 RTT average is clearly smaller than the global IPv6 RTT average. However, the IPv6 RTT average is still higher than that of IPv4 but this difference can be explained by the current structure of the Internet: where even if the peer has native IPv6 connectivity, it does not mean that the IPv6 peering agreements of his/her ISP are as good as their IPv4 ones. This inconvenience should be reduced by the IPv6 deployment evolution. The average of the standard deviation of IPv6 peers is particularly high which implies that different attempts gave highly different results.

	Native IPv6 connectivity	All IPv6 connectivity	IPv4
Average	493.1	812.8	324.8
Standard Deviation	646.3	890.8	403.6

Table 5: Native IPv6 RTT Information in msec

[4.6.](#) Comparaison of Hops

We will analyze the TTL and Hop Limit on the path from the remote peers to us for which we obtained more than 1,300,000 statistics about IPv4 peers and more than 90,000 about IPv6 peers. We observe in Table 6 that the number of routers on the path from IPv4 peers is superior to the number of routers from IPv6 peers but not to that of native IPv6 which is quite coherent. As a matter of fact, as the Hop Limit is not decremented in the tunnel, its value arrives higher at the destination. The fact that there is a higher number of routers between native IPv6 peers and us also seems logical as the number of interconnections in the IPv6 Internet is smaller than in the IPv4 Internet and that may result in a longer path.

Hop count	IPv4	IPv6	Native IPv6
Average	15.10	11.93	18.12

Table 6: Native IPv6 RTT Information in msec

[4.7.](#) MTU Comparaison

We obtained more than 190,000 MTU statistics about IPv6 peers and more than 1,800,000 about IPv4 peers, which gives us a sufficient sample.

Table 7 shows the different MTU results depending on the IP version and whether the path MTU discovery reached the destination. Unsurprisingly, the global IPv6 MTU is far smaller than the IPv4 one as more than 98% of the addresses were formed thanks to a tunneling mechanism. The results of the native IPv6 peers are much higher but are still inferior to those of IPv4.

MTU	IPv4	IPv6	Native IPv6
Average	1,497.6	1,288.9	1,468.0

+-----+-----+-----+-----+

Table 7: Native IPv6 RTT Information in msec

4.8. Monthly Evolution

The monthly analysis done in 2009, shown in Table 8, allows us to see this decrease more clearly with 1.7% lost between the month of June and July. As July is a holiday month, the sharing habits may be different. As a matter of fact, most downloaders are students. However, this hypothesis does not explain why the IPv6 traffic decreases. The fact that Universities, which may have IPv6 connectivity, may be closed during the holidays and thus prevent students from using their bandwidth cannot explain this reduction as the percentage of IPv6 traffic coming from Universities is very low. Nevertheless, this decline from the Universities exists as IPv6 traffic in July diminished by more than three times when compared to results in June.

Once more the peers that were solely discovered via DHT are not taken into account in the IPv6 evolution analysis.

	May 2009	June 2009	July 2009	October 2009
%-age of IPv6 peers	4.50	4.79	3.10	2.30

Table 8: Monthly Evolution of IPv6 vs IPv4 peers

4.9. Protocols Used to Discover Peers

In this section we will look at the effectiveness of the different protocols to discover peers. As we can see in Table 9, PEX is globally far more effective to find peers than other means. It is even more efficient for IPv6 in comparison with the performance of trackers.

Source	IPv4	IPv6	Total
Tracker	1,550,023	37,745	1,587,768
PEX	3,558,648	167,606	3,726,254
DHT	843,353	0	843,353
LSD	0	0	0

Table 9: Number of Valid Peers Discovered by Protocol

Concerning DHT, we have never been able to get IPv6 peers by this means because of incompatibilities between BitTorrent clients and because of problems in the LibTorrent Rasterbar library. But if we look at the Table 9, the DHT results are not as good as the other means.

The good performance of PEX can be explained by its exponential way of functioning. As a matter of fact, via a small number of peers a large number of peers can be reached. However, as already explained, the proportion of disposable information is very important.

We did not receive any peers via LSD either since no other computer was running a BitTorrent client in the local network. Nevertheless, it was functioning during all our study.

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., "Kademlia : A Peer-to-peer Information System Based on the XOR Metric".
- [TFE28] SixXs, "List of Tunnel Brokers", <<http://www.sixxs.net/tools/aiccu/brokers/>>.
- [TFE29] IPv6 Task Force, "List of Tunnel Brokers", <<http://www.ipv6tf.org/index.php?page=using/connectivity/test>>.
- [TFE5] Cohen, B., "The BitTorrent Protocol Specification", 2008, <http://www.bittorrent.org/beps/bep_0003.html>.

[8.2.](#) Informative References

- [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.
- [RFC4212] Blinov, M. and C. Adams, "Alternative Certificate Formats for the Public-Key Infrastructure Using X.509 (PKIX) Certificate Management Protocols", [RFC 4212](#), October 2005.

- [RFC4380] Huitema, C., "Teredo: Tunneling IPv6 over UDP through Network Address Translations (NATs)", [RFC 4380](#), February 2006.
- [TFE23] Norberg, Arvid., "LibTorrent Rasterbar Library", <<http://www.rasterbar.com/products/libtorrent/>>.
- [TFE3] Schulze, H., "Internet Study 2008/2009", 2009.

Defecche & Vyncke

Expires April 19, 2010

[Page 17]

Internet-Draft

IPv6 Traffic in BitTorrent

October 2009

- [THESIS] Defecche, M., "Measure and Analysis of IPv6 Traffic in Peer-to-Peer Networks. Master Thesis", August 2009.

Authors' Addresses

Martin Defecche
University of Liege
Institut Montefiore
Bat. B28 Reseaux informatiques
Grande Traverse 10
Liege 4000
Belgium

Email: martin.defecche@gmail.com

Eric Vyncke
Cisco Systems
De Kleetlaan 6a
Diegem 1831
Belgium

Phone: +32 2 778 4677
Email: evyncke@cisco.com

