

PPSP
Internet-Draft
Intended status: Standards Track
Expires: April 23, 2014

L. Deng
J. Peng
China Mobile
Y. Zhang
CoolPad
Y. Huang
Huawei
October 20, 2013

Efficient Chunk Availability Compression for PPSP
draft-deng-ppsp-bfbitmap-03.txt

Abstract

Bloom filters are proposed to be used in compressing chunk availability information, periodically exchanged between peers and the tracker through PPSP-TP and PPSP protocols, to reduce relevant cost and enhance scalability.

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 April 23, 2014.

Copyright and License Notice

Copyright (c) 2013 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. Code Components extracted from this document must

include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	2
2. Background on Bloom Filter	3
3. BF-based Chunk Availability Exchange	5
3.1. A non-BF PPSP session	5
3.2. A PPSP Session with BF-bitmaps	6
3.3. Summary	8
3.3.1. Base TP protocol	8
3.3.2. Extended TP protocol	8
3.4. Peer protocol	9
4. Open Issues	10
5. Security Considerations	11
6. IANA Considerations	11
7. References	11
7.1. Normative References	11
7.2. Informative References	11
Authors' Addresses	12

1. Introduction

As it is pointed out by [I-D.ietf-ppsp-problem-statement], current P2P streaming practices often use a "bitmap" message to exchange chunk availability. The message is of kilobytes in size and exchanged frequently, e.g., an interval of several seconds or less.

To begin with, in a mobile environment with scarce bandwidth, the message size may need to be shortened or it may require more efficient methods for expressing and distributing chunk availability information, which is different from wire-line P2P streaming.

Even in a wire-line P2P streaming application, frequent exchange of large volume of bitmap information, is among the key factors that set a limit to the system's efficiency and scalability [P2P-limit].

Therefore, the following requirements for PPSP protocols in terms of chunk availability exchange are stated in [I-D.ietf-ppsp-problem-statement] :

PPSP.TP.REQ-3: The tracker protocol MUST take the frequency of messages and efficient use of bandwidth into consideration, when communicating chunk availability information.

PPSP.PP.REQ-7: The peer protocol MUST take the frequency of messages and efficient use of bandwidth into consideration, when communicating chunk information.

In this draft, we propose an efficient bitmap compression scheme for chunk availability information in PPSP protocols. Given the Bloom Filters' wide applications in Internet and demonstrated efficiency with highly compacted data structure and low complexity and cost in terms of information storage, transportation and computation, it is expected to relieve a PPSP implementer from the dilemma between "the frequency of messages" (i.e. the timely exchange of information that contributes to better user experience) and "efficient use of bandwidth" (i.e. the limit of a single node/peer that holds the system's overall scalability by throat).

2. Background on Bloom Filter

Bloom Filter (or BF for short) was first introduced in 1970s [BF-bloom], which makes use of multiple hashing functions to build a mapping from a set of elements to a compact binary array, to realize highly efficient member queries with a tolerably low error rate of wrongly reported hits. Despite their extraordinary efficiency in terms of storage reduction and query acceleration, BFs suffer from the fact that there is possibility that a non-member of the set be wrongly taken as a member after the query. However, research [BF-analysis] shows that the odds that a BF-based membership query makes an erroneous hit can be suppressed to near zero, by a tactful configuration of various system parameters, including the hash functions used, the number of hash functions to be used, and the length of the bit array.

```
-----
BF(set S, integer m, hash set H)
1 filter=allocate m bits initialized to 0;
2 for each element xi in S do
3   for each hash functions hi in H do
4     filter[hi(xi)]=1;
5 return filter;
-----

MT(element elm, BF filter, integer m, hash set H)
1 for each hash functions hi in H do
2   if (filter[hi(elm)]!=1)
3     return false;
4 return true;
-----

ST(BF query, BF filter)
```

```

1 temp=query OR filter;
2 if (temp!=filter)
3     return false;
4 return true;
-----

```

Figure 1: Basic algorithms for BF-bitmap

As shown in Figure 1, the $BF(S,m)$ algorithm takes a n -membered subset $S=\{x_1, x_2, \dots, x_n\}$ from a universal set U as input, and outputs a m -bit binary array B as a compacted representation of S . In order to do that, it makes use of k independent random hash functions, each of which maps a member to a marked bit in B (i.e $h_j: U \rightarrow [1, m]$, $j=1 \dots k$). The BF algorithm is highly efficient in the following aspects:

- o It is quite simple and straightforward to generate the BF representation of a set S , $B=BF(S)$: initially, all the bits in B is set to 0; then, for each member x of the set S , mark each bit in B , to which a hash function maps x (as shown in Figure 1 as the BF algorithm).
- o It is highly efficient to check whether or not a given element x is in any BF-represented set $B=BF(S)$: for each hash function h_j , check the value of $B[h_j(x)]$ against 1. It is always safe to exclude the element x out of set S , once there is a zero-valued hash bit. Otherwise it is assumed that x is a member of S (the MT algorithm in Figure 1).
- o It is also highly efficient to check whether or not a given element set S is contained by another set S' if they are both represented as BF-bitmaps, say $query(=BF(S))$ and $filter(=BF(S'))$ for instance. It is always safe to return "false" to the requestor if there's any marked bits in $BF(S')$ and not marked in $BF(S)$, which can be realized by two simple bitwise operations(as shown by the ST algorithm in Figure 1).

For instance, given a 2GB movie file, the original bitmap for a sharing peer would be 1024-bit (if the system is using 2MB-sized segments). By simply using 4 uniform random hash functions and a 128-bit BF-bitmap, the possibility of erroneous hits by MT algorithm would be lower than 3%.

As for a simple illustration, the 4 hash functions may be established through the MD5 message-digest algorithm [RFC1321], a widely used cryptographic hash function that produces a 128-bit (16-byte) hash value from an arbitrary binary input. MD5 has been utilized in a wide variety of security applications, and is also commonly used to check data integrity.

Specifically, the processing of 4 hash functions is as follows: use the MD5 algorithm to turn a given chunk_ID into a 128-bit binary array, further separate the 128 bits into 4 arrays (32-bit each), and finally divide each of them using 128 to yield 4 integers in the range of [1,m].

3. BF-based Chunk Availability Exchange

We first construct a general message flow (shown in Figure 2) from PPSP protocols, and then discuss how to integrate BF-bitmap algorithm with it.

3.1. A non-BF PPSP session

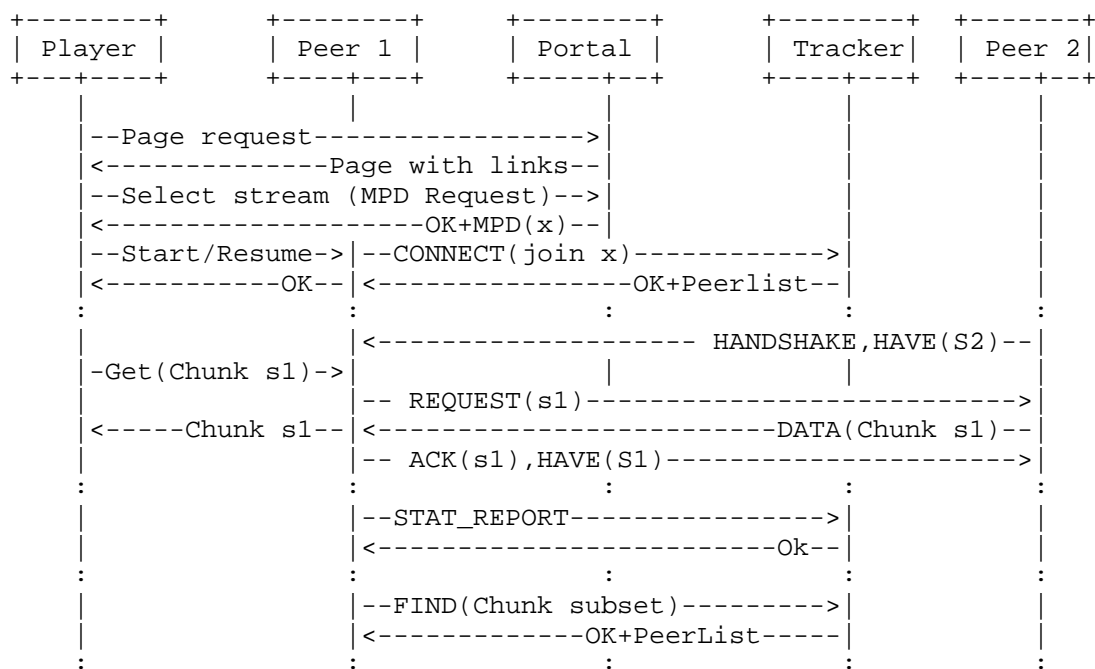


Figure 2: A typical PPSP session for watching a streaming content.

When a peer wants to receive streaming of a selected content (Leech mode):

1. Peer connects to a connection tracker (which may be located through a web portal) and joins a swarm.

2. Peer acquires a list of other peers in the swarm from the connection tracker (via the tracker protocol) through the CONNECT message.
3. Peer exchanges its content availability with the peers on the obtained peer list (via peer protocol) through the HAVE message.
4. Peer requests content from the identified peers (via peer protocol) through the REQUEST-DATA messages.
5. Peer periodically reports its status and chunk availability with the tracker (via the tracker protocol) through the STAT_REPORT message.
6. Peer acquires a list of other peers for a specific subset of media chunks in the swarm from the connection tracer (via the tracker protocol) through the FIND message.

3.2. A PPSP Session with BF-bitmaps

This document proposes to employ bloom filter algorithms in compressing chunk availability information exchanged and stored between peers and the tracker through the PPSP-TP protocols and PPSP protocol. Relevant extensions to the current protocols are summarized as follows: (as shown in Figure 3)

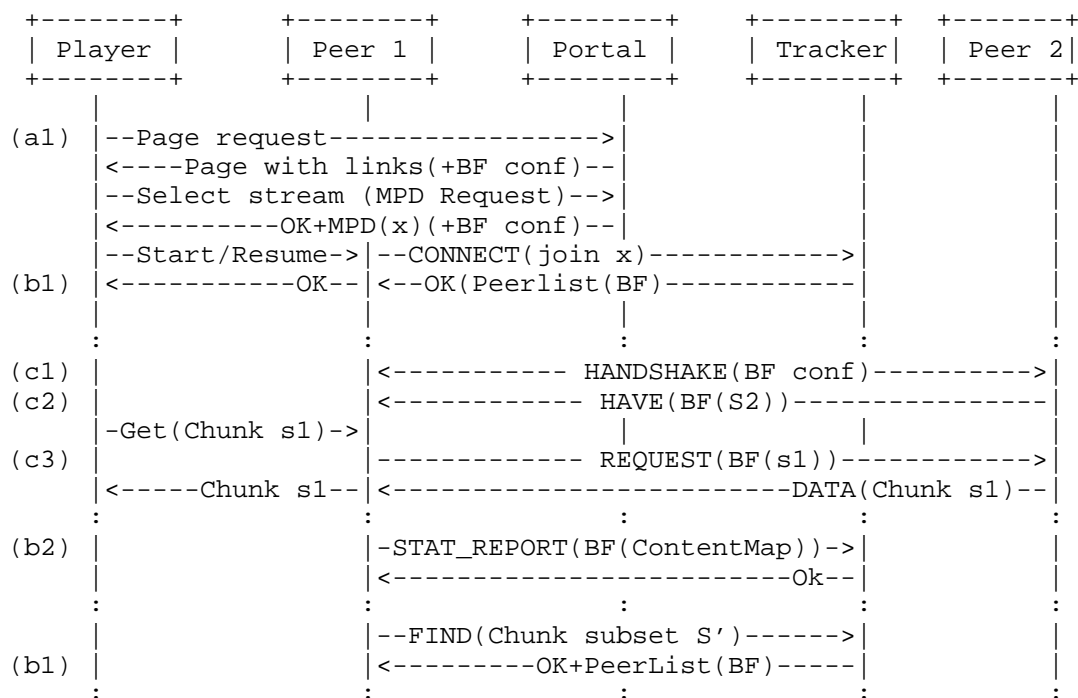


Figure 3: A typical PPSP session with BF-bitmaps.

a. Configuration Setup:

- * (a1) Configuration Setup: m , The length of the output bit array and H , the hash functions in use, are system level parameters that should be configured globally: (a1) the BF configurations (or BF conf for short) be stored at the web portal and published to a requesting peer through the web page or MPD file transaction, which could be incorporated to the "Installation and Initial Setup" in PPSP TP protocol [I-D.ietf-ppsp-base-tracker-protocol].

b. Integration to the extended TP protocol

[I-D.ietf-huang-extended-tracker-protocol]: In the extended TP protocol, a new "ContentGroup" element is added to the format of Request. It MAY be present in requests referencing content, i.e., CONNECT and FIND, if the request includes a content scope. The relevant interactions are described as follows:

- * (b1) In response to a CONNECT(join)/FIND request from a peer, the tracker may accompany the returned peer list with each recommended peer's BF-formed chunk availability bitmap, as the initial guidance for the requestor to start looking for neighbors in the same swarm. The additional cost for bearing the chunk-level availability information is constant ($O(m)$) for each peer in the returned peer list.
- * (b2) STAT_REPORT: Peers use the $BF(S, m, H)$ algorithm for compressing the subset of locally stored and integrity verified chunks (set S) in terms of a given swarm-ID, whenever reporting or updating its chunk availability information with the tracker. As the length of each BF-bitmap is constant ($O(m)$), this will greatly reduce the tracker's resource expenditure in communicating and storing such information for a large peer population.

c. Integration to the peer protocol

[I-D.ietf-ietf-ppsp-peer-protocol]:

- * (c1) HANDSHAKE: Peers exchange their local settings for the chunk compression schemes via HANDSHAKE messages.
- * (c2) HAVE: Peers use the $BF(S, m, H)$ algorithm for compressing the subset of locally stored and integrity verified chunks (e.g. set S_2 for Peer 2 in Figure 3) in terms of a given swarm-ID, whenever sharing its chunk availability information with another peer. The length of each BF-bitmap is constant ($O(m)$).
- * (c3) REQUEST: For a downloading peer to decide which neighbor to request for a given chunk_ID s , it uses the member query algorithm $MT(s, bf, m, H)$ on each neighbor's BF-bitmap bf . The

computation cost for this member check is constant ($O(m)$). It is also optional for a requesting peer to use BF-bitmap to indicate its data request to another peer, if needed.

3.3. Summary

3.3.1. Base TP protocol

In view of the wish that the "configuration processes for the PPSP Tracking facility, the service Portal and content sources are not standardized, enabling all the flexibility for implementers. But as there could have been different options to realize the chunk availability information (i.e. the chunk bitmap), it should be systematically configured so as to be mutually understandable to both parties. Therefore, we propose to add the following sentence to [I-D.ietf-ppsp-base-tracker-protocol] Section 5.1.1 "Installation and Initial Setup": "In case of a peer or the tracker wishes to exchange further information about the available peers in a flexible way, e.g. the chunk availability information of a specific peer in the same swarm could be represented in a various ways, there should be a way of negotiation/indication about the specific method/parameters in use, e.g. in the MPD file downloaded by the requesting peer from the web portal."

3.3.2. Extended TP protocol

The "StatisticsGroup" element in STAT_REPORT message has been extended to contain contentmap information, as shown in Figure 4).

Element Name or Attribute Name	Use	Description
StatisticsGroup	0...1	Provides statistic data on peer and content.
Stat @property	1...N M	Groups statistics property data. The property to be reported property values and elements in Table 5 of [I-D.ietf-ppsp-base-tracker-protocol]
Representation @id	0...N CM	Describes a component of content. Unique identifier for this Representation.
SegmentInfo @startIndex	1...N CM	Provides segment information by segment range. The chunkmap can be encoded in Base64 [RFC4648]. The index of the first media

@endIndex	CM	segment in the chunkmap report for this Representation. The index of the last media segment in the chunkmap report for this Representation.
@chunkmapSize	CM	Size of chunkmap reported.

Figure 4: Semantics of ContentGroup.

In order to incorporate BF-based representation, we propose to

- o use the "SegmentInfo" field for the BF-formatted bitmap; and
- o the three attributes of "startIndex", "endIndex" and "chunkmapSize" are meaningless in this case and are allowed to be absent.

3.4. Peer protocol

In peer protocols, the bitmap representation schemes are called "chunk addressing schemes", whose choice and configuration is included in the protocol options information exchanged between peers in communication initiation via HANDSHAKE messages. As shown by Figure 5), there are currently five options defined for bitmap representations.

Method	Description
0	32-bit bins
1	64-bit byte ranges
2	32-bit chunk ranges
3	64-bit bins
4	64-bit chunk ranges
5-255	Unassigned

Figure 5: PPSP Peer Chunk Addressing Methods.

After downloading and verifying a chunk, a peer updates its local chunk availability informations to each neighboring peer in the same swarm using HAVE message, which consists of a single chunk specification that states that the sending peer has those chunks and successfully checked their integrity. In order to incorporate BF-based representation, we propose to

- o add a defined value (e.g. 5) from "unassigned" value range of the PPSP peer chunk addressing methods for the BF-formatted bitmap;
and
- o use HAVE/REQUEST message to convey the BF-format array for the overall local chunk bitmap.

4. Open Issues

As the BF-based scheme is highly controllable and stable in terms of resource efficiency, it is viable to incorporate the contentmap information about each returned peer in the Peerlist without introducing considerable extra overhead. The current definition of PeerGroup element in the extended TP protocol needs to be extended to include "SegmentInfo" field as defined in StatisticsGroup element and used accordingly. However, the necessity of introducing other peer specific information into the TP protocol via Peerlist is still subject to discussion. In particular, this info exchange is far less frequent and accurate as compared to the peer protocols's HANDSHAKE and HAVE messages.

For the sake of simplicity, the above integration with tracker protocol does not consider the case of employing more than one bitmap compression algorithms for a single swarm. Meanwhile, it is noted that there are considerable contents in the peer protocol dealing with chunk addressing schemes inter-working between peers employing different schemes. However, in BF-based scheme, as for the irreversible nature of hash functions, it is not feasible to translate a BF-formatted array back into an original bitmap. Therefore, in order to keep compatible and open to peers using other addressing schemes, it is required that there be an original bitmap maintained by the local peer to allow the ability of communicating with peers using another addressing scheme.

Even if the chunk addressing scheme can be uniquely specified in a MPD file for a given swarm, there is still possibility that a peer who cannot support the relevant addressing scheme comes along and contact the tracker for joining intention. Therefore, there is need for such configuration exchange channel via TP protocol to allow mismatch detection and error handling. A possible solution could be to add an algorithm identifier attribute in the "SwarmId" field, and keep the value definition for different schemes consistent with the peer protocol's "Chunking Addressing Method" field.

5. Security Considerations

TBA

6. IANA Considerations

None.

7. References

7.1. Normative References

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

7.2. Informative References

- [BF-analysis] Broder, A. and M. Mitzenmacher, "Network applications of Bloom Filters: a survey", Internet Mathematics Vol. 1, No. 4, pp. 485-509, 2004.
- [BF-bloom] Bloom, B., "Space/time trade-offs in hash coding with allowable errors.", Communications of ACM Vol. 13, No. 7, pp. 422-426, 1970.
- [I-D.ietf-huang-extended-tracker-protocol] Huang, R., Zong, N., Cruz, R., Nunes, ., and J. Taveira, "PPSP Tracker Protocol-Extended Protocol", draft-ietf-huang-extended-tracker-protocol-02 (work in progress), February 2013.
- [I-D.ietf-ietf-ppsp-peer-protocol] Bakker, A., Petrocco, R., and V. Grishchenko, "Peer-to-Peer Streaming Peer Protocol (PPSPP)", draft-ietf-ppsp-peer-protocol-06 (work in progress), February 2013.
- [I-D.ietf-p2psip-base] Jennings, C., Lowekamp, B., Rescorla, E., Baset, S., and H. Schulzrinne, "REsource LOcation And Discovery (RELOAD) Base Protocol", draft-ietf-p2psip-base-26 (work in progress), February 2013.
- [I-D.ietf-ppsp-base-tracker-protocol]

Cruz, R., Nunes, M., Gu, Y., Xia, J., and J. Taveira,
"PPSP Tracker Protocol-Base Protocol (PPSP-TP/1.0)",
draft-ietf-ppsp-base-tracker-protocol-00 (work in
progress), February 2013.

[I-D.ietf-ppsp-problem-statement]
Zhang, Y. and N. Zong, "Problem Statement and Requirements
of Peer-to-Peer Streaming Protocol (PPSP)", draft-ietf-
ppsp-problem-statement-12 (work in progress), January
2013.

[P2P-limit]
Feng, C., Li, B., and B. Li, "Understanding the
performance gap between pull-based mesh streaming
protocols and fundamental limits", in Proc. of IEEE
INFOCOM , 2009.

[RFC1321] Rivest, . and . Newport, "RFC 1321: The MD5 message-digest
algorithm", April 1992.

Authors' Addresses

Lingli Deng
China Mobile

Email: denglingli@chinamobile.com

Jin Peng
China Mobile

Email: pengjin@chinamobile.com

Yunfei Zhang
CoolPad

Email: hishigh@gmail.com

Yihong Huang
Huawei

Email: rachel.huang@huawei.com

PPSP
INTERNET-DRAFT
Intended Status: Standards Track
Expires: April 24, 2014

Rui S. Cruz
Mario S. Nunes
IST/INESC-ID/INOV
Yingjie Gu
Jinwei Xia
Huawei
Joao P. Taveira
IST/INOV
Deng Lingli
China Mobile
October 21, 2013

PPSP Tracker Protocol-Base Protocol (PPSP-TP/1.0)
draft-ietf-ppsp-base-tracker-protocol-02

Abstract

This document specifies the base Peer-to-Peer Streaming Protocol-Tracker Protocol (PPSP-TP/1.0), an application-layer control (signaling) protocol for the exchange of meta information between trackers and peers. The specification outlines the architecture of the protocol and its functionality, and describes message flows, message processing instructions, message formats, formal syntax and semantics. The PPSP Tracker Protocol enables cooperating peers to form content streaming overlay networks to support near real-time Structured Media content delivery (audio, video, associated timed text and metadata), such as adaptive multi-rate, layered (scalable) and multi-view (3D) videos, in live, time-shifted and on-demand modes.

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/lid-abstracts.html>

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

Copyright and License Notice

Copyright (c) 2013 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. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1	Introduction	5
1.1	Terminology	5
2	Operation and Protocol Architecture Overview	8
2.1	Operation	8
2.2	Enrollment and Bootstrap	9
2.3	Architectural and Functional View	11
2.3.1	Messaging Model	11
2.3.2	Request/Response model	11
2.4	State Machines and Flows of the Protocol	13
2.4.1	Normal Operation	15
2.4.2	Error Conditions	15
3	Protocol Specification	17
3.1	Request/Response Syntax and Format	17
3.2	Semantics of PPSPTrackerProtocol elements	19
3.3	Request element in request Messages	22
3.4	Response element in response Messages	22
4	Request/Response Processing	24
4.1	CONNECT Request	24
4.2	FIND Request	28
4.3	STAT_REPORT Request	30
4.4	Error and Recovery conditions	31
5	Operations and Manageability	33
5.1	Operational Considerations	33
5.1.1	Installation and Initial Setup	33
5.1.2	Migration Path	34
5.1.3	Requirements on Other Protocols and Functional Components	34
5.1.4	Impact on Network Operation	34
5.1.5	Verifying Correct Operation	34
5.2	Management Considerations	34
5.2.1	Interoperability	34
5.2.2	Management Information	35
5.2.3	Fault Management	35
5.2.4	Configuration Management	35
5.2.5	Accounting Management	35
5.2.6	Performance Management	36
5.2.7	Security Management	36
6	Security Considerations	37
6.1	Authentication between Tracker and Peers	37
6.2	Content Integrity protection against polluting peers/trackers	37
6.3	Residual attacks and mitigation	38
6.4	Pro-incentive parameter trustfulness	38
7	Guidelines for Extending PPSP-TP	39
7.1	Forms of PPSP-TP Extension	39
7.2	Issues to Be Addressed in PPSP-TP Extensions	41

8	IANA Considerations	43
9	Acknowledgments	43
10	References	44
10.1	Normative References	44
10.2	Informative References	44
	Appendix A. Revision History	47
	Appendix B. PPSP-TP Message Syntax for HTTP/1.1	48
	B.1 Header Fields	49
	B.2 Methods	49
	B.3 Message Bodies	50
	B.4 Message Response Codes	50
	Appendix C. Use Scenarios	52
	C.1 Additional Terminology	52
	C.2 Functional Entities	53
	C.3 Streaming Capabilities	54
	C.4 NAT Traversal	54
	C.5 Content Information Metadata	54
	C.6 Authentication, Confidentiality, Integrity	55
	Appendix D. Implementation Options	55
	Authors' Addresses	56

1 Introduction

The Peer-to-Peer Streaming Protocol (PPSP) is composed of two protocols: the PPSP Tracker Protocol and the PPSP Peer Protocol [RFC6972] specifies that the Tracker Protocol should standardize format/encoding of information and messages between PPSP peers and PPSP trackers and also defines the requirements.

The PPSP Tracker Protocol provides communication between trackers and peers, by which peers send meta information to trackers, report streaming status and obtain peer lists from trackers.

The PPSP architecture requires PPSP peers able to communicate with a tracker in order to participate in a particular streaming content swarm. This centralized tracker service is used by PPSP peers for content registration and location.

The signaling and the media data transfer between PPSP peers is not in the scope of this specification.

This document describes the base PPSP Tracker protocol and how it satisfies the requirements for the IETF Peer-to-Peer Streaming Protocol, in order to derive the implications for the standardization of the PPSP streaming protocols and to identify open issues and promote further discussion.

1.1 Terminology

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 [RFC2119].

ABSOLUTE TIME: Absolute time is expressed as ISO 8601 [ISO.8601.2004] timestamps, using zero UTC offset (GMT). Fractions of a second may be indicated. Example for December 25, 2010 at 14h56 and 20.25 seconds: basic format 20101225T145620.25Z or extended format 2010-12-25T14:56:20.25Z.

CHUNK: A Chunk is a basic unit of data organized in P2P streaming for storage, scheduling, advertisement and exchange among peers.

CHUNK ID: A unique resource identifier for a SEGMENT CHUNK. The identifier type depends on the addressing scheme used, i.e., an integer, an HTTP-URL and possibly a byte-range, and is described in the MPD.

CONNECTION TRACKER: The node running the tracker service to which the PPSP peer will connect when it wants to get registered and join

the PPSP system.

LEECHER: A Peer that has not yet completed the transfer of all Chunks of the media content.

LIVE STREAMING: It refers to a scenario where all the audiences receive streaming content for the same ongoing event. It is desired that the lags between the play points of the audiences and streaming source be small.

MEDIA PRESENTATION DESCRIPTION (MPD): Formalized description for a media presentation, i.e., describes the structure of the media, namely, the Representations, the codecs used, the segments (Chunks), and the corresponding addressing scheme.

METHOD: The method is the primary function that a request from a peer is meant to invoke on a tracker. The method is carried in the request message itself.

ONLINE TIME: Online Time shows how long the peer has been in the P2P streaming system since it joined. This value indicates the stability of a peer, and can be calculated by tracker when necessary.

PEER: A PEER refers to a participant in a P2P streaming system that not only receives streaming content, but also caches and streams streaming content to other participants.

PEER ID: The identifier of a PEER such that other PEERs, or the TRACKER, can refer to the PEER by using its ID. The Peer ID is mandatory, can take the form of a universal unique identifier (UUID), defined in [RFC4122], and can be bound to a network address of the peer, i.e., an IP address, or a uniform resource identifier/locator (URI/URL) that uniquely identifies the corresponding peer in the network. The Peer ID and any required security certificates are obtained from an offline enrollment server.

PEER LIST: A list of PEERs which are in a same SWARM maintained by the TRACKER. A PEER can fetch the PEER LIST of a SWARM from the TRACKER or from other PEERs in order to know which PEERs have the required streaming content.

PPSP: The abbreviation of Peer-to-Peer Streaming Protocols. PPSP refer to the primary signaling protocols among various P2P streaming system components, including the TRACKER and the PEER.

PPSP-TP: The abbreviation of Peer-to-Peer Streaming Protocols - Tracker Protocol.

REPRESENTATION: Structured collection of one or more media components.

REQUEST: A message sent from a peer to a tracker, for the purpose of invoking a particular operation.

RESPONSE: A message sent from a tracker to a peer, for indicating the status of a request sent from the peer to the tracker.

SEEDER: A Peer that holds and shares the complete media content.

SEGMENT: The segment is a basic unit of partitioned streaming media, which is used by a peer for the purpose of storage, advertisement and exchange among peers.

SEGMENT CHUNK: For Structured Media contents, a segment may correspond to a set of nested and dependent CHUNKs, a set of correlated additive CHUNKs or a set of alternate REPRESENTATION CHUNKs.

SWARM: A SWARM refers to a group of PEERs who exchange data to distribute CHUNKs of the same content (e.g. video/audio program, digital file, etc.) at a given time.

SWARM ID: The identifier of a SWARM containing a group of PEERs sharing a common streaming content. The Swarm-ID may use a universal unique identifier (UUID), e.g., a 64 or 128 bit datum to refer to the content resource being shared among peers.

SUPER-NODE: A SUPER-NODE is a special kind of PEER deployed by ISPs. This kind of PEER is more stable with higher computing, storage and bandwidth capabilities than normal PEERs.

TRACKER: A TRACKER refers to a directory service that maintains a list of PEERs participating in a specific audio/video channel or in the distribution of a streaming file. Also, the TRACKER answers PEER LIST queries received from PEERs. The TRACKER is a logical component which can be centralized or distributed.

TRANSACTION ID: The identifier of a REQUEST from the PEER to the TRACKER. Used to disambiguate RESPONSES that may arrive in a different order of the corresponding REQUESTs.

VIDEO-ON-DEMAND (VoD): It refers to a scenario where different audiences may watch different parts of the same recorded streaming with downloaded content.

2 Operation and Protocol Architecture Overview

2.1 Operation

The functional entities related to PPSP protocols are the Client Media Player, the service Portal, the tracker and the peers. The complete description of Client Media Player and service Portal is not discussed here, as not in the scope the specification. The functional entities directly involved in the PPSP Tracker Protocol are trackers and peers (which may support different capabilities).

The Client Media Player is a logical entity providing direct interface to the end user at the client device, and includes the functions to select, request, decode and render contents. The Client Media Player may interface with the local peer application using request and response standard formats for HTTP Request and Response messages [RFC2616].

The service Portal is a logical entity typically used for client enrollment and content information publishing, searching and retrieval.

A Peer corresponds to a logical entity (typically in a user device) that actually participates in sharing a media content. Peers are organized in (various) swarms corresponding each swarm to the group of peers streaming a certain content at any given time.

The Tracker is a logical entity that maintains the lists of peers storing segments for a specific Live media channel or on-demand media streaming content, answers queries from peers and collects information on the activity of peers. While a tracker may have an underlying implementation consisting of more than one physical node, logically the tracker can most simply be thought of as a single element, and in this document it will be treated as a single logical entity.

The Tracker Protocol is not used to exchange actual content data (either on-demand or Live streaming) with peers, but information about which peers can provide the content.

When a peer wants to receive streaming of a selected content (Leech mode):

1. Peer connects to a connection tracker and joins a swarm.
2. Peer acquires a list of other peers in the swarm from the connection tracker.
3. Peer exchanges its content availability with the peers on the obtained peer list (via peer protocol).

4. Peer identifies the peers with desired content.
5. Peer requests content from the identified peers (peer protocol).

When a peer wants to share streaming contents (Seeder mode) with other peers:

1. Peer connects to the connection tracker.
2. Peer sends information to the connection tracker about the swarms it belongs to (joined swarms).

After having been disconnected due to some termination condition, a peer can resume previous activity by connecting and re-joining the corresponding swarm(s).

2.2 Enrollment and Bootstrap

In order to join an existing P2P streaming service and to participate in content sharing, any peer must first locate a tracker, using the following (not exhaustive) typical methods (illustrated in Figures 1 and 2):

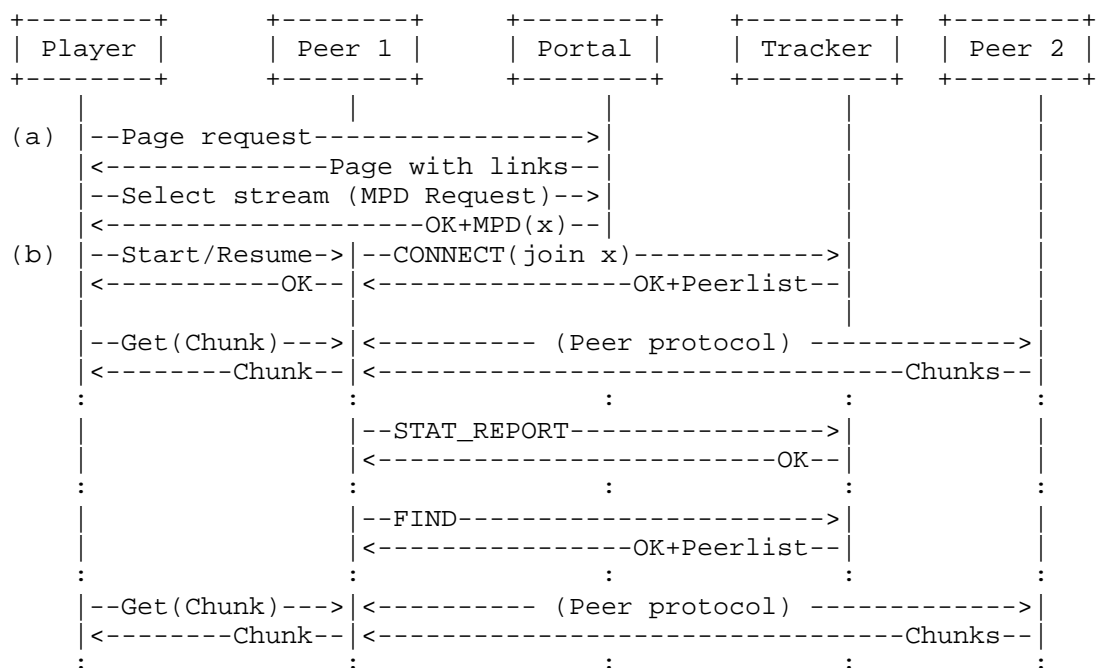


Figure 1: A typical PPSP session for watching a streaming content.

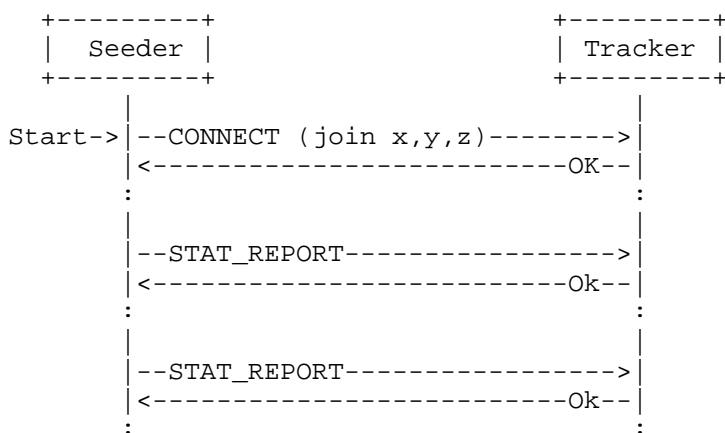


Figure 2: A typical PPSP session for a streaming Seeder.

1. From a content service provider provisioning mechanism: this is the typical case used for the provider Seeders (edge caches and/or Media Servers), where some "Start" activation signal triggers the process for contents x, y, z (Figure 2).
2. From a web page: a Publishing and Searching Portal may provide tracker addressing information to end users.
3. From the Media Presentation Description (MPD) file of a content: this meta-info file must contain information about the address of one or more trackers (that can be grouped by tiers of priority) which are controlling the swarm for that media content, as illustrated in Figure 1 for a content x.

As illustrated in Figure 1, a peer may initiate a stream using the above method 3 starting at point (a), or resume a previously initiated stream, using also method 3 but starting at point (b).

In order to be able to bootstrap, a peer must first obtain a Peer ID (identifier of the peer) and any required security certificates or authorization tokens from an enrollment service (end user registration). The specification of the format of the Peer ID is not in the scope of this document.

The specification of the mechanisms used by the Client Media Player and the peer (to signal start/resume streams or request media chunks), obtain a Peer ID, security certificates or tokens are not in the scope of this document.

2.3 Architectural and Functional View

The PPSP Tracker Protocol architecture is intended to be compatible with the web infrastructure. PPSP-TP is designed with a layered approach i.e., a PPSP-TP Request/Response layer, a Message layer and a Transport layer. The PPSP-TP Request/Response layer deals with the interactions between tracker and peers using Request and Response codes (see Figure 3).

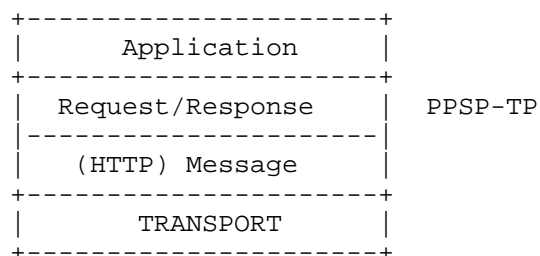


Figure 3: Abstract layering of PPSP-TP.

The Message layer deals with the framing format, for encoding and transmitting the data through the underlying transport protocol, and the asynchronous nature of the interactions between tracker and peers.

The Transport layer is responsible for the actual transmission of requests and responses over network transports, including the determination of the connection to use for a request or response message when using a connection-oriented transport like TCP [RFC0793], or TLS [RFC5246] over it.

2.3.1 Messaging Model

The messaging model of PPSP-TP aligns with HTTP protocol and the semantics of its messages, currently in version 1.1 [RFC2616], but intended to support future versions of HTTP. The exchange of messages of PPSP-TP is envisioned to be performed over a stream-oriented reliable transport protocol, like TCP [RFC0793].

Appendix B describes the message syntax when using HTTP/1.1.

2.3.2 Request/Response model

PPSP-TP is a text-based protocol, uses the UTF-8 character set [RFC3629] and the protocol messages are either requests from peers to a tracker service, or responses from a tracker service to peers.

PPSP-TP Request and Response semantics are carried as entities (header and body) in messages which correspond to either HTTP request methods or HTTP response codes, respectively.

Requests are sent, and responses returned to these requests. A single request generates a single response (neglecting fragmentation of messages in transport).

The response codes are consistent with HTTP response codes, however, not all HTTP response codes are used for the PPSP-TP (section 3 and Appendix B.4).

The Request Messages of the base protocol are listed in Table 1:

+-----+	
	PPSP-TP/1.0
	Req. Messages
+-----+	
	CONNECT
	FIND
	STAT_REPORT
+-----+	

Table 1: Request Messages

CONNECT: This Request message is an "action signal" used when a peer registers in the tracker (or if already registered) to notify it about the participation in named swarm(s). The tracker records the Peer ID, connect-time (referenced to the absolute time), peer IP addresses (and associated location information), link status and Peer Mode for the named swarm(s). The tracker also changes the content availability of the valid named swarm(s), i.e., changes the peers lists of the corresponding swarm(s) for the requester Peer ID. On receiving a CONNECT message, the tracker first checks the peer mode type (SEED/LEECH) for the specified swarm(s) and then decides the next steps (more details are referred in section 4.1)

FIND: This Request message is an "action signal" used by peers to request to the tracker, whenever needed, a list of peers active in the named swarm. On receiving a FIND message, the tracker finds the peers, listed in content status of the specified swarm that can satisfy the requesting peer's requirements, returning the list to the requesting peer. To create the peer list, the tracker may take peer status, capabilities and peers priority into consideration. Peer priority may be determined by network topology preference,

operator policy preference, etc.

STAT_REPORT: This Request message is an "information signal" that allows an active peer to send status (and optionally statistic data) to the tracker to signal continuing activity. This request message **MUST** be sent periodically to the tracker while the peer is active in the system.

2.4 State Machines and Flows of the Protocol

The state machine for the tracker is very simple, as shown in Figure 4. Peer ID registrations represent a dynamic piece of state maintained by the network.




Figure 4: Tracker State Machine

When there are no peers connected in the tracker, the state machine is in the INIT state.

When the "first" peer connects for registration with its Peer ID, the state machine moves from INIT to STARTED. As long as there is at least one active registration of a Peer ID, the state machine remains in the STARTED state. When the "last" Peer ID is removed, the state machine transitions to TERMINATED. From there, it immediately transitions back to the INIT state. Because of that, the TERMINATED state here is transient.

In addition to the tracker state machine, each peer is modeled with its own transaction state machine (Figure 5), instantiated per Peer ID in the tracker, and deleted when it is removed.

Unlike the tracker state machine, which exists even when no Peer IDs are registered, the "per-Peer-ID" transaction state machine is instantiated only when the Peer ID starts registration in the tracker, and is deleted when the Peer ID is de-registered/removed. This allows for an implementation optimization whereby the tracker can destroy the objects associated with the "per-Peer-ID" transaction state machine once it enters the TERMINATE state (Figure 5).

When a new Peer ID is added, the corresponding "per-Peer-ID" state machine is instantiated, and it moves into the PEER REGISTERED state.

Because of that, the START state here is transient.

When the Peer ID is no longer bound to a registration, the "per-Peer-ID" state machine moves to the TERMINATE state, and the state machine is destroyed.

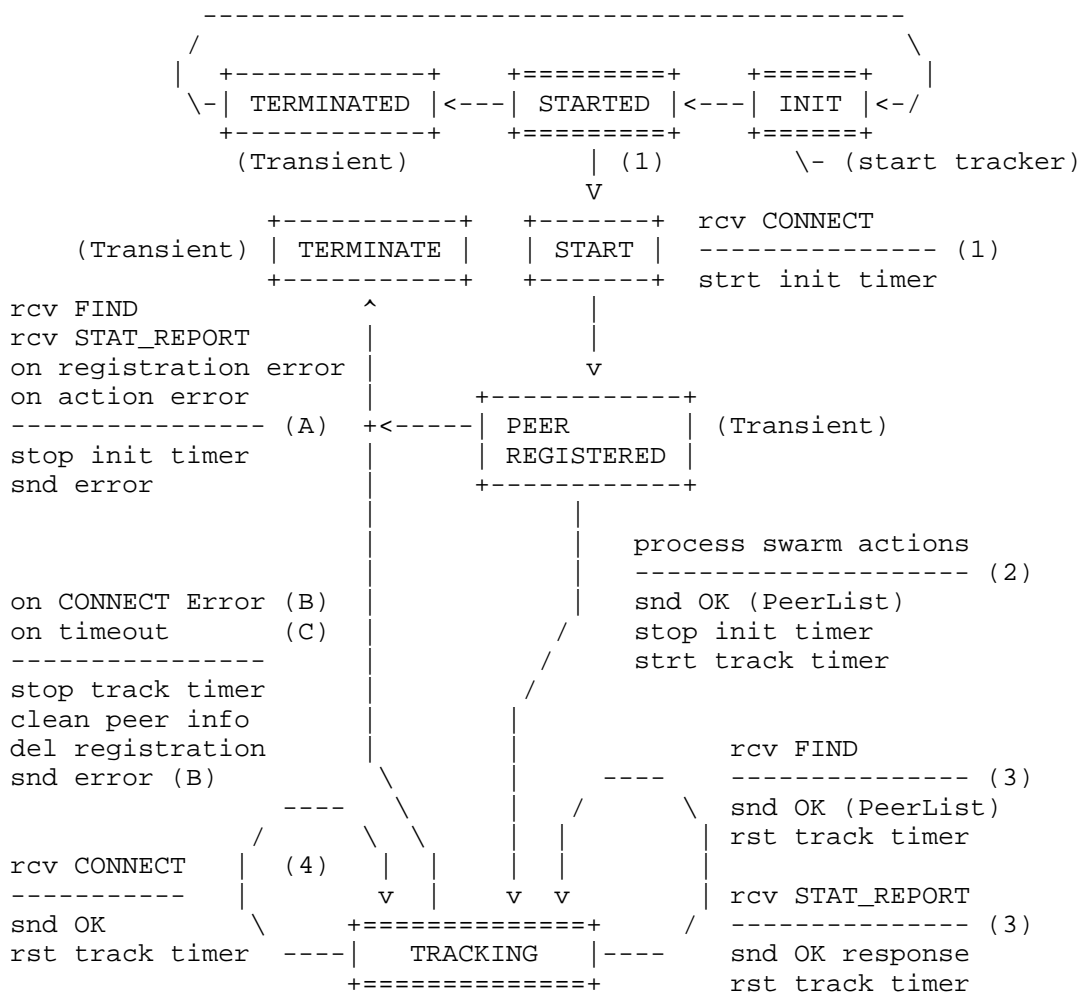


Figure 5: Per-Peer-ID Transaction State Machine and Flow Diagram

During the lifetime of streaming activity of a peer, the "per-Peer-ID" transaction state machine progresses from one state to another in response to various events. The events that may potentially advance the state include:

- o Reception of CONNECT, FIND and STAT_REPORT messages, or
- o Timeout events.

The state diagram in Figure 5 illustrates state changes, together with the causing events and resulting actions. Specific error conditions are not shown in the state diagram.

2.4.1 Normal Operation

On normal operation the process consists of the following steps:

- 1) When a Peer wants to access the system it needs to register on a tracker by sending a CONNECT message asking for the swarm(s) it wants to join. This CONNECT request from a new Peer ID triggers the instantiation of a "per-Peer-ID" State Machine. In the START state of the new "per-Peer-ID" SM, the tracker initiates the registration of the Peer ID and associated information (IP addresses), starts the "init timer" and moves to PEER REGISTERED state.
- 2) In PEER REGISTERED state, if Peer ID is valid, the tracker either a) processes the requested action(s) for the valid swarm information contained in the CONNECT request and in case of success the tracker stops the "init timer", starts the "track timer" and sends the response to the peer (the response MAY contain the appropriate list of peers for the joining swarm(s), as detailed in section 4.1, or b) moves the valid FIND request to TRACKING state.
- 3) In TRACKING state, STAT_REPORT or FIND messages received from that Peer ID will reset the "track timer" and are respectively responded with a) a successful condition, b) a successful condition containing the appropriate list of peers for the named swarm (section 4.2).
- 4) While TRACKING, a CONNECT message received from that Peer ID with valid swarm actions information (section 4.1) resets the "track timer" and is responded with a successful condition.

2.4.2 Error Conditions

Peers MUST NOT generate protocol elements that are invalid. However, several situations of a peer may lead to abnormal conditions in the interaction with the tracker. The situations may be related with peer malfunction or communications errors. The tracker reacts to the abnormal situations depending on its current state related to a Peer ID, as follows:

- A) At PEER REGISTERED state, when a CONNECT Request only contains invalid swarm actions (section 4.1), the tracker responds with error code 403 Forbidden, deletes the registration, transition to TERMINATE state for that Peer ID and the SM is destroyed.

At the PEER REGISTERED state, if the Peer ID is considered invalid (in the case of a CONNECT request or in the case of FIND or STAT_REPORT requests received from an unregistered Peer ID), the tracker responds with either error codes 401 Unauthorized or 403 Forbidden (described in section 4.4), transitions to TERMINATE state for that Peer ID and the SM is destroyed.

- B) At the TRACKING state (while the "track timer" has not expired) receiving a CONNECT message from that Peer ID with invalid swarm actions (section 4.1) is considered an error condition. The tracker responds with error code 403 Forbidden (described in section 3), stops the "track timer", deletes the registration, transitions to TERMINATE state for that Peer ID and the SM is destroyed.
- C) In TRACKING state, without receiving messages from the peer, on timeout (track timer) the tracker cleans all the information associated with the Peer ID in all swarms it was joined, deletes the registration, transitions to TERMINATE state for that Peer ID and and the SM is destroyed.

NOTE: These situations may correspond to a malfunction at the peer or to malicious conditions. Therefore, as preventive measure, the tracker proceeds to TERMINATE state for the Peer ID by de-registering the peer and cleaning all peer information.

3 Protocol Specification

3.1 Request/Response Syntax and Format

The message-body for Requests and Responses requiring it, correspond to a XML document [XML].

The XML message-body MUST begin with an XML declaration line specifying the version of XML being used and indicating the character encoding, which SHOULD be "UTF-8". The root element MUST begin with an <PPSPTrackerProtocol> element. This element identifies the start of an PPSP-TP protocol element, the namespace used within the protocol, and the location of the protocol schema.

The generic format of a Request is the following:

```
<?xml version="1.0" encoding="UTF-8"?>
<PPSPTrackerProtocol xmlns="TBD" schemaLocation="TBD" version="1.0">
  <Request></Request>
  <TransactionID></TransactionID>
  <PeerID></PeerID>
  <SwarmID></SwarmID>
  <PeerNum></PeerNum>
  <PeerGroup></PeerGroup>
  <StatisticsGroup></StatisticsGroup>
</PPSPTrackerProtocol>
```

The generic format of a Response is the following:

```
<?xml version="1.0" encoding="UTF-8"?>
<PPSPTrackerProtocol xmlns="TBD" schemaLocation="TBD" version="1.0">
  <Response></Response>
  <TransactionID></TransactionID>
  <PeerGroup></PeerGroup>
</PPSPTrackerProtocol>
```

The Request element MUST be present in requests and corresponds to the request method type for the message.

The Response element MUST be present in responses and corresponds to the response method type of the message.

The element TransactionID MUST be present in requests to uniquely identify the transaction. Responses to completed transactions use the same TransactionID as the request they correspond to.

The element SwarmID MUST be present in requests to identify the actions to be taken in the specified swarms.

The version of PPSP-TP being used is indicated by the attribute @version of the root element, specified in a XML Schema ([XMLSchema.1] and [XMLSchema.2]) with XML namespaces [XMLNameSpace] to identify protocol grammars.

All Request messages MUST contain a PeerID element to uniquely identify the peer (Peer ID) in the network.

The PeerID information may be present on the following levels:

- On PPSPTrackerProtocol level in PPSPTrackerProtocol.PeerID element. For details refer to subsection 3.2 Table 2.
- On PeerGroup level in PeerGroup.PeerInfo.PeerID element (attribute @swarID). For details refer to subsection 3.2 Table 3.

The SwarmID element MUST be present in CONNECT and FIND Requests and SHOULD be present in STAT_REPORT Requests on the following levels:

- On PPSPTrackerProtocol level in PPSPTrackerProtocol.SwarmID element. For details refer to subsection 3.2 Table 2.
- On StatisticsGroup level in StatisticsGroup.Stat.SwarmID element. For details refer to subsection 3.2 Table 4.

The PeerNum element MAY be present in CONNECT and FIND requests and MAY contain the attribute @abilityNAT to inform the tracker on the preferred type of peers, in what concerns their NAT traversal situation, to be returned in a peer list.

The StatisticsGroup element MAY be present in STAT_REPORT requests. Details of usage in subsection 4.2.

The semantics of the attributes and elements within a PPSPTrackerProtocol root element is described in subsection 3.2.

Request and Response processing is provided in section 4 for each message.

3.2 Semantics of PPSPTrackerProtocol elements

The semantics of PPSPTrackerProtocol elements and attributes are described in the following tables.

Element Name or Attribute Name	Use	Description
PPSPTrackerProtocol	1	The root element
@version	M	Provides the version of PPSP-TP
Request	0...1	Provides the request method and MUST be present in Request
Response	0...1	Provides the response method and MUST be present in Response
TransactionID	M	Root transaction Identification
Result	0...N	Result of @action MUST be present in Responses
@transactionID	CM	Identifier of the @action
PeerID	0...1	Peer Identifier MUST be present in Request
SwarmID	0...N	Swarm Identifier MUST be present in Requests
@action	CM	Must be set to JOIN or LEAVE
@peerMode	CM	Mode of Peer participation in the swarm, LEECH or SEED
@transactionID	CM	Identifier for the @action
PeerNUM	0...1	Maximum peers in PeerList, =<30
@abilityNAT	OP	Type of NAT traversal peers, as No-NAT, STUN, TURN or PROXY
@concurrentLinks	OP	Concurrent connectivity level of peers, HIGH, LOW or NORMAL
@onlineTime	OP	Availability or online duration of peers, HIGH or NORMAL
@uploadBWlevel	OP	Upload bandwidth capability of peers, HIGH or NORMAL
PeerGroup	0...1	Information on peers (Table 3)
StatisticsGroup	0...1	Statistic data of peer (Table 4)
Legend: Use for attributes: M=Mandatory, OP=Optional, CM=Conditionally Mandatory Use for elements: minOccurs...maxOccurs (N=unbounded) Elements are represented by their name (case-sensitive) Attribute names (case-sensitive) are preceded with an @		

Table 2: Semantics of the base PPSPTrackerProtocol.

Element Name or Attribute Name	Use	Description
PeerGroup	0...1	Contains description of peers
PeerInfo	1...N	Provides information on a peer
@swarmID	0...1	Swarm Identifier
PeerID	0...1	Peer Identifier
PeerAddress	0...N	IP Address information
@addrType	M	Type of IP address, which can be ipv4 or ipv6
@priority	CM	The priority of this interface
@type	CM	Describes the address for NAT traversal, which can be HOST REFLEXIVE or PROXY
@connection	OP	Access type (3G, ADSL, etc.)
@asn	OP	Autonomous System Number
@ip	M	IP address value
@port	M	IP service port value
@peerProtocol	OP	PPSP Peer Protocol supported
Legend: Use for attributes: M=Mandatory, OP=Optional, CM=Conditionally Mandatory Use for elements: minOccurs...maxOccurs (N=unbounded) Elements are represented by their name (case-sensitive) Attribute names (case-sensitive) are preceded with an @		

Table 3: Semantics of PeerGroup.

If STUN-like functions are enabled in the tracker and a PPSP-ICE method [RFC5245] is used the attributes @type and @priority MUST be returned with the transport address candidates in responses to CONNECT requests.

The @asn attribute MAY be used to inform about the network location, in terms of Autonomous System, for each of the active public network interfaces of the peer.

The @connection attribute is informative on the type of access network of the respective interface.

Element Name or Attribute Name	Use	Description
StatisticsGroup	0...1	Provides statistic data on peer and content
Stat	1...N	Groups statistics property data
@property	M	The property to be reported Property values and elements in Table 5
Legend: Use for attributes: M=Mandatory, OP=Optional, CM=Conditionally Mandatory Use for elements: minOccurs...maxOccurs (N=unbounded) Elements are represented by their name (case-sensitive) Attribute names (case-sensitive) are preceded with an @		

Table 4: Semantics of StatisticsGroup.

The Stat element is used to describe several properties relevant to the P2P network. These properties can be related with stream statistics and peer status information. Each Stat element will correspond to a @property type and several Stat blocks can be reported in a single STAT_REPORT message, corresponding to some or all the swarms the peer is actively involved.

"Property Name" Element Name or Attribute Name	Use	Description
"StreamStatistics"		Property for swarm statistics
SwarmID	0...1	Swarm Identifier
UploadedBytes	0...1	Bytes sent to swarm
DownloadedBytes	0...1	Bytes received from swarm
AvailBandwidth	0...1	Upstream Bandwidth available
Legend: Use for attributes: M=Mandatory, OP=Optional, CM=Conditionally Mandatory Use for elements: minOccurs...maxOccurs (N=unbounded) Elements are represented by their name (case-sensitive) Attribute names (case-sensitive) are preceded with an @		

Table 5: Basic StatisticsGroup property blocks.

Other properties may be defined, related, for example, with incentives and reputation mechanisms like "peer online time", or connectivity conditions like physical "link status", etc.

For that purpose, the Stat element may be extended to provide additional scheme specific information for new @property groups, new sibling elements and new attributes (guidelines in section 7).

3.3 Request element in request Messages

Table 6 defines the valid string representations for the requests. These values MUST be treated as case-sensitive.

XML Request Methods String Values
CONNECT
FIND
STAT_REPORT

Table 6: Valid Strings for Request element of requests.

3.4 Response element in response Messages

Table 7 defines the valid string representations for Response messages that require message-body. These values MUST be treated as case-sensitive.

Response messages not requiring message-body only use the standard HTTP Status-Code and Reason-Phrase (appended, if appropriate, with detail phrase, as described in section 4.4).

XML Response Method String Values	HTTP Status-Code and Reason-Phrase
SUCCESSFUL	200 OK
AUTHENTICATION REQUIRED	401 Unauthorized

Table 7: Valid Strings for Response element of responses.

SUCCESSFUL: indicates that the request has been processed properly and the desired operation has completed. The body of the response message includes the requested information and MUST include the same TransactionID of the corresponding request.

In CONNECT Request: returns information about the successful registration of the peer and/or of each swarm @action requested. MAY additionally return the list of peers corresponding to the join @action requested.

In FIND Request: returns the list of peers corresponding to the requested scope.

In STAT_REPORT Request: confirms the success of the requested operation.

AUTHENTICATION REQUIRED: authentication is required for the peer to make the request.

4 Request/Response Processing

When a PPSP-TP message is received some basic processing is performed, regardless of the message type.

Upon reception, a message is examined to ensure that it is properly formed. The receiver **MUST** check that the HTTP message itself is properly formed, and if not, appropriate standard HTTP errors **MUST** be generated. The receiver must also verify that the XML body is properly formed. In case of error due to malformed messages appropriate responses **MUST** be returned, as described in section 4.4.

4.1 CONNECT Request

This method is used when a peer registers to the system and/or requests swarm actions.

The peer **MUST** properly form the XML message-body, set the Request method to **CONNECT**, generate and set the TransactionIDs, set the PeerID with the identifier of the peer and include SwarmID action(s) to be performed. The peer **MAY** also include the IP addresses of its network interfaces in the **CONNECT** message.

The following example of the message-body of a **CONNECT** Request corresponds to a peer that wants to start (or re-start) sharing its previously streamed contents (peerMode is of **SEED**). Note for this case that the peer also requests from the Tracker an appropriate list of peers (PeerNum element) already active in the swarm, i.e., a list of 15 peers having STUN capabilities in terms of NAT. In the case of a Super-Node peer of an ISP, the **CONNECT** request would be similar but, optionally not including the PeerNum element:

```
<?xml version="1.0" encoding="UTF-8"?>
<PPSPTrackerProtocol xmlns="TBD" schemaLocation="TBD" version="1.0">
  <Request>CONNECT</Request>
  <PeerID>656164657220</PeerID>
  <PeerNum abilityNAT="STUN">15</PeerNum>
  <SwarmID action="JOIN" peerMode="SEED"
    transactionID="12345.1">1111</SwarmID>
  <SwarmID action="JOIN" peerMode="SEED"
    transactionID="12345.2">2222</SwarmID>
  <TransactionID>12345.0</TransactionID>
</PPSPTrackerProtocol>
```

Another example of the message-body of a **CONNECT** Request corresponds to a peer Leecher requesting join to a swarm, in order to start receiving the stream, and providing optional information on the addresses of its network interface(s):

```

<?xml version="1.0" encoding="UTF-8"?>
<PPSPTrackerProtocol xmlns="TBD" schemaLocation="TBD" version="1.0">
  <Request>CONNECT</Request>
  <PeerID>656164657221</PeerID>
  <PeerNum abilityNAT="STUN">5</PeerNum>
  <SwarmID action="JOIN" peerMode="LEECH"
    transactionID="12345.1">1111</SwarmID>
  <TransactionID>12345.0</TransactionID>
  <PeerGroup>
    <PeerInfo>
      <PeerAddress addrType="ipv4" ip="192.0.2.2" port="80"
        priority="1" peerProtocol="PPSP-PP" />
      <PeerAddress addrType="ipv6" ip="2001:db8::2" port="80"
        priority="2" peerProtocol="PPSP-PP" />
    </PeerInfo>
  </PeerGroup>
</PPSPTrackerProtocol>

```

The next example of the message-body of a CONNECT Request corresponds to a peer Leecher "leaving" a previously joined swarm and requesting join to a new swarm:

```

<?xml version="1.0" encoding="UTF-8"?>
<PPSPTrackerProtocol xmlns="TBD" schemaLocation="TBD" version="1.0">
  <Request>CONNECT</Request>
  <PeerID>656164657221</PeerID>
  <PeerNum abilityNAT="STUN">5</PeerNum>
  <SwarmID action="LEAVE" peerMode="LEECH"
    transactionID="12345.1">1111</SwarmID>
  <SwarmID action="JOIN" peerMode="LEECH"
    transactionID="12345.2">2222</SwarmID>
  <TransactionID>12345.0</TransactionID>
</PPSPTrackerProtocol>

```

When receiving a well-formed CONNECT Request message, the tracker MAY, when applicable, start by pre-processing the peer authentication information (provided as Authorization scheme and token in the HTTP message) to check whether it is valid and that it can connect to the service, then proceed to register the peer in the service and perform the swarm actions requested. In case of success a Response message with a corresponding response value of SUCCESSFUL will be generated.

The element SwarmID MUST be present with cardinality 1 to N, each containing the request for @action, the @peerMode of the peer and the child @transactionID for that swarm. The @peerMode element MUST be set to the type of participation of the peer in the swarm (SEED or LEECH).

The valid sets of SwarmID @action combinations for the CONNECT Request logic in PPSP-TP base protocol are summarized in Table 8 (referring to the tracker "per-Peer-ID" state machine in section 2.4).

SwarmID Elements	@peerMode value	@action value	Initial State	Final State	Request validity
1	LEECH	JOIN	START	TRACKING	Valid
1	LEECH	LEAVE	START	TERMINATE	Invalid
1	LEECH	LEAVE	TRACKING	TERMINATE	Valid
1 1	LEECH LEECH	JOIN LEAVE	START	TERMINATE	Invalid
1 1	LEECH LEECH	JOIN LEAVE	TRACKING	TRACKING	Valid
N	SEED	JOIN	START	TRACKING	Valid
N	SEED	JOIN	TRACKING	TERMINATE	Invalid
N	SEED	LEAVE	TRACKING	TERMINATE	Valid

Table 8: Validity of @action combinations in CONNECT Request.

The element PeerInfo, if present, MAY contain multiple PeerAddress child elements with attributes @addrType, @ip, @port and @peerProtocol, and optionally @priority and @type (if PPSP-ICE NAT traversal techniques are used) corresponding to each of the network interfaces the peer wants to advertise.

The element PeerNum indicates to the tracker the number of peers to be returned in a list corresponding to the indicated properties, being @abilityNAT for NAT traversal (considering that PPSP-ICE NAT traversal techniques may be used), and optionally @concurrentLinks, @onlineTime and @uploadBWlevel for the preferred capabilities. If STUN-like function is enabled in the tracker, the response MAY include the peer reflexive address.

The response MUST have the same TransactionID values as the corresponding request and actions.

The next example illustrates a Response for the CONNECT Request:

```

<?xml version="1.0" encoding="UTF-8"?>
<PPSPTrackerProtocol xmlns="TBD" schemaLocation="TBD" version="1.0">
  <Response>SUCCESSFUL</Response>
  <TransactionID>
    <Result transactionID="12345.0">200 OK</Result>
    <Result transactionID="12345.1">200 OK</Result>
    <Result transactionID="12345.2">200 OK</Result>
  </TransactionID>
  <PeerGroup>
    <PeerInfo>
      <PeerID>656164657221</PeerID>
      <PeerAddress addrType="ipv4" ip="198.51.100.1" port="80"
        priority="1"
        type="REFLEXIVE"
        connection="3G"
        asn="64496" />
    </PeerInfo>
    <PeerInfo swarmID="2222">
      <PeerID>956264622298</PeerID>
      <PeerAddress addrType="ipv4" ip="198.51.100.22" port="80"
        asn="64496" peerProtocol="PPSP-PP" />
    </PeerInfo>
    <PeerInfo swarmID="2222">
      <PeerID>3332001256741</PeerID>
      <PeerAddress addrType="ipv4" ip="198.51.100.201" port="80"
        asn="64496" peerProtocol="PPSP-PP" />
    </PeerInfo>
  </PeerGroup>
</PPSPTrackerProtocol>

```

The Response MUST include a PeerGroup with PeerInfo data of the requesting peer public IP address. If STUN-like function is enabled in the tracker, the PeerAddress includes the attribute @type with a value of REFLEXIVE, corresponding to the transport address "candidate" of the peer. The PeerGroup MAY also include PeerInfo data corresponding to the Peer IDs and public IP addresses of the selected active peers in the requested swarm. The tracker MAY also include the attribute @asn with network location information of the transport address, corresponding to the Autonomous System Number of the access network provider of the referenced peer.

In case the @peerMode is SEED, the tracker responds with a SUCCESSFUL response and enters the peer information into the corresponding swarm activity. In case the @peerMode is LEECH (or if the peer Seeder includes a PeerNum element in the request) the tracker will search and select an appropriate list of peers satisfying the conditions set by the requesting peer. The peer list returned MUST contain the Peer IDs and the corresponding IP Addresses. To create the peer list, the

tracker may take peer status and network location information into consideration, to express network topology preferences or Operators' policy preferences, with regard to the possibility of connecting with other IETF efforts such as ALTO [I.D.ietf-alto-protocol].

IMPLEMENTATION NOTE: If no PeerNum attributes are present in the request the tracker MAY return a random sample from the peer population.

4.2 FIND Request

This method allows peers to request to the tracker, whenever needed, a new peer list for the swarm.

The peer MUST properly form the XML message-body, set the request method to FIND, set the PeerID with the identifier of the peer, set the SwarmID with the identifier of the swarm the peer is interested in.

The peer MUST generate and set the TransactionID for the request.

An example of the message-body of a FIND Request is the following:

```
<?xml version="1.0" encoding="UTF-8"?>
<PPSPTrackerProtocol xmlns="TBD" schemaLocation="TBD" version="1.0">
  <Request>FIND</Request>
  <PeerID>656164657221</PeerID>
  <SwarmID>1111</SwarmID>
  <TransactionID>12345</TransactionID>
  <PeerNum abilityNAT="STUN"
    concurrentLinks="HIGH"
    onlineTime="NORMAL"
    uploadBWlevel="NORMAL">5</PeerNum>
</PPSPTrackerProtocol>
```

The FIND request MAY include a PeerNum element to indicate to the tracker the maximum number of peers to be returned in a list corresponding to the indicated properties, being @abilityNAT for NAT traversal (considering that PPSP-ICE NAT traversal techniques may be used), and optionally @concurrentLinks, @onlineTime and @uploadBWlevel for the preferred capabilities.

When receiving a well-formed FIND Request the tracker processes the information to check if it is valid. In case of success a response message with a Response value of SUCCESSFUL will be generated and the tracker will include the appropriate list of peers satisfying the conditions requested. The peer list returned MUST contain the Peer

IDs and the corresponding IP Addresses.

The tracker may take peer status and network location information into consideration when selecting the peer list to return, to express network topology preferences or Operators' policy preferences, with regard to the possibility of connecting with other IETF efforts such as ALTO [I.D.ietf-alto-protocol].

An example of a Response message for the FIND Request is the following:

```
<?xml version="1.0" encoding="UTF-8"?>
<PPSPTrackerProtocol xmlns="TBD" schemaLocation="TBD" version="1.0">
  <Response>SUCCESSFUL</Response>
  <TransactionID>12345</TransactionID>
  <PeerGroup>
    <PeerInfo swarmID="1111">
      <PeerID>956264622298</PeerID>
      <PeerAddress addrType="ipv4" ip="198.51.100.22" port="80"
        asn="64496" />
    </PeerInfo>
    <PeerInfo swarmID="1111">
      <PeerID>3332001256741</PeerID>
      <PeerAddress addrType="ipv4" ip="198.51.100.201" port="80"
        asn="64496" />
    </PeerInfo>
  </PeerGroup>
</PPSPTrackerProtocol>
```

The Response MUST include a PeerGroup with PeerInfo data that includes the public IP addresses of the selected active peers in the swarm.

The tracker MAY also include the attribute @asn with network location information of the transport addresses of the peers, corresponding to the Autonomous System Numbers of the access network provider of each peer in the list.

The response MAY also include a PeerGroup with PeerInfo data that includes the requesting peer public IP address. If STUN-like function is enabled in the tracker, the PeerAddress includes the attribute @type with a value of REFLEXIVE, corresponding to the transport address "candidate" of the peer.

An example of a Response message for the FIND Request including the requesting peer public IP address is the following:

```
<?xml version="1.0" encoding="UTF-8"?>
<PPSPTrackerProtocol xmlns="TBD" schemaLocation="TBD" version="1.0">
  <Response>SUCCESSFUL</Response>
  <TransactionID>12345</TransactionID>
  <PeerGroup>
    <PeerInfo>
      <PeerID>656164657221</PeerID>
      <PeerAddress addrType="ipv4" ip="198.51.100.1" port="80"
        priority="1"
        type="REFLEXIVE"
        connection="3G"
        asn="64496" />
    </PeerInfo>
    <PeerInfo swarmID="1111">
      <PeerID>956264622298</PeerID>
      <PeerAddress addrType="ipv4" ip="198.51.100.22" port="80"
        asn="64496" />
    </PeerInfo>
    <PeerInfo swarmID="1111">
      <PeerID>3332001256741</PeerID>
      <PeerAddress addrType="ipv4" ip="198.51.100.201" port="80"
        asn="64496" />
    </PeerInfo>
  </PeerGroup>
</PPSPTrackerProtocol>
```

IMPLEMENTATION NOTE: If no PeerNum attributes are present in the request the tracker MAY return a random sample from the peer population.

4.3 STAT_REPORT Request

This method allows peers to send status and statistic data to trackers. The method is initiated by the peer, periodically while active.

The peer MUST properly form the XML message-body, set the Request method to STAT_REPORT, set the PeerID with the identifier of the peer, and generate and set the TransactionID.

The report MAY include a StatisticsGroup containing multiple Stat elements describing several properties relevant to the P2P network. These properties can be related with stream statistics and peer status information. Other properties may be defined (guidelines in section 7.1) related for example, with incentives and reputation mechanisms. In case no StatisticsGroup is included, the STAT_REPORT is used as a "keep-alive" message to prevent the Tracker from de-registering the peer when track timer expires.

An example of the message-body of a STAT_REPORT Request is the following:

```
<?xml version="1.0" encoding="UTF-8"?>
<PPSPTrackerProtocol xmlns="TBD" schemaLocation="TBD" version="1.0">
  <Request>STAT_REPORT</Request>
  <PeerID>656164657221</PeerID>
  <TransactionID>12345</TransactionID>
  <StatisticsGroup>
    <Stat property="StreamStatistics">
      <SwarmID>1111</SwarmID>
      <UploadedBytes>512</UploadedBytes>
      <DownloadedBytes>768</DownloadedBytes>
      <AvailBandwidth>1024000</AvailBandwidth>
    </Stat>
  </StatisticsGroup>
</PPSPTrackerProtocol>
```

If the request is valid the tracker processes the received information for future use, and generates a response message with a Response value of SUCCESSFUL.

The response MUST have the same TransactionID value as the request.

An example of a Response message for the START_REPORT Request is the following:

```
<?xml version="1.0" encoding="UTF-8"?>
<PPSPTrackerProtocol xmlns="TBD" schemaLocation="TBD" version="1.0">
  <Response>SUCCESSFUL</Response>
  <TransactionID>12345</TransactionID>
</PPSPTrackerProtocol>
```

4.4 Error and Recovery conditions

If the peer fails to read the tracker response, the same Request with identical content, including the same TransactionID, SHOULD be repeated, if the condition is transient.

The TransactionID on a Request can be reused if and only if all of the content is identical, including Date/Time information. Details of the retry process (including time intervals to pause, number of retries to attempt, and timeouts for retrying) are implementation dependent.

The tracker SHOULD be prepared to receive a Request with a repeated TransactionID.

Error situations resulting from the Normal Operation or from abnormal conditions (section 2.4.2) MUST be responded with the adequate response codes, as described here:

If the message is found to be incorrectly formed, the receiver MUST respond with a 400 (Bad Request) response with an empty message-body. The Reason-Phrase SHOULD identify the syntax problem in more detail, for example, "Missing Content-Type header field".

If the version number of the protocol is for a version the receiver does not supports, the receiver MUST respond with a 400 (Bad Request) with an empty message-body. Additional information SHOULD be provided in the Reason-Phrase, for example, "PPSP Version #.#".

If the length of the received message does not matches the Content-Length specified in the message header, or the message is received without a defined Content-Length, the receiver MUST respond with a 411 (Length Required) response with an empty message-body.

If the Request-URI in a Request message is longer than the tracker is willing to interpret, the tracker MUST respond with a 414 (Request-URI Too Long) response with an empty message-body.

In the PEER REGISTERED and TRACKING states of the tracker, certain requests are not allowed (section 2.4.2). The tracker MUST respond with a 403 (Forbidden) response with an empty message-body. The Reason-Phrase SHOULD identify the error condition in more detail, for example, "Action not allowed".

If the tracker is unable to process a Request message due to unexpected condition, it SHOULD respond with a 500 (Internal Server Error) response with an empty message-body.

If the tracker is unable to process a Request message for being in an overloaded state, it SHOULD respond with a 503 (Service Unavailable) response with an empty message-body.

5 Operations and Manageability

This section provides the operational and managements aspects that are required to be considered in implementations of the PPSP Tracker Protocol. These aspects follow the recommendations expressed in RFC 5706 [RFC5706].

5.1 Operational Considerations

The PPSP-TP provides communication between trackers and peers and is conceived as a "client-server" mechanism, allowing the exchange of information about the participant peers sharing multimedia streaming contents.

The "serving" component, i.e., the tracker, is a logical entity that can be envisioned as a centralized service (implemented in one or more physical nodes), or a fully distributed service.

The "client" component can be implemented at each peer participating in the streaming of contents.

5.1.1 Installation and Initial Setup

Content providers wishing to use PPSP for content distribution should setup at least a PPSP Tracker and a service Portal (public web server) to publish links of the content descriptions, for access to their on-demand or live original contents sources. Content/Service providers should also create conditions to generate PEER IDs and any required security certificates, as well as CHUNK IDs and SWARM IDs for each streaming content. The configuration processes for the PPSP Tracking facility, the service Portal and content sources are not standardized, enabling all the flexibility for implementers.

The SWARM IDs of available contents, as well as the addresses of the PPSP Tracking facility, can be distributed to end-users in various ways, but it is common practice to include both the SWARM ID and the corresponding PPSP Tracker addresses (as URLs) in the MPD of the content, which is obtainable (a link) from the service Portal.

End-users browse and search for the desired contents in the service Portal, selecting by clicking the links of the corresponding MPDs. This action typically launches the Client Media Player (with PPSP awareness) which will then, using PPSP-TP, contact the PPSP Tracker to join the corresponding swarm and obtain the transport addresses of other PPSP peers in order to start streaming the content.

5.1.2 Migration Path

Since there is no previous standard protocol providing similar functionality, this specification does not details a migration path.

5.1.3 Requirements on Other Protocols and Functional Components

For security reasons, when using PPSP Peer protocol with PPSP-TP, the mechanisms described in section 6.1 should be observed.

5.1.4 Impact on Network Operation

As the messaging model of PPSP-TP aligns with HTTP protocol and the semantics of its messages, the impact on Network Operation is similar to using HTTP.

5.1.5 Verifying Correct Operation

The correct operation of PPSP-TP can be verified both at the Tracker and at the peer by logging the behavior of PPSP-TP. Additionally, the PPSP Tracker collects the status of the peers including peer's activity, and such information can be used to monitor and obtain the global view of the operation.

5.2 Management Considerations

The management considerations for PPSP-TP are similar to other solutions using HTTP for large-scale content distribution. The PPSP Tracker can be realized by geographically distributed tracker nodes or multiple server nodes in a data center. As these nodes are akin to WWW nodes, their configuration procedures, detection of faults, measurement of performance, usage accounting and security measures can be achieved by standard solutions and facilities.

5.2.1 Interoperability

Interoperability refers to allowing information sharing and operations between multiple devices and multiple management applications. For PPSP-TP, distinct types of devices host PPSP-TP servers (Trackers) and clients (Peers). Therefore, support for multiple standard schema languages, management protocols and information models, suited to different purposes, was considered in the PPSP-TP design. Specifically, management functionalities for PPSP-TP devices can be achieved with Simple Network Management Protocol (SNMP) [RFC3410], syslog [RFC5424] and NETCONF [RFC6241].

5.2.2 Management Information

PPSP Trackers may implement SNMP management interfaces, namely the Application Management MIB [RFC2564] without the need to instrument the Tracker application itself. The channel, connections and transaction objects of the the Application Management MIB can be used to report the basic behavior of the PPSP Tracker service.

The Application Performance Measurement MIB (APM-MIB) [RFC3729] and the Transport Performance Metrics MIB (TPM-MIB) [RFC4150] can be used with PPSP-TP, providing adequate metrics for the analysis of performance for transaction flows in the network, in direct relationship to the transport of PPSP-TP.

The Host Resources MIB [RFC2790] can be used to supply information on the hardware, the operating system, and the installed and running software on a PPSP Tracker host.

The TCP-MIB [RFC4022] can additionally be considered for network monitoring.

Logging is an important functionality for PPSP-TP server (Tracker) and client (Peer), done via syslog [RFC5424].

5.2.3 Fault Management

As PPSP Tracker failures can be mainly attributed to host or network conditions, the facilities previously described for verifying the correct operation of PPSP-TP and the management of PPSP Tracker servers, appear sufficient for PPSP-TP fault monitoring.

5.2.4 Configuration Management

PPSP Tracker deployments, when realized by geographically distributed tracker nodes or multiple server nodes in a data center, may benefit from a standard way of replicating atomic configuration updates over a set of server nodes. This functionality can be provided via NETCONF [RFC6241].

5.2.5 Accounting Management

PPSP-TP implementations, namely for content provider environments, can benefit from accounting standardization efforts as defined in [RFC2975], in terms of resource consumption data, for the purposes of capacity and trend analysis, cost allocation, auditing, and billing.

5.2.6 Performance Management

Being transaction-oriented, PPSP-TP performance, in terms of availability and responsiveness, can be measured with the facilities of the APM-MIB [RFC3729] and the TPM-MIB [RFC4150].

5.2.7 Security Management

Standard SNMP notifications for PPSP Tracker management and syslog messages [RFC5424] can be used, to alert operators to the conditions identified in the security considerations (section 6).

The statistics collected about the operation of PPSP-TP can be used for detecting attacks, such as the receipt of malformed messages, messages out of order, or messages with invalid timestamps.

6 Security Considerations

P2P streaming systems are subject to attacks by malicious/unfriendly peers/trackers that may eavesdrop on signaling, forge/deny information/knowledge about streaming content and/or its availability, impersonating to be another valid participant, or launch DoS attacks to a chosen victim.

No security system can guarantee complete security in an open P2P streaming system where participants may be malicious or uncooperative. The goal of security considerations described here is to provide sufficient protection for maintaining some security properties during the tracker-peer communication even in the face of a large number of malicious peers and/or eventual distrustful trackers (under the distributed tracker deployment scenario).

Since the protocol uses HTTP to transfer signaling most of the same security considerations described in RFC 2616 also apply [RFC2616].

6.1 Authentication between Tracker and Peers

To protect the PPSP-TP signaling from attackers pretending to be valid peers (or peers other than themselves) all messages received in the tracker SHOULD be received from authorized peers. For that purpose a peer SHOULD enroll in the system via a centralized enrollment server. The enrollment server is expected to provide a proper Peer ID for the peer and information about the authentication mechanisms. The specification of the enrollment method and the provision of identifiers and authentication tokens is out of scope of this specification.

A channel-oriented security mechanism should be used in the communication between peers and tracker, such as the Transport Layer Security (TLS) to provide privacy and data integrity.

Due to the transactional nature of the communication between peers and tracker the method for adding authentication and data security services can be the OAuth 2.0 Authorization [RFC6749] with bearer token, which provides the peer with the information required to successfully utilize an access token to make protected requests to the tracker [RFC6750].

6.2 Content Integrity protection against polluting peers/trackers

Malicious peers may disclaim ownership of popular content to the tracker but try to serve polluted (i.e., decoy content or even virus/trojan infected contents) to other peers.

This kind of pollution can be detected by incorporating integrity verification schemes for published shared contents. As content chunks are transferred independently and concurrently, a correspondent chunk-level integrity verification **MUST** be used, checked with signed fingerprints received from authentic origin.

6.3 Residual attacks and mitigation

To mitigate the impact of Sybil attackers, impersonating a large number of valid participants by repeatedly acquiring different peer identities, the enrollment server **SHOULD** carefully regulate the rate of peer/tracker admission.

There is no guarantee that peers honestly report their status to the tracker, or serve authentic content to other peers as they claim to the tracker. It is expected that a global trust mechanism, where the credit of each peer is accumulated from evaluations for previous transactions, may be taken into account by other peers when selecting partners for future transactions, helping to mitigate the impact of such malicious behaviors. A globally trusted tracker **MAY** also take part of the trust mechanism by collecting evaluations, computing credit values and providing them to joining peers.

6.4 Pro-incentive parameter trustfulness

Property types for STAT_REPORT messages (such as those of property="StreamStatistics" in Table 5 of section 3.2) may consider additional pro-incentive parameters (guidelines for extension in section 7), which can enable the tracker to improve the performance of the whole P2P streaming system. Trustworthiness of these pro-incentive parameters is critical to the effectiveness of the incentive mechanisms. Furthermore, both the amount of uploaded and downloaded data should be reported to the tracker to allow checking if there is any inconsistency between the upload and download report, and establish an appropriate credit/trust system.

One such solution could be a reputation-incentive mechanism, based on the notions of reputation, social awareness and fairness. The mechanism would promote cooperation among participants (via each peer's reputation) based on the history of past transactions, such as, count of chunk requests (sent, received) in a swarm, contribution time of the peer, cumulative uploaded and downloaded content, JOIN and LEAVE timestamps, attainable rate, etc.

Alternatively, exchange of cryptographic receipts signed by receiving peers can be used to attest to the upload contribution of a peer to the swarm, as suggested in [Contracts].

7 Guidelines for Extending PPSP-TP

Extension mechanisms allow designers to add new features or to customize existing features of a protocol for different operating environments [RFC6709].

Since the beginning that extensibility has been a design goal of PPSP-TP, as the protocol is represented in the Extensible Markup Language [XML].

A powerful feature of XML is the extensibility, but this feature also provides multiple opportunities to make poor design decisions.

Extending a protocol implies either the addition of features without changing the protocol itself or the addition of new elements creating new versions of an existing schema and therefore new versions of the protocol.

In PPSP-TP it means that an extension **MUST NOT** alter an existing protocol schema as the changes would result in a new version of an existing schema, not an extension of an existing schema, typically non-backwards-compatible.

Additionally, a designer **MUST** remember that extensions themselves **MAY** also be extensible.

Extensions **MUST** adhere to the principles described in this section in order to be considered valid.

Extensions **MAY** be documented as Internet-Draft and RFC documents if there are requirements for coordination, interoperability, and broad distribution.

Extensions need not be published as Internet-Draft or RFC documents if they are intended for operation in a closed environment or are otherwise intended for a limited audience.

7.1 Forms of PPSP-TP Extension

PPSP-TP extensions **MUST** be specified in XML. This ensures that parsers capable of processing PPSP-TP structures will also be capable of processing PPSP-TP extensions. Guidelines for the use of XML in IETF protocols can be found in RFC 3470 [RFC3470].

In PPSP-TP two extension mechanisms can be used: a Request-Response Extension or a Protocol-level Extension.

- o Request-Response Extension: Adding elements or attributes to an

existing element mapping in the schema is the simplest form of extension. This form should be explored before any other. This task can be accomplished by extending an existing element mapping.

For example, an element mapping for the StatisticsGroup (Section 3.2, Table 5) can be extended to include additional elements needed to express status information about the activity of the peer, such as OnlineTime for the Stat element:

```
<Stat property="StreamStatistics">  
  <OnlineTime>3600</OnlineTime>  
</Stat>
```

Another example also for the StatisticsGroup is for additional @property attributes and elements, such as those necessary to report segment chunks availability (maps of Chunk IDs):

```
<Stat property="ContentMap">  
  <SegmentInfo chunkmapSize="25">  
    A/8D/wP/A/8D/wP/A/8D/wP/A/8D/wP/....  
  </SegmentInfo>  
</Stat>
```

- o Protocol-level Extension: If there is no existing element mapping that can be extended to meet the requirements and the existing PPSP-TP Request and Response message structures are insufficient, then extending the protocol should be considered in order to define new operational Requests and Responses.

For example, to enhance the level of control and the granularity of the operations, a new version of the protocol with new messages (JOIN, DISCONNECT), a retro-compatible change in semantics of an existing CONNECT Request/Response and an extension in STAT_REPORT could be considered.

As illustrated in Figure 6, the peer would use an enhanced CONNECT Request to perform the initial registration in the system. Then it would JOIN a first swarm as Seeder, later JOIN a second swarm as Leecher, and then DISCONNECT from the latter swarm but keeping as Seeder for the first one. When deciding to leave the system, the peer DISCONNECTs gracefully from it:

```

+-----+
|  Peer  |
+-----+
|
|  --CONNECT----->
|  <-----OK--
|  --JOIN(swarm_a;SEED)----->
|  <-----OK--
|  :
|  --STAT_REPORT(activity)----->
|  <-----Ok--
|  :
|  --JOIN(swarm_b;LEECH)----->
|  <-----OK+PeerList--
|  :
|  --STAT_REPORT(ChunkMap_b)----->
|  <-----Ok--
|  :
|  --DISCONNECT(swarm_b)----->
|  <-----Ok--
|  :
|  --STAT_REPORT(activity)----->
|  <-----Ok--
|  :
|  --DISCONNECT----->
|  <-----Ok(BYE)--

```

Figure 6: Example of a session for a PPSP-TP extended version.

7.2 Issues to Be Addressed in PPSP-TP Extensions

There are several issues that all extensions should take into consideration.

- Overview of the Extension: It is RECOMMENDED that extensions to PPSP-TP have a protocol overview section that discusses the basic operation of the extension. The most important processing rules for the elements in the message flows SHOULD also be mentioned.
- Backward Compatibility: One of the most important issues to consider is whether the new extension is backward compatible with the base PPST-TP.
- Syntactic Issues: Extensions that define new Request/Response methods SHOULD use all capitals for the method name, keeping with a long-standing convention in many protocols, such as HTTP. Method names are case sensitive in PPSP-TP. Method names SHOULD be shorter than 16 characters and SHOULD attempt to convey the

general meaning of the Request or Response.

- Semantic Issues: PPSP-TP extensions MUST clearly define the semantics of the extensions. Specifically, the extension MUST specify the behaviors expected from both the Peer and the Tracker in processing the extension, with the processing rules in temporal order of the common messaging scenario.

Processing rules generally specify actions to be taken on receipt of messages and expiration of timers.

The extension SHOULD specify procedures to be taken in exceptional conditions that are recoverable. Handling of unrecoverable errors does not require specification.

- Security Issues: Being security an important component of any protocol, designers of PPSP-TP extensions need to carefully consider security requirements, namely authorization requirements and requirements for end-to-end integrity.
- Examples of Usage: The specification of the extension SHOULD give examples of message flows and message formatting and include examples of messages containing new syntax. Examples of message flows should be given to cover common cases and at least one failure or unusual case.

8 IANA Considerations

There are presently no IANA considerations with this document.

9 Acknowledgments

The authors would like to thank many people for for their help and comments, particularly: Zhang Yunfei, Liao Hongluan, Roni Even, Bhumip Khasnabish, Wu Yichuan, Peng Jin, Chi Jing, Zong Ning, Song Haibin, Chen Wei, Zhijia Chen, Christian Schmidt, Lars Eggert, David Harrington, Henning Schulzrinne, Kangheng Wu, Martin Stiemmerling, Jianyin Zhang, Johan Pouwelse and Arno Bakker.

Rui Cruz, Mario Nunes and Joao Taveira are partially supported by the SARACEN project [SARACEN], a research project of the European Union 7th Framework Programme (contract no. ICT-248474).

The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the SARACEN project, the European Commission, Huawei or China Mobile.

This I-D document was prepared using NroffEdit version 2.08 (<http://aaa-sec.com/nroffedit/>).

10 References

10.1 Normative References

- [RFC0793] Postel, J., "Transmission Control Protocol", STD 7, RFC 793, September 1981.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2616] Fielding, R., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., and T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1", RFC 2616, June 1999.
- [RFC3629] Yergeau, F., "UTF-8, a transformation format of ISO 10646", STD 63, RFC 3629, November 2003.
- [RFC5322] Resnick, P., Ed., "Internet Message Format", RFC 5322, October 2008.
- [XML] Bray, T., Paoli, J., Sperberg-McQueen, C., Maler, E. and F. Yergeau, "Extensible Markup Language (XML) 1.0 (Fifth Edition)", W3C xml, November 2008, <<http://www.w3.org/TR/xml/>>.
- [XMLSchema.1] Thompson, H., Beech, D., Maloney, M. and N. Mendelsohn, "XML Schema Part 1: Structures Second Edition", W3C xmlschema-1, October 2004, <<http://www.w3.org/TR/xmlschema-1/>>.
- [XMLSchema.2] Biron, P. and A. Malhotra, "XML Schema Part 2: Datatypes Second Edition", W3C xmlschema-2, October 2004, <<http://www.w3.org/TR/xmlschema-2/>>.
- [XMLNamespace] Bray, T., Hollander, D., Layman, A., Tobin, R. and H. Thompson, "Namespaces in XML", W3C REC-xml-names, December 2009, <<http://www.w3.org/TR/REC-xml-names/>>.
- [EXI] Schneider, J. and T. Kamiya, "Efficient XML Interchange (EXI) Format 1.0", W3C exi, March 2011, <<http://www.w3.org/TR/exi/>>.

10.2 Informative References

- [RFC1952] Deutsch, P., "GZIP file format specification version 4.3", RFC 1952, May 1996.
- [RFC2564] Kalbfleisch, C., Krupczak, C., Presuhn, R., and J.

- Saperia, "Application Management MIB", RFC 2564, May 1999.
- [RFC2790] Waldbusser, S. and P. Grillo, "Host Resources MIB", RFC 2790, March 2000.
- [RFC2818] Rescorla, E., "HTTP Over TLS", RFC 2818, May 2000.
- [RFC2975] Aboba, B., Arkko, J., and D. Harrington, "Introduction to Accounting Management", RFC 2975, October 2000.
- [RFC3410] Case, J., Mundy, R., Partain, D., and B. Stewart, "Introduction and Applicability Statements for Internet-Standard Management Framework", RFC 3410, December 2002.
- [RFC3470] Hollenbeck, S., Rose, M., and L. Masinter, "Guidelines for the Use of Extensible Markup Language (XML) within IETF Protocols", BCP 70, RFC 3470, January 2003.
- [RFC3729] Waldbusser, S., "Application Performance Measurement MIB", RFC 3729, March 2004.
- [RFC4022] Raghunarayan, R., Ed., "Management Information Base for the Transmission Control Protocol (TCP)", RFC 4022, March 2005.
- [RFC4122] Leach, P., Mealling, M., and R. Salz, "A Universally Unique Identifier (UUID) URN Namespace", RFC 4122, July 2005.
- [RFC4150] Dietz, R. and R. Cole, "Transport Performance Metrics MIB", RFC 4150, August 2005.
- [RFC5245] Rosenberg, J., "Interactive Connectivity Establishment (ICE): A Protocol for Network Address Translator (NAT) Traversal for Offer/Answer Protocols", RFC 5245, April 2010.
- [RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, August 2008.
- [RFC5424] Gerhards, R., "The Syslog Protocol", RFC 5424, March 2009.
- [RFC5706] Harrington, D., "Guidelines for Considering Operations and Management of New Protocols and Protocol Extensions", RFC 5706, November 2009.
- [RFC6241] Enns, R., Ed., Bjorklund, M., Ed., Schoenwaelder, J., Ed., and A. Bierman, Ed., "Network Configuration Protocol

- (NETCONF)", RFC 6241, June 2011.
- [RFC6709] Carpenter, B., Aboba, B., Ed., and S. Cheshire, "Design Considerations for Protocol Extensions", RFC 6709, September 2012.
- [RFC6749] Hardt, D., Ed., "The OAuth 2.0 Authorization Framework", RFC 6749, October 2012.
- [RFC6750] Jones, M. and D. Hardt, "The OAuth 2.0 Authorization Framework: Bearer Token Usage", RFC 6750, October 2012.
- [RFC6972] Zhang, Y. and N. Zong, "Problem Statement and Requirements of the Peer-to-Peer Streaming Protocol (PPSP)", RFC 6972, July 2013.
- [I-D.ietf-httpbis-http2] Belshe, M., Twist, Peon, R., Thomson, M., and A. Melnikov, "Hypertext Transfer Protocol version 2.0", draft-ietf-httpbis-http2-06 (work in progress), August 2013.
- [ISO.IEC.23009-1] ISO/IEC, "Information technology -- Dynamic adaptive streaming over HTTP (DASH) -- Part 1: Media presentation description and segment formats", ISO/IEC DIS 23009-1, Aug. 2011.
- [I.D.ietf-alto-protocol] Alimi, R., Penno, R. and Y. Yang, "ALTO Protocol", draft-ietf-alto-protocol-20, (work in progress), October 2013.
- [I-D.pantos-http-live-streaming] Pantos, R. and W. May, "HTTP Live Streaming", draft-pantos-http-live-streaming-12 (work in progress), October 2013.
- [SARACEN] "SARACEN Project Website", <http://www.saracen-p2p.eu/>.
- [Contracts] Piatek, M., Venkataramani, A., Yang, R., Zhang, D. and A. Jaffe, "Contracts: Practical Contribution Incentives for P2P Live Streaming", in NSDI '10: USENIX Symposium on Networked Systems Design and Implementation, April 2010.

Appendix A. Revision History

- 00 2013-02-14 Initial version.
- 01 2013-02-14 Minor revision.
 - * Terminology definitions related with streaming media contents moved to Appendix C. - Use Scenarios and Assumptions.
 - * Terminology updated to follow reviewed PPSP Problem Statement [RFC6972] definitions.
 - Definition of JOIN TIME, PEER-PEER MESSAGES and TRACKER-PEER MESSAGES removed.
 - + Definitions of PEER LIST, CHUNK ID, PPSP-TP, SEGMENT CHUNK, TRANSACTION ID added.
 - + Definitions of SEGMENT updated.
 - + Descriptions of Client Media Player and service Portal, related to PPSP protocols, were added to section 2.1.
 - + In section 2.2, "not exhaustive" and "typical" was added to first paragraph.
 - + In section 2.3.2, the description of CONNECT was updated to clarify the information recorded and maintained by the tracker related with Peer addressing, location and activity mode in the swarms.
 - + In section 3.1 and Table 2, PeerNum element justified to be limited to 30 entries.
 - + Section 6.4 enhanced with complementary information on incentive mechanisms.
 - + Section 5 completed.
- 02 2013-10-21 Minor revision.
 - + In section 2.2, Figure 1 was complemented to include the FIND request.
 - In section 2.3.1, references to HTTP/2.0 and framing formats were removed.
 - + In section 2.3.2, Table 1 was updated with FIND request message and the corresponding description added in text.
 - + In section 2.4, Figure 5 and the text were updated to reflect the FIND request message.
 - + In sections 2.4.1 and 2.4.2, the PEER REGISTERED and TRACKING state descriptions were modified to include the FIND request message.
 - In sections 3.1, references to EXI were removed and text was modified to reflect the FIND request message.
 - + In sections 3.3, Table 1 was updated with FIND request message.
 - + In sections 3.4, the text was updated to include the response to FIND requests.
 - + Section 4.2 renamed to 4.3, and new section 4.2 describes in detail the processing of FIND requests. All references to previous section 4.2 updated to 4.3.
 - + Appendix D added.

Appendix B. PPSP-TP Message Syntax for HTTP/1.1

PPSP-TP messages use the generic message format of RFC 5322 [RFC5322] for transferring the payload of the message (Requests and Responses).

PPSP-TP messages consist of a start-line, one or more header fields, an empty line indicating the end of the header fields, and, when applicable, a message-body.

The start-line, each message-header line, and the empty line MUST be terminated by a carriage-return line-feed sequence (CRLF). Note that the empty line MUST be present even if the message-body is not.

The PPSP-TP message and header field syntax is identical to HTTP/1.1 [RFC2616].

A Request message is a standard HTTP/1.1 message starting with a Request-Line generated by the HTTP client peer. The Request-Line contains a method name, a Request-URI, and the protocol version separated by a single space (SP) character.

Request-Line =
 Method SP Request-URI SP HTTP-Version CRLF

A Request message example is the following:

```
<Method> /<Resource> HTTP/1.1
Host: <Host>
Content-Lenght: <ContentLenght>
Content-Type: <ContentType>
Authorization: <AuthToken>

[Request_Body]
```

The HTTP Method token and Request-URI (the Resource) identifies the resource upon which to apply the operation requested.

The Response message is also a standard HTTP/1.1 message starting with a Status-Line generated by the tracker. The Status-Line consists of the protocol version followed by a numeric Status-Code and its associated Reason-Phrase, with each element separated by a single SP character.

Status-Line =
 HTTP-Version SP Status-Code SP Reason-Phrase CRLF

A Response message example is the following:

```
HTTP/1.1 <Status-Code> <Reason-Phrase>  
Content-Lenght: <ContentLenght>  
Content-Type: <ContentType>  
Content-Encoding: <ContentCoding>
```

```
[Response_Body]
```

The Status-Code element is a 3-digit integer result code that indicates the outcome of an attempt to understand and satisfy a request.

The Reason-Phrase element is intended to give a short textual description of the Status-Code.

B.1 Header Fields

The header fields are identical to HTTP/1.1 header fields in both syntax and semantics.

Some header fields only make sense in requests or responses. If a header field appears in a message not matching its category (such as a request header field in a response), it MUST be ignored.

The Host request-header field in the request message follows the standard rules for the HTTP/1.1 Host header.

The Content-Type entity-header field MUST be used in requests and responses containing message-bodies to define the Internet media type of the message-body.

The Content-Encoding entity-header field MAY be used in response messages with "gzip" compression scheme [RFC1952] for faster transmission times and less network bandwidth usage.

The Content-Length entity-header field MUST be used in messages containing message-bodies to locate the end of each message in a stream.

The Authorization header field in the request message allows a peer to authenticate itself with a tracker, containing authentication information.

B.2 Methods

PPSP-TP uses HTTP/1.1 POST method token for all request messages.

B.3 Message Bodies

PPSP-TP requests MUST contain message-bodies.

PPSP-TP responses MAY include a message-body.

If the message-body has undergone any encoding such as compression, then this MUST be indicated by the Content-Encoding header field; otherwise, Content-Encoding MUST be omitted.

The character set of the message body is indicated as part of the Content-Type header-field, and the default value for PPSP-TP messages is "UTF-8".

B.4 Message Response Codes

The response codes in PPSP-TP response messages are consistent with HTTP/1.1 response status-codes. However, not all HTTP/1.1 response status-codes are appropriate for PPSP-TP, and only those that are appropriate are given here. Other HTTP/1.1 response codes SHOULD NOT be used in PPSP-TP.

The class of the response is defined by the first digit of the Status-Code. The last two digits do not have any categorization role. For this reason, any response with a Status-Code between 200 and 299 is referred to as a "2xx response", and similarly to the other supported classes:

2xx: Success -- the action was successfully received, understood, and accepted;

4xx: Peer Error -- the request contains bad syntax or cannot be fulfilled at this tracker;

5xx: Tracker Error -- the tracker failed to fulfill an apparently valid request;

The valid response codes are the following (Status-Code Reason-Phrase):

200 OK -- The request has succeeded. The information returned with the response describes or contains the result of the action;

400 Bad Request -- The request could not be understood due to malformed syntax.

401 Unauthorized -- The request requires authentication.

403 Forbidden -- The tracker understood the request, but is refusing to fulfill it. The request SHOULD NOT be repeated.

- 404 Not Found -- This status is returned if the tracker did not find anything matching the Request-URI.
- 408 Request Timeout -- The peer did not produce a request within the time that the tracker was prepared to wait.
- 411 Length Required -- The tracker refuses to accept the request without a defined Content-Length. The peer MAY repeat the request if it adds a valid Content-Length header field containing the length of the message-body in the request message.
- 414 Request-URI Too Long -- The tracker is refusing to service the request because the Request-URI is longer than the tracker is willing to interpret. This rare condition is likely to occur when the tracker is under attack by a client attempting to exploit security holes.
- 500 Internal Server Error -- The tracker encountered an unexpected condition which prevented it from fulfilling the request.
- 503 Service Unavailable -- The tracker is currently unable to handle the request due to a temporary overloading or maintenance condition.

Appendix C. Use Scenarios

This section is tutorial in nature and does not contain any normative statements.

This section describes some aspects of the use of PPSP-TP. The examples were chosen to illustrate the basic operation, but not to limit what PPSP-TP may be used for.

C.1 Additional Terminology

ADAPTIVE STREAMING: Multiple alternate representations (different qualities and bitrates) of the same media content co-exist for the same streaming session; each alternate representation corresponds to a different media quality level; peers can choose among the alternate representations for decode and playback.

BASE LAYER: The playable representation level in Scalable Video Coding (SVC) required by all upper level Enhancements Layers for proper decoding of the video.

COMPLEMENTARY REPRESENTATION: Representation in a set of content representations which have inter-representation dependencies and which, when combined, result in a single representation for decoding and presentation.

CONTINUOUS MEDIA: Media with an inherent notion of time, for example, speech, audio, video, timed text or timed metadata.

ENHANCEMENT LAYER: Enhancement differential quality level (complementary representation) in Scalable Video Coding (SVC) used to produce a higher quality, higher definition video in terms of space (i.e., image resolution), time (i.e., frame rate) or Signal-to-Noise Ratio (i.e., fidelity) when combined with the playable Base Layer.

MEDIA COMPONENT: An encoded version of one individual media type such as audio, video or timed text with specific attributes, e.g., bandwidth, language, or resolution.

SCALABLE STREAMING: With Multiple Description Coding (MDC), multiple additive descriptions (that can be independently played-out) to refine the quality of the video when combined together. With Scalable Video Coding (SVC), nested dependent enhancement layers (hierarchical levels of quality), refine the quality of lower layers, from the lowest level (the playable Base Layer). With Multiple View Coding (MVC), multiple views nested in dependent enhancement layers (levels of quality) allow the video to be played in compatible 3D rendering devices when the views are combined together.

C.2 Functional Entities

The functional entities related to PPSP protocols are the Client Media Player, the service Portal, the tracker and the peers.

The Client Media Player is a logical entity providing direct interface to the end user at the client device, and includes the functions to select, request, decode and render contents. The Client Media Player may interface with the local peer application using request and response standard formats for HTTP Request and Response messages [RFC2616].

The service Portal is a logical entity typically used for client enrollment and content information publishing, searching and retrieval.

The tracker is a logical entity that maintains the lists of PPSP active peers storing and exchanging a specific media content. The tracker also stores the status of active peers in swarms, to help in the selection of appropriate peers for a requesting peer. The tracker can be realized by geographically distributed tracker nodes or multiple server nodes in a data center, increasing the content availability, the service robustness and the network scalability or reliability. The management and locating of content index information are totally internal behaviors of the tracker system, which is invisible to the PPSP Peer.

The peer is also a logical entity in the client device embedding the P2P core engine, with a client serving side interface to respond to Client Media Player requests and a network side interface to exchange data and PPSP signaling with trackers and other peers.

The streaming technique is chunk-based, i.e., client peers obtain media chunks from serving peers and handle the buffering that is necessary for the playback processes during the download of the media chunks.

In Live streaming, all end users are interested in a specific media coming from an ongoing program, which means that all respective peers share nearly the same streaming content at a given point of time. Peers may store the live media for further distribution (known as time-shift TV), where the stored media is distributed in a VoD-like manner.

In VoD, different end users watch different parts of the recorded media content during a past event. In this case, each respective peer obtains from other peers the information on media chunks they store and then get the required media from a selected set of those

peers. While watching VoD, an end user can also switch to any place of the content, e.g., skip the credits part, or skip the part that it is not interested in. In this case the respective participating peer may not store all the content segments. From the whole swarm point of view, the participating peers typically store different parts of content.

C.3 Streaming Capabilities

This section is tutorial in nature and does not contain any normative statements.

The process used for media streaming distribution assumes a segment (chunk) transfer scheme whereby the original content (that can be encoded using adaptive or scalable techniques) is chopped into small segments having the following representations:

1. Adaptive - alternate representations with different qualities and bitrates; a single representation is non-adaptive;
2. Scalable description levels - multiple additive descriptions (i.e., addition of descriptions refine the quality of the video);
3. Scalable layered levels - nested dependent layers corresponding to several hierarchical levels of quality, i.e., higher enhancement layers refine the quality of the video of lower layers.
4. Scalable multiple views - views correspond to mono (2D) and stereoscopic (3D) videos, with several hierarchical levels of quality.

These streaming distribution techniques support dynamic variations in video streaming quality while ensuring support for a plethora of end user devices and network connections.

C.4 NAT Traversal

It is assumed that all trackers must be in the public Internet and have been placed there deliberately. This document will not describe NAT Traversal mechanisms but the protocol facilitates flexible NAT Traversal techniques, such as those based on ICE [RFC5245], considering that the tracker node may provide NAT traversal services, as a STUN-like tracker. Being a STUN-like tracker, it can discover the reflexive candidate addresses of a peer and make them available in responses to other requesting peers.

C.5 Content Information Metadata

Multimedia contents may consist of several media components (for example, audio, video, and timed text), each of which might have different characteristics.

The representations of a media content correspond to encoded alternatives of the same media component, varying from other representations by bitrate, resolution, number of channels, or other characteristics. Each representation consists of one or multiple segments. Segments are the media stream transport chunks in temporal sequence.

These characteristics may be described in a Media Presentation Description (MPD) file. It is envisioned that the content information metadata used in PPSP may align with MPD formats, such as ISO/IEC 23009-1 [ISO.IEC.23009-1] and [I-D.pantos-http-live-streaming].

C.6 Authentication, Confidentiality, Integrity

Channel-oriented security can be used in the communication between peers and tracker, such as the Transport Layer Security (TLS) to provide privacy and data integrity. HTTP/1.1 over TLS (HTTPS) [RFC2818] is the preferred approach for preventing disclosure of peer critical information via the communication channel.

Due to the transactional nature of the communication between peers and tracker a method for adding authentication and data security services via replaceable mechanisms may be employed. One such method is the OAuth 2.0 Authorization [RFC6749] with bearer token, providing the peer with the information required to successfully utilize the access token to make protected requests to the tracker [RFC6750].

Appendix D. Implementation Options

With newer versions of HTTP, the framing formats used for encoding and transmitting the data over the wire, e.g., the encapsulation proposed for HTTP/2.0 [I-D.ietf-httpbis-http2] may be used in PPSP-TP.

Similarly, for compact on the wire representation, PPSP-TP Requests and Responses can be represented in a binary form using, for example, Efficient XML Interchange (EXI) Format [EXI].

When using the above binary format for both the transport protocol and the PPSP messages the returned Peer List from a CONNECT or a FIND Request should be restricted to a number of entries that makes the total message size smaller than 1 KiB, allowing it to fit inside a single IP packet. The RECOMMENDED approach is to limit the list to 30 entries (less than 1KiB in size), as this size has a high likelihood of traveling across the Internet without fragmentation, or be transmitted as a single IP packet over an Ethernet network (1500 byte frame size).

Authors' Addresses

Rui Santos Cruz
IST/INESC-ID/INOV
Phone: +351.939060939
Email: rui.cruz@ieee.org

Gu Yingjie
Email: guyingjie@gmail.com

Mario Serafim Nunes
IST/INESC-ID/INOV
Rua Alves Redol, n.9
1000-029 LISBOA, Portugal
Phone: +351.213100256
Email: mario.nunes@inov.pt

Jinwei Xia
Huawei
Nanjing, Baixia District 210001, China
Phone: +86-025-86622310
Email: xiajinwei@huawei.com

Joao P. Taveira
IST/INOV
Email: joao.silva@inov.pt

Deng Lingli
China Mobile

PPSP
Internet-Draft
Intended status: Standards Track
Expires: April 24, 2014

A. Bakker
Vrije Universiteit Amsterdam
R. Petrocco
V. Grishchenko
Technische Universiteit Delft
October 21, 2013

Peer-to-Peer Streaming Peer Protocol (PPSPP)
draft-ietf-ppsp-peer-protocol-08

Abstract

The Peer-to-Peer Streaming Peer Protocol (PPSPP) is a protocol for disseminating the same content to a group of interested parties in a streaming fashion. PPSPP supports streaming of both pre-recorded (on-demand) and live audio/video content. It is based on the peer-to-peer paradigm, where clients consuming the content are put on equal footing with the servers initially providing the content, to create a system where everyone can potentially provide upload bandwidth. It has been designed to provide short time-till-playback for the end user, and to prevent disruption of the streams by malicious peers. PPSPP has also been designed to be flexible and extensible. It can use different mechanisms to optimize peer uploading, prevent freeriding, and work with different peer discovery schemes (centralized trackers or Distributed Hash Tables). It supports multiple methods for content integrity protection and chunk addressing. Designed as a generic protocol that can run on top of various transport protocols, it currently runs on top of UDP using LEDBAT for congestion control.

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 April 24, 2014.

Copyright Notice

Copyright (c) 2013 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. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	6
1.1. Purpose	6
1.2. Requirements Language	7
1.3. Terminology	7
2. Overall Operation	9
2.1. Example: Joining a Swarm	9
2.2. Example: Exchanging Chunks	10
2.3. Example: Leaving a Swarm	11
3. Messages	11
3.1. HANDSHAKE	11
3.2. HAVE	12
3.3. DATA	12
3.4. ACK	12
3.5. INTEGRITY	13
3.6. SIGNED_INTEGRITY	13
3.7. REQUEST	13
3.8. CANCEL	14
3.9. CHOKE and UNCHOKE	14
3.10. Peer Address Exchange	14
3.10.1. PEX_REQ and PEX_RES Messages	14
3.11. Channels	16
3.12. Keep Alive Signalling	17
4. Chunk Addressing Schemes	17
4.1. Start-End Ranges	18
4.1.1. Chunk Ranges	18
4.1.2. Byte Ranges	18
4.2. Bin Numbers	18
4.3. In Messages	20
4.3.1. In HAVE Messages	20
4.3.2. In ACK Messages	20

5.	Content Integrity Protection	21
5.1.	Merkle Hash Tree Scheme	21
5.2.	Content Integrity Verification	22
5.3.	The Atomic Datagram Principle	23
5.4.	INTEGRITY Messages	24
5.5.	Discussion and Overhead	24
5.6.	Automatic Detection of Content Size	25
5.6.1.	Peak Hashes	26
5.6.2.	Procedure	28
6.	Live Streaming	28
6.1.	Content Authentication	28
6.1.1.	Sign All	29
6.1.2.	Unified Merkle Tree	30
6.1.2.1.	Signed Munro Hashes	30
6.1.2.2.	Munro Signature Calculation	33
6.1.2.3.	Procedure	33
6.1.2.4.	Secure Tune In	33
6.2.	Forgetting Chunks	34
7.	Protocol Options	35
7.1.	End Option	35
7.2.	Version	35
7.3.	Minimum Version	36
7.4.	Swarm Identifier	36
7.5.	Content Integrity Protection Method	37
7.6.	Merkle Tree Hash Function	37
7.7.	Live Signature Algorithm	38
7.8.	Chunk Addressing Method	38
7.9.	Live Discard Window	39
7.10.	Supported Messages	40
8.	UDP Encapsulation	40
8.1.	Chunk Size	40
8.2.	Datagrams and Messages	41
8.3.	Channels	42
8.4.	HANDSHAKE	43
8.5.	HAVE	44
8.6.	DATA	45
8.7.	ACK	45
8.8.	INTEGRITY	46
8.9.	SIGNED_INTEGRITY	47
8.10.	REQUEST	48
8.11.	CANCEL	48
8.12.	CHOKe and UNCHOKe	49
8.13.	PEX_REQ, PEX_RESv4, PEX_RESv6 and PEX_REScert	49
8.14.	KEEPALIVE	51
8.15.	Detecting a Dead Peer	51
8.16.	Flow and Congestion Control	52
8.17.	Example of Operation	52
9.	Extensibility	57

9.1. Chunk Picking Algorithms	57
9.2. Reciprocity Algorithms	57
10. Acknowledgements	57
11. IANA Considerations	58
11.1. PPSP Peer Protocol Message Type Registry	58
11.2. PPSP Peer Protocol Option Registry	58
11.3. PPSP Peer Protocol Version Number Registry	58
11.4. PPSP Peer Protocol Content Integrity Protection Method Registry	58
11.5. PPSP Peer Protocol Merkle Hash Tree Function Registry . .	58
11.6. PPSP Peer Protocol Chunk Addressing Method Registry . . .	58
12. Manageability Considerations	58
12.1. Operations	59
12.1.1. Installation and Initial Setup	59
12.1.2. Requirements on Other Protocols and Functional Components	60
12.1.3. Migration Path	60
12.1.4. Impact on Network Operation	60
12.1.5. Verifying Correct Operation	60
12.1.6. Configuration	61
12.2. Management Considerations	61
12.2.1. Management Interoperability and Information	62
12.2.2. Fault Management	62
12.2.3. Configuration Management	62
12.2.4. Accounting Management	63
12.2.5. Performance Management	63
12.2.6. Security Management	63
13. Security Considerations	63
13.1. Security of the Handshake Procedure	63
13.1.1. Protection Against Attack 1	64
13.1.2. Protection Against Attack 2	65
13.1.3. Protection Against Attack 3	65
13.2. Secure Peer Address Exchange	66
13.2.1. Protection against the Amplification Attack	66
13.2.2. Example: Tracker as Certification Authority	67
13.2.3. Protection Against Eclipse Attacks	68
13.3. Support for Closed Swarms (PPSP.SEC.REQ-1)	68
13.4. Confidentiality of Streamed Content (PPSP.SEC.REQ-2+3) . .	68
13.5. Strength of the Hash Function for Merkle Hash Trees . . .	69
13.6. Limit Potential Damage and Resource Exhaustion by Bad or Broken Peers (PPSP.SEC.REQ-4+6)	69
13.6.1. HANDSHAKE	69
13.6.2. HAVE	69
13.6.3. DATA	70
13.6.4. ACK	70
13.6.5. INTEGRITY and SIGNED_INTEGRITY	70
13.6.6. REQUEST	71
13.6.7. CANCEL	71

13.6.8. CHOKe	71
13.6.9. UNCHOKe	71
13.6.10. PEX_RES	72
13.6.11. Unsolicited Messages in General	72
13.7. Exclude Bad or Broken Peers (PPSP.SEC.REQ-5)	72
14. References	72
14.1. Normative References	72
14.2. Informative References	73
Appendix A. Revision History	77
Authors' Addresses	95

1. Introduction

1.1. Purpose

This document describes the Peer-to-Peer Streaming Peer Protocol (PPSPP), designed for disseminating the same content to a group of interested parties in a streaming fashion. PPSPP supports streaming of both pre-recorded (on-demand) and live audio/video content. It is based on the peer-to-peer paradigm where clients consuming the content are put on equal footing with the servers initially providing the content, to create a system where everyone can potentially provide upload bandwidth.

PPSPP has been designed to provide short time-till-playback for the end user, and to prevent disruption of the streams by malicious peers. Central in this design is a simple method of identifying content based on self-certification. In particular, content in PPSPP is identified by a single cryptographic hash that is the root hash in a Merkle hash tree calculated recursively from the content [MERKLE][ABMRKL]. This self-certifying hash tree allows every peer to directly detect when a malicious peer tries to distribute fake content. The tree can be used for both static and live content. Moreover, it ensures only a small amount of information is needed to start a download and to verify incoming chunks of content, thus ensuring short start-up times.

PPSPP has also been designed to be extensible for different transports and use cases. Hence, PPSPP is a generic protocol which can run directly on top of UDP, TCP, or other protocols. As such, PPSPP defines a common set of messages that make up the protocol, which can have different representations on the wire depending on the lower-level protocol used. When the lower-level transport allows, PPSPP can also use different congestion control algorithms.

At present, PPSPP is set to run on top of UDP using LEDBAT for congestion control [RFC6817]. Using LEDBAT enables PPSPP to serve the content after playback (seeding) without disrupting the user who may have moved to different tasks that use its network connection.

PPSPP is also flexible and extensible in the mechanisms it uses to promote client contribution and prevent freeriding, that is, how to deal with peers that only download content but never upload to others. It also allows different schemes for chunk addressing and content integrity protection, if the defaults are not fit for a particular use case. In addition, it can work with different peer discovery schemes, such as centralized trackers or fast Distributed Hash Tables [JIM11]. Finally, in this default setup, PPSPP maintains only a small amount of state per peer. A reference implementation of

PPSPP over UDP is available [SWIFTIMPL].

1.2. Requirements Language

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 [RFC2119].

1.3. Terminology

message

The basic unit of PPSPP communication. A message will have different representations on the wire depending on the transport protocol used. Messages are typically multiplexed into a datagram for transmission.

datagram

A sequence of messages that is offered as a unit to the underlying transport protocol (UDP, etc.). The datagram is PPSPP's Protocol Data Unit (PDU).

content

Either a live transmission, a pre-recorded multimedia asset, or a file.

chunk

The basic unit in which the content is divided. E.g. a block of N kilobyte.

chunk ID

Unique identifier for a chunk of content (e.g. an integer). Its type depends on the chunk addressing scheme used.

chunk specification

An expression that denotes one or more chunk IDs.

chunk addressing scheme

Scheme for identifying chunks and expressing the chunk availability map of a peer in a compact fashion.

chunk availability map

The set of chunks a peer has successfully downloaded and checked the integrity of.

bin

A number denoting a specific binary interval of the content (i.e., one or more consecutive chunks) in the bin numbers chunk addressing scheme (see Section 4).

content integrity protection scheme

Scheme for protecting the integrity of the content while it is being distributed via the peer-to-peer network. I.e. methods for receiving peers to detect whether a requested chunk has been maliciously modified by the sending peer.

hash

The result of applying a cryptographic hash function, more specifically a modification detection code (MDC) [HAC01], such as SHA-1 [FIPS180-3], to a piece of data.

Merkle hash tree

A tree of hashes whose base is formed by the hashes of the chunks of content, and its higher nodes are calculated by recursively computing the hash of the concatenation of the two child hashes (see Section 5.1).

root hash

The root in a Merkle hash tree calculated recursively from the content (see Section 5.1).

swarm

A group of peers participating in the distribution of the same content.

swarm ID

Unique identifier for a swarm of peers, in PPSPP a sequence of bytes. When Merkle hash trees are used for content integrity protection, the identifier is the so-called root hash of the content (video-on-demand). For live streaming, the swarm ID is a public key.

tracker

An entity that records the addresses of peers participating in a swarm, usually for a set of swarms, and makes this membership information available to other peers on request.

choking

When a peer A is choking peer B it means that A is currently not willing to accept requests for content from B.

seeding

Peer A is said to be seeding when A has downloaded a static content asset completely and is now offering it for others to download.

leeching

Peer A is said to be leeching when A has not completely downloaded a static content asset yet or is not offering to upload it to others.

channel

A logical connection between two peers. The channel concept allows peers to use the same transport address for communicating with different peers.

channel ID

Unique, randomly chosen identifier for a channel, local to each peer. So the two peers logically connected by a channel each have a different channel ID for the channel.

In this document the prefixes kilo, mega, etc. denote base 1024.

2. Overall Operation

The basic unit of communication in PPSPP is the message. Multiple messages are multiplexed into a single datagram for transmission. A datagram (and hence the messages it contains) will have different representations on the wire depending on the transport protocol used (see Section 8).

The overall operation of PPSPP is illustrated in the following examples. The examples assume that UDP is used for transport, the Merkle Hash Tree scheme is used for content integrity protection, and that a specific policy is used for selecting which chunks to download.

2.1. Example: Joining a Swarm

Consider a user who wants to watch a video. To play the video, the user clicks on the play button of a HTML5 <video> element shown in his PPSPP-enabled browser. Imagine this element has a PPSPP URL (to be defined elsewhere) identifying the video as its source. The browser passes this URL to its PPSPP protocol handler. Let's call this protocol handler peer A. Peer A parses the URL to retrieve the transport address of a PPSPP tracker and swarm ID of the content. The tracker address may be optional in the presence of a decentralized tracking mechanism.

Peer A now registers with the tracker following the PPSPP tracker protocol [I-D.ietf-ppsp-base-tracker-protocol] and receives the IP address and port of peers already in the swarm, say B, C, and D. Peer A now sends a datagram containing a HANDSHAKE message to B, C, and D.

This message conveys protocol options, in particular, peer A includes the ID of the swarm as the destination peers can listen for multiple swarms on the same transport address.

Peer B and C respond with datagrams containing a HANDSHAKE message and one or more HAVE messages. A HAVE message conveys (part of) the chunk availability of a peer and thus contains a chunk specification that denotes what chunks of the content peer B, resp. C have. Peer D sends a datagram with a HANDSHAKE and HAVE messages, but also with a CHOKE message. The latter indicates that D is not willing to upload chunks to A at present.

2.2. Example: Exchanging Chunks

In response to B and C, A sends new datagrams to B and C containing REQUEST messages. A REQUEST message indicates the chunks that a peer wants to download, and thus contains a chunk specification. The REQUEST messages to B and C refer to disjunct sets of chunks. B and C respond with datagrams containing HAVE, DATA and, in this example, INTEGRITY messages. In the Merkle hash tree content protection scheme (see Section 5.1), the INTEGRITY messages contain all cryptographic hashes that peer A needs to verify the integrity of the content chunk sent in the DATA message. Using these hashes peer A verifies that the chunks received from B and C are correct. It also updates the chunk availability of B and C using the information in the received HAVE messages. In addition, it passes the chunks of video to the user's browser for rendering.

After processing, A sends a datagram containing HAVE messages for the chunks it just received to all its peers. In the datagram to B and C it includes an ACK message acknowledging the receipt of the chunks, and adds REQUEST messages for new chunks. ACK messages are not used when a reliable transport protocol is used. When e.g. C finds that A obtained a chunk (from B) that C did not yet have, C's next datagram includes a REQUEST for that chunk.

Peer D also sends HAVE messages to A when it downloads chunks from other peers. When D is willing to accept REQUESTs from A, D sends a datagram with an UNCHOKE message to inform A. If B or C decide to choke A they sending a CHOKE message and A should then re-request from other peers. B and C may continue to send HAVE, REQUEST, or periodic KEEPALIVE messages such that A keeps sending them HAVE messages.

Once peer A has received all content (video-on-demand use case) it stops sending messages to all other peers that have all content (a.k.a. seeders). Peer A can also contact the tracker or another source again to obtain more peer addresses.

2.3. Example: Leaving a Swarm

To leave a swarm in a graceful way, peer A sends a specific HANDSHAKE message to all its peers (see Section 8.4) and deregisters from the tracker following the (PPSP) tracker protocol. Peers receiving the datagram should remove A from their current peer list. If A crashes ungracefully, peers should remove A from their peer list when they detect it no longer sends messages (see Section 8.15).

3. Messages

In general, no error codes or responses are used in the protocol; absence of any response indicates an error. Invalid messages are discarded, and further communication with the peer SHOULD be stopped. The rationale is that it is sufficient to classify peers as either good (i.e., responding with chunks) or bad and only use the good ones. This behavior allows a peer to deal with slow, crashed and (silent) malicious peers.

Multiple messages are multiplexed into a single datagram for transmission. Messages in a single datagram MUST be processed in the strict order in which they appear in the datagram.

For the sake of simplicity, one swarm of peers deals with one content asset (e.g. file) only. Retrieval of a collections of files can be done either by using multiple swarms or by using an external storage mapping from the linear byte space of a single swarm to different files, transparent to the protocol.

3.1. HANDSHAKE

The initiating peer and the addressed peer MUST send a HANDSHAKE message as the first message in the first datagrams they exchange. The payload of the HANDSHAKE message is a channel ID (see Section 3.11) and a sequence of protocol options. Example options are the content integrity protection scheme used and an option to specify the swarm identifier. The complete set of protocol options are specified in Section 7.

After the handshakes are exchanged, the initiator knows that the peer really responds. Hence, the second datagram the initiator sends MAY already contain some heavy payload, e.g. DATA messages. To minimize the number of initialization round-trips, the first two datagrams exchanged MAY also contain some minor payload, e.g. HAVE messages to indicate the current progress of a peer or a REQUEST (see Section 3.7), but MUST NOT include any DATA message.

3.2. HAVE

The HAVE message is used to convey which chunks a peer has available for download. The set of chunks it has available may be expressed using different chunk addressing and availability map compression schemes, described in Section 4. HAVE messages can be used both for sending a complete overview of a peer's chunk availability as well as for updates to that set.

In particular, whenever a receiving peer P has successfully checked the integrity of a chunk, or interval of chunks, it SHOULD send a HAVE message to all peers Q₁..Q_n it wants to interact with in the near future. A policy in peer P determines when the HAVE is sent. P may send it directly, or peer P may wait until either it has other data to send to Q_i, or until it has received and checked multiple chunks. The policy will depend on how urgent it is to distribute this information to the other peers. This urgency is generally determined in turn by the chunk picking policy (see Section 9.1). In general, the HAVE messages can be piggybacked onto other messages. Peers that do not receive HAVE messages are effectively prevented from downloading the newly available chunks, hence the HAVE message can be used as a method of choking.

The HAVE message MUST contain the chunk specification of the received and verified chunks. A receiving peer MUST NOT send a HAVE message to peers for which the handshake procedure is still incomplete, see Section 13.1. A peer SHOULD NOT send a HAVE message to peers that have the complete content already (e.g. in video-on-demand scenarios).

3.3. DATA

The DATA message is used to transfer chunks of content. The DATA message MUST contain the chunk ID of the chunk and chunk itself. A peer MAY send the DATA messages for multiple chunks in the same datagram. The DATA message MAY contain additional information if needed by the specific congestion control mechanism used. At present PPSPP uses LEDBAT [RFC6817] for congestion control, which requires the current system time to be sent along with the DATA message, so the current system time MUST be included.

3.4. ACK

ACK messages MUST be sent to acknowledge received chunks if PPSPP is run over an unreliable transport protocol. ACK messages MAY be sent if a reliable transport protocol is used. In the former case, a receiving peer that has successfully checked the integrity of a chunk, or interval of chunks C MUST send an ACK message containing a

chunk specification for C. As LEDBAT is used, an ACK message **MUST** contain the one-way delay, computed from the peer's current system time received in the DATA message. A peer **MAY** delay sending ACK messages as defined in the LEDBAT specification.

3.5. INTEGRITY

The INTEGRITY message carries information required by the receiver to verify the integrity of a chunk. Its payload depends on the content integrity protection scheme used. When the Merkle Hash Tree scheme is used, an INTEGRITY message **MUST** contain a cryptographic hash of a subtree of the Merkle hash tree and the chunk specification that identifies the subtree.

As a typical example, when a peer wants to send a chunk and Merkle hash trees are used, it creates a datagram that consists of several INTEGRITY messages containing the hashes the receiver needs to verify the chunk and the actual chunk itself encoded in a DATA message. What are the necessary hashes and the exact rules for encoding them into datagrams is specified in Section 5.3, and Section 5.4, respectively.

3.6. SIGNED_INTEGRITY

The SIGNED_INTEGRITY message carries digitally signed information required by the receiver to verify the integrity of a chunk in live streaming. It logically contains a chunk specification, a timestamp and a digital signature. Its exact payload depends on the live content integrity protection scheme used, see Section 6.1.

3.7. REQUEST

While bulk download protocols normally do explicit requests for certain ranges of data (i.e., use a pull model, for example, BitTorrent [BITTORRENT]), live streaming protocols quite often use a request-less push model to save round trips. PPSP supports both models of operation.

The REQUEST message is used to request one or more chunks from another peer. A REQUEST message **MUST** contain the specification of the chunks the requester wants to download. A peer receiving a REQUEST message **MAY** send out the requested chunks (by means of DATA messages). When peer Q receives multiple REQUESTs from the same peer P, peer Q **SHOULD** process the REQUESTs in the order received. Multiple REQUEST messages **MAY** be sent in one datagram, for example, when a peer wants to request several rare chunks at once.

When live streaming via a push model, a peer receiving REQUESTs also

MAY send some other chunks in case it runs out of requests or for some other reason. In that case the only purpose of REQUEST messages is to provide hints and coordinate peers to avoid unnecessary data retransmission.

3.8. CANCEL

When downloading on demand or live streaming content, a peer can request urgent data from multiple peers to increase the probability of it being delivered on time. In particular, when the specific chunk picking algorithm (see Section 9.1), detects that a request for urgent data might not be served on time, a request for the same data can be sent to a different peer. When a peer P decides to request urgent data from a peer Q, peer P SHOULD send a CANCEL message to all the peers to which the data has been previously requested. The CANCEL message contains the specification of the chunks P no longer wants to request. In addition, when peer Q receives a HAVE message for the urgent data from peer P, peer Q MUST also cancel the previous REQUEST(s) from P. In other words, the HAVE message acts as an implicit CANCEL.

3.9. CHOKE and UNCHOKE

Peer A can send a CHOKE message to peer B to signal it will no longer be responding to REQUEST messages from B, for example, because A's upload capacity is exhausted. Peer A MAY send a subsequent UNCHOKE message to signal that it will respond to new REQUESTs from B again (A SHOULD discard old requests). When peer B receives a CHOKE message from A it MUST NOT send new REQUEST messages and it cannot expect answers to any outstanding ones, as the transfer of chunks is choked. The CHOKE and UNCHOKE messages are informational as responding to REQUESTs is OPTIONAL, see Section 3.7.

3.10. Peer Address Exchange

3.10.1. PEX_REQ and PEX_RES Messages

Peer address exchange messages (or PEX messages for short) are common in many peer-to-peer protocols. They allow peers to exchange the transport addresses of the peers they are currently interacting with, thereby reducing the need to contact a central tracker (or DHT) to discover new peers. The strength of this mechanism is therefore that it enables decentralized peer discovery: after an initial bootstrap no central tracker is needed anymore. Its weakness is that it enables a number of attacks, so it should not be used outside a benign environment unless extra security measures are in place.

PPSP supports peer-address exchange in benign and potentially

hostile environments, as an OPTIONAL feature (not mandatory to implement). The general mechanism works as follows. To obtain some peer addresses a peer A MAY send a PEX_REQ message to peer B. Peer B MAY respond with one or more PEX_RES messages. PPSPP supports three types of PEX_RES reply messages, each containing the address of a single peer Ci. The address in the PEX_RES message MUST be of a peer B has exchanged messages with in the last 60 seconds to guarantee liveliness. Upon receipt, peer A may contact any or none of the returned peers Ci. Alternatively, peers MAY ignore PEX_REQ and PEX_RES messages if uninterested in obtaining new peers or because of security considerations (rate limiting) or any other reason. The PEX messages can be used to construct a dedicated tracker peer.

As indicated, there are three types of PEX_RES messages: PEX_RESv4 containing a single IPv4 address and port, PEX_RESv6 containing a single IPv6 address and port, and a PEX_REScert message. The PEX_RESv4 and PEX_RESv6 MUST only be used in a benign environment, as they provide no guarantees that the host addressed actually participates in a PPSPP swarm.

To use PEX in PPSPP in a potentially hostile environment, such as the Internet, three conditions must be met:

1. Peer transport addresses must be relatively stable.
2. PEX_REScert messages must be used instead of PEX_RESv4 and PEX_RESv6.
3. A peer must not obtain all its peer addresses through PEX.

The full security analysis for PEX messages can be found in Section 13.2. A PEX_REScert message carries a swarm-membership certificate rather than an IP address and port. A membership certificate for peer C states that peer C at address (ipC,portC) is part of swarm S at time T and is cryptographically signed by an issuer. The receiver A can check the certificate for a valid signature by a trusted issuer, the right swarm and liveliness and only then consider contacting C. These swarm-membership certificates correspond to signed node descriptors in secure decentralized peer sampling services [SPS].

Several designs are possible for the security environment for these membership certificates. That is, there are different designs possible for who signs the membership certificates and how public keys are distributed. Section 13.2.2 describes an example where a central tracker acts as the Certification Authority.

In a potentially hostile environment, peers must also ensure that


they do not end up interacting only with malicious peers when using the peer-address exchange feature. To this extent, peers **MUST** ensure that part of their connections are to peers whose addresses came from a trusted and secured tracker (see Section 13.2.3).

Once a PPSPP implementation has obtained a list of peers (either via PEX, from a central tracker or via a DHT), it has to determine which peers to actually contact. In this process, a PPSPP implementation can benefit from information by network or content providers to help improve network usage and boost PPSPP performance. How a P2P system like PPSPP can perform these optimizations using the ALTO protocol is described in detail in [I-D.ietf-alto-protocol], Section 7.

3.11. Channels

It is increasingly complex for peers to enable communication between each other due to NATs and firewalls. Therefore, PPSPP uses a multiplexing scheme, called channels, to allow multiple swarms to use the same transport address. Channels loosely correspond to TCP connections and each channel belongs to a single swarm, as illustrated in Figure 1. As with TCP connections, a channel is identified by a unique identifier local to the peer at each end of the connection (cf. TCP port), which **MUST** be randomly chosen. In other words, the two peers connected by a channel use different IDs to denote the same channel. The IDs are different and random for security reasons, see Section 13.1.

In the PPSPP-over-UDP encapsulation (Section 8.3), when a channel C has been established between peer A and peer B, the datagrams containing messages from A to B are prefixed with the four byte channel ID allocated by peer B, and vice versa for datagrams from B to A. The channel IDs used are exchanged as part of the handshake procedure, see Section 8.4. In that procedure, the channel ID with value 0 is used for the datagram that initiates the handshake. PPSPP can be used in combination with STUN [RFC5389].



Network stack of a PPSPP peer that is reachable on UDP port 6778 and is connected via channel 481 to one peer in swarm A and two peers in swarm B via channels 836 and 372, respectively. Channel ID 0 is special and is used for handshaking.

Figure 1

3.12. Keep Alive Signalling

A peer **SHOULD** send a "keep alive" message periodically to each peer it wants to interact with in the future, but has no other messages to send them at present. Periodically sending "keep alive" messages prevents other peers from closing the connection after a predefined time interval of 3 minutes, as described in Section 8.15. PPSPP does not define an explicit message type for "keep alive" messages. In the PPSPP-over-UDP encapsulation they are implemented as simple datagrams consisting of a 4-byte channel ID only, see Section 8.3 and Section 8.4.

4. Chunk Addressing Schemes

PPSPP can use different methods of chunk addressing, that is, support different ways of identifying chunks and different ways of expressing the chunk availability map of a peer in a compact fashion.

All peers in a swarm **MUST** use the same chunk addressing method.

4.1. Start-End Ranges

A chunk specification consists of a single (start specification,end specification) pair that identifies a range of chunks (end inclusive). The start and end specifications can use one of multiple addressing schemes. Two schemes are currently defined, chunk ranges and byte ranges.

4.1.1. Chunk Ranges

The start and end specification are both chunk identifiers. Chunk identifiers are 32-bit or 64-bit unsigned integers. A PPSP peer MUST support this scheme.

```

      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
~                               Start chunk (32 or 64)                               ~
+-----+-----+-----+-----+-----+-----+-----+-----+
~                               End chunk (32 or 64)                               ~
+-----+-----+-----+-----+-----+-----+-----+-----+

```

4.1.2. Byte Ranges

The start and end specification are 64-bit byte offsets in the content. The support for this scheme is OPTIONAL.

```

      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
|                               Start byte offset (64)                               |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                               End byte offset (64)                               |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                               |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

4.2. Bin Numbers


PPSP introduces a novel method of addressing chunks of content called "bin numbers" (or "bins" for short). Bin numbers allow the addressing of a binary interval of data using a single integer. This reduces the amount of state that needs to be recorded per peer and the space needed to denote intervals on the wire, making the protocol light-weight. In general, this numbering system allows PPSP to work with simpler data structures, e.g. to use arrays instead of binary

trees, thus reducing complexity. The support for this scheme is OPTIONAL.

In bin addressing, the smallest binary interval is a single chunk (e.g. a block of bytes which may be of variable size), the largest interval is a complete range of 2^{63} chunks. In a novel addition to the classical scheme, these intervals are numbered in a way which lays them out into a vector nicely, which is called bin numbering, as follows. Consider an chunk interval of width W . To derive the bin numbers of the complete interval and the subintervals, a minimal balanced binary tree is built that is at least W chunks wide at the base. The leaves from left-to-right correspond to the chunks $0..W-1$ in the interval, and have bin number $I*2$ where I is the index of the chunk (counting beyond $W-1$ to balance the tree). The bin number of higher level nodes P in the tree is calculated as follows:

$$\text{binP} = (\text{binL} + \text{binR}) / 2$$

where binL is the bin of node P 's left-hand child and binR is the bin of node P 's right-hand child. Given that each node in the tree represents a subinterval of the original interval, each such subinterval now is addressable by a bin number, a single integer. The bin number tree of an interval of width $W=8$ looks like this:



The bin number tree of an interval of width $W=8$

Figure 2

So bin 7 represents the complete interval, bin 3 represents the interval of chunk $0..3$, bin 1 represents the interval of chunks 0 and 1, and bin 2 represents chunk C1. The special numbers $0xFFFFFFFF$ (32-bit) or $0xFFFFFFFFFFFFFFFF$ (64-bit) stands for an empty interval,

and 0x7FFF...FFF stands for "everything".

When bin numbering is used, the ID of a chunk is its corresponding (leaf) bin number in the tree and the chunk specification in HAVE and ACK messages is equal to a single bin number (32-bit or 64-bit), as follows.

```

      0               1               2               3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
~                               Bin number (32 or 64)                               ~
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

4.3. In Messages

4.3.1. In HAVE Messages

When a receiving peer has successfully checked the integrity of a chunk or interval of chunks it MUST send a HAVE message to all peers it wants to interact with. The latter allows the HAVE message to be used as a method of choking. The HAVE message MUST contain the chunk specification of the biggest complete interval of all chunks the receiver has received and checked so far that fully includes the interval of chunks just received. So the chunk specification MUST denote at least the interval received, but the receiver is supposed to aggregate and acknowledge bigger intervals, when possible.

As a result, every single chunk is acknowledged a logarithmic number of times. That provides some necessary redundancy of acknowledgments and sufficiently compensates for unreliable transport protocols.

Implementation note:

To record which chunks a peer has in the state that an implementation keeps for each peer, an implementation MAY use the efficient "binmap" data structure, which is a hybrid of a bitmap and a binary tree, discussed in detail in [BINMAP].

4.3.2. In ACK Messages

PPSPP peers MUST use ACK messages to acknowledge received chunks if an unreliable transport protocol is used. When a receiving peer has successfully checked the integrity of a chunk or interval of chunks C it MUST send a ACK message containing the chunk specification of its biggest, complete interval covering C to the sending peer (see HAVE).

5. Content Integrity Protection

PPSPP can use different methods for protecting the integrity of the content while it is being distributed via the peer-to-peer network. More specifically, PPSPP can use different methods for receiving peers to detect whether a requested chunk has been maliciously modified by the sending peer. In benign environments, content integrity protection can be disabled.

For static content, PPSPP currently defines one method for protecting integrity, called the Merkle Hash Tree scheme. If PPSPP operates over the Internet, this scheme **MUST** be used. If PPSPP operates in a benign environment this scheme **MAY** be used. So the scheme is mandatory-to-implement, to satisfy the requirement of strong security for an IETF protocol [RFC3365]. An extended version of the scheme is used to efficiently protect dynamically generated content (live streams), as explained below and in Section 6.1.


The Merkle Hash Tree scheme can work with different chunk addressing schemes. All it requires is the ability to address a range of chunks. In the following description abstract node IDs are used to identify nodes in the tree. On the wire these are translated to the corresponding range of chunks in the chosen chunk addressing scheme.

5.1. Merkle Hash Tree Scheme

PPSPP uses a method of naming content based on self-certification. In particular, content in PPSPP is identified by a single cryptographic hash that is the root hash in a Merkle hash tree calculated recursively from the content [ABMRKL]. This self-certifying hash tree allows every peer to directly detect when a malicious peer tries to distribute fake content. It also ensures only a small amount of information is needed to start a download (the root hash and some peer addresses). For live streaming a dynamic tree and a public key are used, see below.

The Merkle hash tree of a content asset that is divided into N chunks is constructed as follows. Note the construction does not assume chunks of content to be fixed size. Given a cryptographic hash function, more specifically a modification detection code (MDC) [HAC01], such as SHA1, the hashes of all the chunks of the content are calculated. Next, a binary tree of sufficient height is created. Sufficient height means that the lowest level in the tree has enough nodes to hold all chunk hashes in the set, as with bin numbering. The figure below shows the tree for a content asset consisting of 7 chunks. As before with the content addressing scheme, the leaves of the tree correspond to a chunk and in this case are assigned the hash of that chunk, starting at the left-most leaf. As the base of the

tree may be wider than the number of chunks, any remaining leaves in the tree are assigned an empty hash value of all zeros. Finally, the hash values of the higher levels in the tree are calculated, by concatenating the hash values of the two children (again left to right) and computing the hash of that aggregate. If the two children are empty hashes, the parent is an empty all zeros hash as well (to save computation). This process ends in a hash value for the root node, which is called the "root hash". Note the root hash only depends on the content and any modification of the content will result in a different root hash.



The Merkle hash tree of an interval of width W=8

Figure 3

5.2. Content Integrity Verification

Assuming a peer receives the root hash of the content it wants to download from a trusted source, it can check the integrity of any chunk of that content it receives as follows. It first calculates the hash of the chunk it received, for example chunk C4 in the previous figure. Along with this chunk it MUST receive the hashes required to check the integrity of that chunk. In principle, these are the hash of the chunk's sibling (C5) and that of its "uncles". A chunk's uncles are the sibling Y of its parent X, and the uncle of that Y, recursively until the root is reached. For chunk C4 its uncles are nodes 13 and 3, marked with * in the figure. Using this information the peer recalculates the root hash of the tree, and compares it to the root hash it received from the trusted source. If they match the chunk of content has been positively verified to be

the requested part of the content. Otherwise, the sending peer either sent the wrong content or the wrong sibling or uncle hashes. For simplicity, the set of sibling and uncles hashes is collectively referred to as the "uncle hashes".

In the case of live streaming the tree of chunks grows dynamically and the root hash is undefined or, more precisely, transient, as long as new data is generated by the live source. Section 6.1.2 defines a method for content integrity verification for live streams that works with such a dynamic tree. Although the tree is dynamic, content verification works the same for both live and predefined content, resulting in a unified method for both types of streaming.

5.3. The Atomic Datagram Principle

As explained above, a datagram consists of a sequence of messages. Ideally, every datagram sent must be independent of other datagrams, so each datagram SHOULD be processed separately and a loss of one datagram must not disrupt the flow of datagrams between two peers. Thus, as a datagram carries zero or more messages, neither messages nor message interdependencies SHOULD span over multiple datagrams.

This principle implies that as any chunk is verified using its uncle hashes the necessary hashes SHOULD be put into the same datagram as the chunk's data. If this is not possible because of a limitation on datagram size, the necessary hashes MUST be sent first in one or more datagrams. As a general rule, if some additional data is still missing to process a message within a datagram, the message SHOULD be dropped.

The hashes necessary to verify a chunk are in principle its sibling's hash and all its uncle hashes, but the set of hashes to send can be optimized. Before sending a packet of data to the receiver, the sender inspects the receiver's previous acknowledgments (HAVE or ACK) to derive which hashes the receiver already has for sure. Suppose, the receiver had acknowledged chunks C0 and C1 (first two chunks of the file), then it must already have uncle hashes 5, 11 and so on. That is because those hashes are necessary to check C0 and C1 against the root hash. Then, hashes 3, 7 and so on must be also known as they are calculated in the process of checking the uncle hash chain. Hence, to send chunk C7, the sender needs to include just the hashes for nodes 14 and 9, which let the data be checked against hash 11 which is already known to the receiver.

The sender MAY optimistically skip hashes which were sent out in previous, still unacknowledged datagrams. It is an optimization trade-off between redundant hash transmission and possibility of collateral data loss in the case some necessary hashes were lost in

the network so some delivered data cannot be verified and thus has to be dropped. In either case, the receiver builds the Merkle tree on-demand, incrementally, starting from the root hash, and uses it for data validation.

In short, the sender **MUST** put into the datagram the missing hashes necessary for the receiver to verify the chunk. The receiver **MUST** remember all the hashes it needs to verify missing chunks that it still wants to download. Note that the latter implies that a hardware-limited receiver **MAY** forget some hashes if it does not plan to announce possession of these chunks to others (i.e., does not plan to send **HAVE** messages.)

5.4. INTEGRITY Messages

Concretely, a peer that wants to send a chunk of content creates a datagram that **MUST** consist of a list of **INTEGRITY** messages followed by a **DATA** message. If the **INTEGRITY** messages and **DATA** message cannot be put into a single datagram because of a limitation on datagram size, the **INTEGRITY** messages **MUST** be sent first in one or more datagrams. The list of **INTEGRITY** messages sent **MUST** contain a **INTEGRITY** message for each hash the receiver misses for integrity checking. A **INTEGRITY** message for a hash **MUST** contain the chunk specification corresponding to the node ID of the hash and the hash data itself. The chunk specification corresponding to a node ID is defined as the range of chunks formed by the leaves of the subtree rooted at the node. For example, node 3 in Figure 3 denotes chunks 0,2,4,6, so the chunk specification should denote that interval. The list of **INTEGRITY** messages **MUST** be sorted in order of the tree height of the nodes, descending. The **DATA** message **MUST** contain the chunk specification of the chunk and chunk itself. A peer **MAY** send the required messages for multiple chunks in the same datagram, depending on the encapsulation.

5.5. Discussion and Overhead

The current method for protecting content integrity in BitTorrent [BITTORRENT] is not suited for streaming. It involves providing clients with the hashes of the content's chunks before the download commences by means of metadata files (called .torrent files in BitTorrent.) However, when chunks are small as in the current UDP encapsulation of PPSP this implies having to download a large number of hashes before content download can begin. This, in turn, increases time-till-playback for end users, making this method unsuited for streaming.

The overhead of using Merkle hash trees is limited. The size of the hash tree expressed as the total number of nodes depends on the

number of chunks the content is divided (and hence the size of chunks) following this formula:

$$nnodes = \text{math.pow}(2, \text{math.log}(nchunks, 2) + 1)$$

In principle, the hash values of all these nodes will have to be sent to a peer once for it to verify all chunks. Hence the maximum on-the-wire overhead is $\text{hashsize} * nnodes$. However, the actual number of hashes transmitted can be optimized as described in Section 5.3. To see a peer can verify all chunks whilst receiving not all hashes, consider the example tree in Section 5.1.

In case of a simple progressive download, of chunks 0,2,4,6, etc. the sending peer will send the following hashes:

Chunk	Node IDs of hashes sent
0	2,5,11
2	- (receiver already knows all)
4	6
6	-
8	10,13 (hash 3 can be calculated from 0,2,5)
10	-
12	14
14	-
Total	# hashes 7

Table 1: Overhead for the example tree

So the number of hashes sent in total (7) is less than the total number of hashes in the tree (16), as a peer does not need to send hashes that are calculated and verified as part of earlier chunks.


5.6. Automatic Detection of Content Size

In PPSPP, the root hash of a static content asset, such as a video file, along with some peer addresses is sufficient to start a download. In addition, PPSPP can reliably and automatically derive the size of such content from information received from the network when fixed sized chunks are used. As a result, it is not necessary to include the size of the content asset as the metadata of the content, in addition to the root hash. Implementations of PPSPP MAY use this automatic detection feature. Note this feature is the only feature of PPSPP that requires that a fixed-sized chunk is used.

5.6.1. Peak Hashes

The ability for a newcomer peer to detect the size of the content depends heavily on the concept of peak hashes. Peak hashes, in general, enable two cornerstone features of PPSPP: reliable file size detection and download/live streaming unification (see Section 6). The concept of peak hashes depends on the concepts of filled and incomplete nodes. Recall that when constructing the binary trees for content verification and addressing the base of the tree may have more leaves than the number of chunks in the content. In the Merkle hash tree these leaves were assigned empty all-zero hashes to be able to calculate the higher level hashes. A filled node is now defined as a node that corresponds to an interval of leaves that consists only of hashes of content chunks, not empty hashes. Reversely, an incomplete (not filled) node corresponds to an interval that contains also empty hashes, typically an interval that extends past the end of the file. In the following figure nodes 7, 11, 13 and 14 are incomplete the rest is filled.

Formally, a peak hash is the hash of a filled node in the Merkle tree, whose sibling is an incomplete node. Practically, suppose a file is 7162 bytes long and a chunk is 1 kilobyte. That file fits into 7 chunks, the tail chunk being 1018 bytes long. The Merkle tree for that file is shown in Figure 4. Following the definition the peak hashes of this file are in nodes 3, 9 and 12, denoted with a *. E denotes an empty hash.



Peak hashes in a Merkle hash tree.

Figure 4

Peak hashes can be explained by the binary representation of the number of chunks the file occupies. The binary representation for 7 is 111. Every "1" in binary representation of the file's packet length corresponds to a peak hash. For this particular file there are indeed three peaks, nodes 3, 9, 12. The number of peak hashes for a file is therefore also at most logarithmic with its size.

A peer knowing which nodes contain the peak hashes for the file can therefore calculate the number of chunks it consists of, and thus get an estimate of the file size (given all chunks but the last are fixed size). Which nodes are the peaks can be securely communicated from one (untrusted) peer A to another B by letting A send the peak hashes and their node IDs to B. It can be shown that the root hash that B obtained from a trusted source is sufficient to verify that these are indeed the right peak hashes, as follows.

Lemma: Peak hashes can be checked against the root hash.

Proof: (a) Any peak hash is always the left sibling. Otherwise, be it the right sibling, its left neighbor/sibling must also be a filled node, because of the way chunks are laid out in the leaves, contradiction. (b) For the rightmost peak hash, its right sibling is zero. (c) For any peak hash, its right sibling might be calculated using peak hashes to the left and zeros for empty nodes. (d) Once the right sibling of the leftmost peak hash is calculated, its parent might be calculated. (e) Once that parent is calculated, we might trivially get to the root hash by concatenating the hash with zeros and hashing it repeatedly.

Informally, the Lemma might be expressed as follows: peak hashes cover all data, so the remaining hashes are either trivial (zeros) or might be calculated from peak hashes and zero hashes.

Finally, once peer B has obtained the number of chunks in the content it can determine the exact file size as follows. Given that all chunks except the last are fixed size B just needs to know the size of the last chunk. Knowing the number of chunks B can calculate the node ID of the last chunk and download it. As always B verifies the integrity of this chunk against the trusted root hash. As there is only one chunk of data that leads to a successful verification the size of this chunk must be correct. B can then determine the exact file size as

$$(\text{number of chunks} - 1) * \text{fixed chunk size} + \text{size of last chunk}$$

5.6.2. Procedure

A PPSPP implementation that wants to use automatic size detection MUST operate as follows. When a peer A sends a DATA message for the first time to a peer B, A MUST first send all the peak hashes for the content, unless B has already signalled earlier in the exchange that it knows the peak hashes by having acknowledged any chunk. If they are needed, the peak hashes MUST be sent as an extra list of uncle hashes for the chunk, before the list of actual uncle hashes of the chunk as described in Section 5.3. The receiver B MUST check the peak hashes against the root hash to determine the approximate content size. To obtain the definite content size peer B MUST download the last chunk of the content from any peer that offers it.

As an example, let's consider a 7162 bytes long file, which fits in 7 chunks of 1 kilobyte, distributed by a peer A. Figure 4 shows the relevant Merkle hash tree. A peer B which only knows the root hash of the file, after successfully connecting to A, requests the first chunk of data, C0 in Figure 4. Peer A replies to B by including in the datagram the following messages in this specific order. First the three peak hashes of this particular file, the hashes of nodes 3, 9 and 12. Second, the uncle hashes of C0, followed by the DATA message containing the actual content of C0. Upon receiving the peak hashes, peer B checks them against the root hash determining that the file is 7 chunks long. To establish the exact size of the file, peer B needs to request and retrieve the last chunk containing data, C6 in Figure 4. Once the last chunk has been retrieved and verified, peer B concludes that it is 1018 bytes long, hence determining that the file is exactly 7162 bytes long.

6. Live Streaming

The set of messages defined above can be used for live streaming as well. In a pull-based model, a live streaming injector can announce the chunks it generates via HAVE messages, and peers can retrieve them via REQUEST messages. Areas that need special attention are content authentication and chunk addressing (to achieve an infinite stream of chunks).

6.1. Content Authentication

For live streaming, PPSPP supports two methods for a peer to authenticate the content it receives from another peer, called "Sign All" and "Unified Merkle Tree".

In the "Sign All" method, the live injector signs each chunk of content using a private key and peers, upon receiving the chunk,

check the signature using the corresponding public key obtained from a trusted source. Support for this method is OPTIONAL.

In the "Unified Merkle Tree" method, PPSPP combines the Merkle Hash Tree scheme for static content with signatures to unify the video-on-demand and live streaming scenarios. The use of Merkle hash trees reduces the number of signing and verification operations, hence providing a similar signature amortization to the approach described in [SIGMCAST]. If PPSPP operates over the Internet, the "Unified Merkle Tree" method MUST be used. If the protocol operates in a benign environment the method MAY be used. So this method is mandatory-to-implement.

In both methods the swarm ID consists of a public key encoded as in a DNSSEC DNSKEY resource record without BASE-64 encoding [RFC4034]. In particular, the swarm ID consists of a 1 byte Algorithm field that identifies the public key's cryptographic algorithm and determines the format of the Public Key field that follows. The value of this Algorithm field is one of the Domain Name System Security (DNSSEC) Algorithm Numbers [IANADNSSECALGNUM]. The RSASHA1 [RFC4034], ECDSAP256SHA256 and ECDSAP384SHA384 [RFC6605] algorithms are MANDATORY to implement.

```

      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
| Algo Number(8)|~
+-----+-----+-----+-----+-----+-----+-----+-----+
~                               DNSSEC Public Key (variable)                               ~
+-----+-----+-----+-----+-----+-----+-----+-----+

```

6.1.1. Sign All

In the "Sign All" method, the live injector signs each chunk of content using a private key and peers, upon receiving the chunk, check the signature using the corresponding public key obtained from a trusted source. In particular, in PPSPP, the swarm ID of the live stream is that public key.

A peer that wants to send a chunk of content creates a datagram that MUST contain a SIGNED_INTEGRITY message with the chunk's signature, followed by a DATA message with the actual chunk. If the SIGNED_INTEGRITY message and DATA message cannot be contained into a single datagram, because of a limitation on datagram size, the SIGNED_INTEGRITY message MUST be sent first in a separate datagram. The SIGNED_INTEGRITY message consists of the chunk specification, the timestamp, and the digital signature.


The digital signature algorithm which is used, is determined by the Live Signature Algorithm protocol option, see Section 7.7. The signature is computed over a concatenation of the on-the-wire representation of the chunk specification, a 64-bit NTP timestamp [RFC5905], and the chunk, in that order. The timestamp is the time signature that was made at the injector in UTC.

6.1.2. Unified Merkle Tree

In this method, the chunks of content are used as the basis for a Merkle hash tree as for static content. However, because chunks are continuously generated, this tree is not static, but dynamic. As a result, the tree does not have a root hash, or more precisely has a transient root hash. A public key therefore serves as swarm ID of the content. It is used to digitally sign updates to the tree, allowing peers to expand it based on trusted information using the following process.

6.1.2.1. Signed Munro Hashes

The live injector generates a number of chunks, denoted `NCHUNKS_PER_SIG`, corresponding to fixed power of 2 (`NCHUNKS_PER_SIG >= 2`), which are added as new leaves to the existing hash tree. As a result of this expansion the hash tree contains a new subtree, that is `NCHUNKS_PER_SIG` chunks wide at the base. The root of this new subtree is referred to as the munro of that subtree, and its hash as the munro hash of the subtree, illustrated in Figure 5. In this figure, node 5 is the new munro, labeled with a \$ sign.



Expanded live tree. With `NCHUNKS_PER_SIG=2`, node 5 is the munro for the new subtree spanning 4 and 6. Node 1 is the munro for the subtree spanning chunks 0 and 2, created in the previous iteration.


Figure 5

Informally, the process now proceeds as follows. The injector now signs only the munro hash of the new subtree using its private key.

Next, the injector announces the existence of the new subtree to its peers using HAVE messages. When a peer, in response to the HAVE messages, requests a chunk from the new subtree, the injector first sends the signed munro hash corresponding to the requested chunk. Afterwards, similar to static content, the injector sends the uncle hashes necessary to verify that chunk, as in Section 5.1. In particular, the injector sends the uncle hashes necessary to verify the requested chunk against the munro hash. This differs from static content, where the verification takes places against the root hash. Finally, the injector sends the actual chunk.

The receiving peer verifies the signature on the signed munro using the swarm ID (a public key), and updates its hash tree. As the peer now knows the munro hash is trusted, it can verify all chunks in the subtree against this munro hash, using the accompanying uncle hashes as in Section 5.1.

To illustrate this procedure, lets consider the next iteration in the process. The injector has generated the current tree shown in Figure 5 and it is connected to several peers that currently have the same tree and all posses chunks 0, 2, 4 and 6. When the injector generates two new chunks, NCHUNKS_PER_SIG=2, the hash tree expands as shown in Figure 6. The two new chunks, 8 and 10, extend the tree on the right side, and to accommodate them a new root is created, node 7. As this tree is wider at the base than the actual number of chunks, there are currently two empty leaves. The munro node for the new subtree is 9, labeled with a \$ sign.



Expanded live tree. With NCHUNKS_PER_SIG=2, node 9 is the munro of the newly added subtree spanning chunks 8 and 10.

Figure 6

The injector now needs to inform its peers of the updated tree, communicating the addition of the new munro hash 9. Hence, it sends a HAVE message with a chunk specification for nodes 8+10 to its peers. As a response, a peer P requests the newly created chunk, e.g. chunk 8, from the injector by sending a REQUEST message. In reply, the injector sends the signed munro hash of node 9 as an INTEGRITY message with the hash of node 9, and a SIGNED_INTEGRITY message with the signature of the hash of node 9. These messages are followed by an INTEGRITY message with the hash of node 10, and a DATA message with chunk 8.

Upon receipt, peer P verifies the signature of the munro and expands its view of the tree. Next, the peer computes the hash of chunk 8 and combines it with the received hash of node 10, computing the expected hash of node 9. He can then verify the content of chunk 8 by computing the computed hash of node 9 with the munro hash of the same node he just received, hence P has successfully verified the integrity of chunk 8.

This procedure requires just one signing operation for every NCHUNKS_PER_SIG chunks created, and one verification operation for every NCHUNKS_PER_SIG received, making it much cheaper than "Sign All". A receiving peer does additionally need to check one or more hashes per chunk via the Merkle Tree scheme, but this has less hardware requirements than a signature verification for every chunk. This approach is similar to signature amortization via Merkle Tree Chaining [SIGMCAST]. The downside of scheme is in an increased latency. A peer cannot download the new chunks until the injector has computed the signature and announced the subtree. A peer MUST check the signature before forwarding the chunks to other peers [POLLIVE].

The number of chunks per signature NCHUNKS_PER_SIG MUST be a fixed power of 2 for simplicity. NCHUNKS_PER_SIG MUST be larger than 1 for performance reasons. There are two related factors to consider when choosing a value for NCHUNKS_PER_SIG. First, the allowed CPU load on clients due to signature verifications, given the expected bitrate of the stream. To achieve a low CPU load in a high bitrate stream, NCHUNKS_PER_SIG should be high. Second, the effect on latency, which increases when NCHUNKS_PER_SIG gets higher, as just discussed. Note how the procedure does not preclude the use of variable-sized chunks.

This method of integrity verification provides an additional benefit. If the system includes some peers that saved the complete broadcast, as soon as the broadcast ends, the content is available as a video-on-demand download using the now stabilized tree and the final root hash as swarm identifier. Peers which saved all the chunks, can now announce the root hash to the tracking infrastructure and instantly

seed the content.

6.1.2.2. Munro Signature Calculation

The digital signature algorithm used is determined by the Live Signature Algorithm protocol option, see Section 7.7. The signature is computed over a concatenation of the on-the-wire representation of the chunk specification of the munro, a 64-bit NTP timestamp [RFC5905], and the munro hash, in that order. The timestamp is the time signature that was made at the injector in UTC.

6.1.2.3. Procedure

Formally, the injector MUST NOT send a HAVE message for chunks in the new subtree until it has computed the signed munro hash for that subtree.

When peer B requests a chunk C from peer A (either the injector or another peer), and peer A decides to reply, it must do so as follows. First, peer A MUST send an INTEGRITY message with the chunk specification for the munro of chunk C and the munro's hash, followed by a SIGNED_INTEGRITY message with the chunk specification for the munro, timestamp and its signature, in a single datagram, unless B indicated earlier in the exchange that it already possess a chunk with the same corresponding munro (by means of HAVE or ACK messages). Following these two messages (if any), peer A MUST send the necessary missing uncles hashes needed for verifying the chunk against its munro hash, and the chunk itself, as described in Section 5.4, sharing datagrams if possible.

6.1.2.4. Secure Tune In

When a peer tunes into a live stream it has to determine what is the last chunk the injector has generated. To facilitate this process in the Unified Merkle Tree scheme, each peer shares its knowledge about the injector's chunks with the others by exchanging their latest signed munro hashes, as follows.

Recall that in PPSPP, when peer A initiates a channel with peer B, peer A sends a first datagram with a HANDSHAKE message, and B responds with a second datagram also containing a HANDSHAKE message (see Section 3.1). When A sends a third datagram to B, and it is received by B both peers know that the other is listening on its stated transport address. B is then allowed to send heavy payload like DATA messages in the fourth datagram. Peer A can already safely do that in the third datagram.

In the Unified Merkle Tree scheme, peer A MUST send its right-most

signed munro hash to B in the third datagram, and in any subsequent datagrams to B, until B indicates that it possess a chunk with the same corresponding munro or a more recent munro (by means of a HAVE or ACK message). B may already have indicated this fact by means of HAVE messages in the second datagram. Conversely, when B sends the fourth datagram or any subsequent datagram to A, B MUST send its right-most signed munro hash, unless A indicated knowledge of it or more recent munros. The right-most signed munro hash of a peer is defined as the munro hash signed by the injector of the right-most subtree of width NCHUNKS_PER_SIG chunks in the peer's Merkle hash tree. Peer A and B MUST NOT send the signed munro hash in the first, respectively, second datagram as it is considered heavy payload.

When a peer receives a SIGNED_INTEGRITY message with a signed munro hash but the timestamp is too old, the peer MUST discard the message. Otherwise it SHOULD use the signed munro to update its hash tree and pick a tune-in point in the live stream. A peer may use the information from multiple peers to pick the tune-in point.

6.2. Forgetting Chunks

As a live broadcast progresses a peer may want to discard the chunks that it already played out. Ideally, other peers should be aware of this fact such that they will not try to request these chunks from this peer. This could happen in scenarios where live streams may be paused by viewers, or viewers are allowed to start late in a live broadcast (e.g., start watching a broadcast at 20:35 whereas it began at 20:30).

PPSPP provides a simple solution for peers to stay up-to-date with the chunk availability of a discarding peer. A discarding peer in a live stream MUST enable the Live Discard Window protocol option, specifying how many chunks/bytes it caches before the last chunk/byte it advertised as being available (see Section 7.9). Its peers SHOULD apply this number as a sliding window filter over the peer's chunk availability as conveyed via its HAVE messages.

Three factors are important when deciding for an appropriate value for this option: the desired amount of playback buffer for peers, the bitrate of the stream and the available resources of the peer. Consider the case of a fresh peer joining the stream. The size of the discard window of the peers it connects to influences how much data it can directly download to establish its prebuffer. If the window is smaller than the desired buffer, the fresh peer has to wait until the peers downloaded more of the stream before it can start playback. As media buffers are generally specified in terms of a number of seconds, the size of the discard window also related to the (average) bitrate of the stream. Finally, if a peer has little

resources to store chunks and metadata it should chose a small discard window.

7. Protocol Options

The HANDSHAKE message in PPSPP can contain the following protocol options. Unless stated otherwise, a protocol option consists of an 8-bit code followed by an 8-bit value. Larger values are all encoded big-endian. Each protocol option is explained in the following subsections.

Code	Description
0	Version
1	Minimum Version
2	Swarm Identifier
3	Content Integrity Protection Method
4	Merkle Hash Tree Function
5	Live Signature Algorithm
6	Chunk Addressing Method
7	Live Discard Window
8	Supported Messages
9-254	Unassigned
255	End Option

Table 2: PPSP Peer Protocol Options

7.1. End Option

A peer MUST conclude the list of protocol options with the end option. Subsequent octets should be considered protocol messages. The code for the end option is 255, and unlike others it has no value octet, so the option's length is 1 octet.

```

 0 1 2 3 4 5 6 7
+---+---+---+---+
| 1 1 1 1 1 1 1 1 |
+---+---+---+---+
```

7.2. Version

A peer MUST include the maximum version of the PPSPP protocol it supports as the first protocol option in the list. The code for this option is 0. Defined values are listed in Table 3.

Version	Description
1	Protocol as described in this document
2-255	Unassigned

Table 3: PPSP Peer Protocol Version Numbers

```

0                               1
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5
+-----+-----+-----+-----+
|0 0 0 0 0 0 0 0 0| Version (8) |
+-----+-----+-----+-----+

```

7.3. Minimum Version

When a peer initiates the handshake it MUST include the minimum version of the PPSP protocol it supports in the list of protocol options, following the Min/max versioning scheme defined in [RFC6709], Section 4.1, strategy 5. The code for this option is 1. Defined values are listed in Table 3.

```

0                               1
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5
+-----+-----+-----+-----+
|0 0 0 0 0 0 0 0 1| Min. Ver. (8) |
+-----+-----+-----+-----+

```

7.4. Swarm Identifier

When a peer initiates the handshake it MUST include a single swarm identifier option. In other cases a peer MAY include a swarm identifier option, as an end-to-end check. This option has the following structure:

```

0                               1                               2                               3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
|0 0 0 0 0 0 0 1 0| Swarm ID Length (16) | ~
+-----+-----+-----+-----+-----+-----+-----+-----+
~ Swarm Identifier (variable) ~
+-----+-----+-----+-----+-----+-----+-----+-----+

```

The Swarm ID Length field contains the length of the single Swarm Identifier that follows in bytes. The Length field is 16 bits wide to allow for large public keys as identifiers in live streaming.

Each PPSPP peer knows the IDs of the swarms it joins so this information can be immediately verified upon receipt.

7.5. Content Integrity Protection Method

A peer MUST include the content integrity method used by a swarm. The code for this option is 3. Defined values are listed in Table 4.

Method	Description
0	No integrity protection
1	Merkle Hash Tree
2	Sign All
3	Unified Merkle Tree
4-255	Unassigned

Table 4: PPSP Peer Content Integrity Protection Methods

The "Merkle Hash Tree" method is the default for static content, see Section 5.1. "Sign All", and "Unified Merkle Tree" are for live content, see Section 6.1, with "Unified Merkle Tree" being the default.

```

0                               1
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5
+-----+-----+
|0 0 0 0 0 0 1 1|   CIPM (8)   |
+-----+-----+
```

7.6. Merkle Tree Hash Function

When the content integrity protection method is "Merkle Hash Tree" this option defining which hash function is used for the tree MUST be included. The code for this option is 4. Defined values are listed in Table 5 (see [FIPS180-3] for the function semantics).

Function	Description
0	SHA1
1	SHA-224
2	SHA-256
3	SHA-384
4	SHA-512
5-255	Unassigned

Table 5: PPSP Peer Protocol Merkle Hash Functions

Implementations MUST support SHA1, see Section 13.5, which is also the default.

```

0                               1
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5
+-----+
|0 0 0 0 0 1 0 0|   MHF (8)   |
+-----+
```

7.7. Live Signature Algorithm

When the content integrity protection method is "Sign All" or "Unified Merkle Tree" this option MUST be defined. The code for this option is 5. The 8-bit value of this option is one of the Domain Name System Security (DNSSEC) Algorithm Numbers [IANADNSSECALGNUM]. The RSASHA1 [RFC4034], ECDSAP256SHA256 and ECDSAP384SHA384 [RFC6605] algorithms are MANDATORY to implement. Default is ECDSAP256SHA256.

```

0                               1
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5
+-----+
|0 0 0 0 0 1 0 1|   LSA (8)   |
+-----+
```

7.8. Chunk Addressing Method

A peer MUST include the chunk addressing method it uses. The code for this option is 6. Defined values are listed in Table 6.

Method	Description
0	32-bit bins
1	64-bit byte ranges
2	32-bit chunk ranges
3	64-bit bins
4	64-bit chunk ranges
5-255	Unassigned

Table 6: PPSP Peer Chunk Addressing Methods

Implementations MUST support "32-bit chunk ranges" and "64-bit chunk ranges". Default is "32-bit chunk ranges".

```

0                                     1
0 1 2 3 4 5 6 7 8 9 0 1 2
+---+---+---+---+---+---+---+---+
|0 0 0 0 0 1 1 0| CAM (8) |
+---+---+---+---+---+---+---+---+
```

7.9. Live Discard Window

A peer in a live swarm MUST include the discard window it uses. The code for this option is 7. The unit of the discard window depends on the chunk addressing method used. For bins and chunk ranges it is a number of chunks, for byte ranges it is a number of bytes. Its data type is the same as for a bin, or one value in a range specification. In other words, its value is a 32-bit or 64-bit integer in big endian format. If this option is used, the Chunk Addressing Method MUST appear before it in the list. This option has the following structure:

```

0                                     1                                     2                                     3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|0 0 0 0 0 1 1 1| Live Discard Window (32 or 64) ~
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
~
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
```

A peer that does not, under normal circumstances, discard chunks MUST set this option to the special value 0xFFFFFFFF (32-bit) or 0xFFFFFFFFFFFFFFFF (64-bit). For example, peers that record a complete broadcast to offer it directly as a static asset after the broadcast ends use these values (see Section 6.1.2). Section 6.2 explains how to determine a value for this option.

7.10. Supported Messages

Peers may support just a subset of the PPSPP messages. For example, peers running over TCP may not accept ACK messages, or peers used with a centralized tracking infrastructure may not accept PEX messages. For these reasons, peers who support only a proper subset of the PPSPP messages MUST signal which subset they support by means of this protocol option. The code for this option is 8. The value of this option is a length octet (SupMsgLen) indicating the length in bytes of the compressed bitmap that follows.

The set of messages supported can be derived from the compressed bitmap by padding it with bytes of value 0 until it is 256 bits in length. Then a 1 bit in the resulting bitmap at position X (numbering left to right) corresponds to support for message type X, see Table 7. In other words, to construct the compressed bitmap, create a bitmap with a 1 for each message type supported and a 0 for a message type that is not, store it as an array of bytes and truncate it to the last non-zero byte.

```

      0               1               2               3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
|0 0 0 0 0 0 1 0| SupMsgLen (8) | ~
+-----+-----+-----+-----+-----+-----+-----+-----+
~                               Supported Messages Bitmap (variable, max 256) ~
+-----+-----+-----+-----+-----+-----+-----+-----+

```

8. UDP Encapsulation

PPSPP implementations MUST use UDP as transport protocol and MUST use LEDBAT for congestion control [RFC6817]. Using LEDBAT enables PPSPP to serve the content after playback (seeding) without disrupting the user who may have moved to different tasks that use its network connection. Future PPSPP versions can also run over other transport protocols, or use different congestion control algorithms.

8.1. Chunk Size

In general, an UDP datagram containing PPSPP messages SHOULD fit inside a single IP packet, so its maximum size depends on the MTU of the network. If the UDP datagram does not fit, its chance of getting lost in the network increases as the loss of a single fragment of the datagram causes the loss of the complete datagram.

The largest message in a PPSPP datagram is the DATA message carrying a chunk of content. So the (maximum) size of a chunk to choose for a

particular swarm depends primarily on the expected MTU. The chunk size should be chosen such that a chunk and its required INTEGRITY messages can generally be carried inside a single datagram, following the Atomic Datagram Principle (Section 5.3). Other considerations are the hardware capabilities of the peers. Having large chunks and therefore less chunks per mebibyte of content reduces processing costs. The chunk addressing schemes can all work with different chunk sizes, see Section 4.

The RECOMMENDED approach is to use fixed-sized chunks of 1024 bytes, as this size has a high likelihood of travelling end-to-end across the Internet without any fragmentation. In particular, with this size a UDP datagram with a DATA message can be transmitted as a single IP packet over an Ethernet network with 1500-byte frames.

A PPSP implementation MAY use a variant of the Packetization Layer Path MTU Discovery (PLPMTUD), described in [RFC4821], for discovering the optimal MTU between sender and destination. As in PLPMTUD, progressively larger probing packets are used to detect the optimal MTU among a link. However, in PPSP, probe packets SHOULD contain actual messages, in particular, multiple DATA messages. By using actual DATA messages as probe packets, the returning ACK messages will confirm the probe delivery, effectively updating the MTU estimate on both ends of the link. To be able to scale up probe packets with sensible increments, a minimum chunk size of 512 bytes SHOULD be used. Smaller chunk sizes lead to an inefficient protocol. An implication is that PPSP supports datagrams over IPv4 of 576 bytes or more only. This variant is not mandatory to implement.

The chunk size used for a particular swarm, or that fact that it is variable MUST be part of the swarm's metadata (which then minimally consists of the swarm ID and the chunk nature and size).

8.2. Datagrams and Messages

When using UDP, the abstract datagram described above corresponds directly to a UDP datagram. Most messages within a datagram have a fixed length, which generally depends on the type of the message. The first byte of a message denotes its type. The currently defined types are:

Msg Type	Description
0	HANDSHAKE
1	DATA
2	ACK
3	HAVE
4	INTEGRITY
5	PEX_RESv4
6	PEX_REQ
7	SIGNED_INTEGRITY
8	REQUEST
9	CANCEL
10	CHOKe
11	UNCHOKe
12	PEX_RESv6
13	PEX_REScert
14-254	Unassigned
255	Reserved

Table 7: PPSP Peer Protocol Message Types

Furthermore, integers are serialized in the network (big-endian) byte order. So consider the example of a HAVE message (Section 3.2) using bin chunk addressing. It has message type of 0x03 and a payload of a bin number, a four-byte integer (say, 1); hence, its on the wire representation for UDP can be written in hex as: "0300000001".

All messages are idempotent or recognizable as duplicates. Idempotent means that processing a message more than once does not lead to a different state from if it was processed just once. In particular, a peer MAY resend DATA, ACK, HAVE, INTEGRITY, PEX *, SIGNED_INTEGRITY, REQUEST, CANCEL, CHOKe and UNCHOKe messages without problems when loss is suspected. When a peer resends a HANDSHAKE message it can be recognized as duplicate by the receiver, because it already recorded the first connection attempt, and be dealt with.

8.3. Channels

As described in Section 3.11 PPSP uses a multiplexing scheme, called channels, to allow multiple swarms to use the same UDP port. In the UDP encapsulation, each datagram from peer A to peer B is prefixed with the channel ID allocated by peer B. The peers learn about eachother's channel ID during the handshake as explained in a moment. A channel ID consists of 4 bytes and MUST be generated following the requirements in [RFC4960] (Sec. 5.1.3).

8.4. HANDSHAKE

A channel is established with a handshake. To start a handshake, the initiating peer needs to know:

1. the IP address of a peer
2. peer's UDP port and
3. the swarm's metadata record which consists of:
 - (a) the swarm ID of the content (see Section 5.1 and Section 6),
 - (b) the chunk size used,
 - (c) the chunk addressing method used,
 - (d) the content integrity protection method used, and
 - (e) the Merkle hash tree function used (if applicable).
 - (f) If automatic content size detection (see Section 5.6) is not used, the content length is also part of the metadata record for static content.

A datagram containing a HANDSHAKE message:

```

      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     Destination Channel ID (32)                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+
| 0 0 0 0 0 0 0 0 |                                     Source Channel ID (32)                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     ~                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     Protocol Options                                     |
|                                     ~                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

where:

Destination Channel ID:

If the message is sent by the initiating peer than it MUST be an all 0-zeros channel ID.

If the message sent by the receiving peer than it MUST consist of the Source Channel ID from the sender's HANDSHAKE message

The octet 0x00: The HANDSHAKE message: 0x00

The Source Channel ID: A locally unused channel ID

Protocol Options: A list of protocol options encoding the swarm's metadata, as defined in Section 7.

A peer SHOULD explicitly close a channel by sending a HANDSHAKE message that MUST contain an all 0-zeros channel ID and a list of protocol options. The list MUST be either empty or contain the maximum version number the sender supports, following the Min/max versioning scheme defined in [RFC6709], Section 4.1.

8.5. HAVE

A HAVE message (type 0x03) consists of a single chunk specification that states that the sending peer has those chunks and successfully checked their integrity. The single chunk specification represents a consecutive range of verified chunks. A bin consists of a single integer, and a chunk or byte range of two integers, of the width specified by the Chunk Addressing protocol options, encoded big endian.

A HAVE message using 32-bit chunk ranges as Chunk Addressing method:

```

      0               1               2               3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
| 0 0 0 0 0 0 1 1 |                               Start chunk (32) |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                   |                               End chunk (32)   |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                   |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

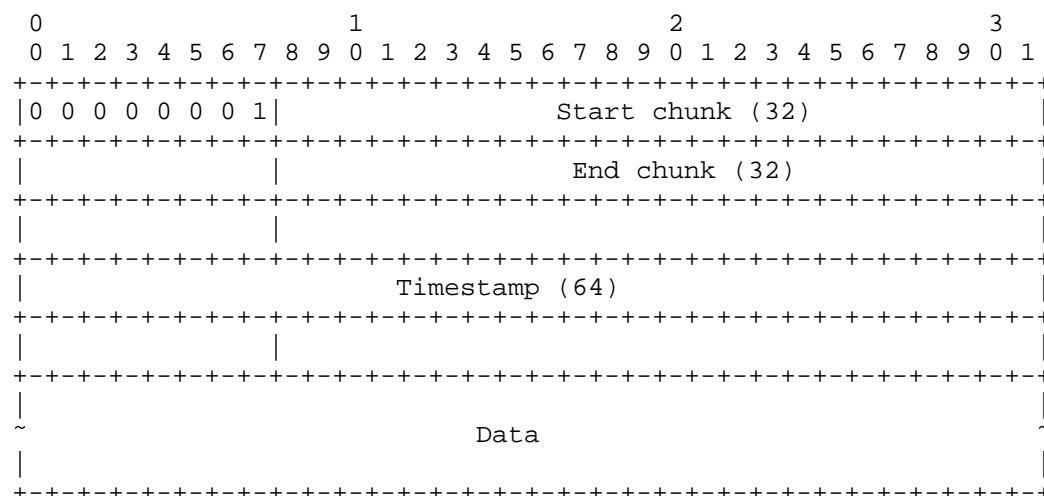
where the first octet is the HAVE message (0x03), followed by the start chunk and the end chunk describing the chunk range.

(received and checked first four kilobytes of a file/stream)

8.6. DATA

A DATA message (type 0x01) consists of a chunk specification, a timestamp and the actual chunk. In case a datagram contains one DATA message, a sender MUST always put the DATA message in the tail of the datagram. A datagram MAY contain multiple DATA messages when the chunk size is fixed and when none of DATA messages carry the last chunk if that is smaller than the chunk size. As the LEDBAT congestion control is used, a sender MUST include a timestamp, in particular, a 64-bit integer representing the current system time with microsecond accuracy. The timestamp MUST be included between chunk specification and the actual chunk.

A DATA message using 32-bit chunk ranges as Chunk Addressing method:




where the first octet is the DATA message (0x01), followed by the start chunk and the end chunk describing the chunk range, the timestamp and the actual data.

8.7. ACK

An ACK message (type 0x02) acknowledges data that was received from its addressee; to comply with the LEDBAT delay-based congestion control an ACK message consists of a chunk specification and a timestamp representing an one-way delay sample. The one-way delay sample is a 64-bit integer with microsecond accuracy, and is computed from the timestamp received from the previous DATA message containing the chunk being acknowledged following the LEDBAT specification.

An ACK message using 32-bit chunk ranges as Chunk Addressing method:




where the first octet is the ACK message (0x02), followed by the start chunk and the end chunk describing the chunk range, and the one-way delay sample.

8.8. INTEGRITY

An INTEGRITY message (type 0x04) consists of a chunk specification and the cryptographic hash for the specified chunk or node. The type and format of the hash depends on the protocol options.

An INTEGRITY message using 32-bit chunk ranges as Chunk Addressing method and SHA1 hash:




where the first octet is the ACK message (0x02), followed by the start chunk and the end chunk describing the chunk range, and the one-way delay sample.

8.9. SIGNED_INTEGRITY

A SIGNED_INTEGRITY message (type 0x07) consists of a chunk specification, a 64-bit NTP timestamp [RFC5905] and a digital signature encoded as a Signature field in a RRSIG record in DNSSEC without the BASE-64 encoding [RFC4034]. The signature algorithm is defined by the Live Signature Algorithm protocol option, see Section 7.7. The plaintext over which the signature is taken depends on the content integrity protection method used, see Section 6.1.

A SIGNED_INTEGRITY message using 32-bit chunk ranges as Chunk Addressing method:



where the first octet is the SIGNED_INTEGRITY message (0x07), followed by the start chunk and the end chunk describing the chunk range, the timestamp, and the Signature.

The length of the digital signature can be derived from the Live Signature Algorithm protocol option and the swarm ID as follows. The first MANDATORY algorithm is RSASHA1. In that case, the swarm ID consists of a 1-byte Algorithm field followed by a RSA public key stored as a tuple (exponent length,exponent,modulus) [RFC3110]. Given the exponent length and the length of the public key tuple in the swarm ID, the length of the modulus in bytes can be calculated. This yields the length of the signature as in RSA this is the length of the modulus [HAC01]. The other MANDATORY algorithms are

ECDSAP256SHA256 and ECDSAP384SHA384 [RFC6605]. For these algorithms the length of the digital signature is 64 and 96 bytes, respectively.

8.10. REQUEST

A REQUEST message (type 0x08) consists of a chunk specification for the chunks the requester wants to download.

A REQUEST message using 32-bit chunk ranges as Chunk Addressing method:

```

      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
| 0 0 0 0 1 0 0 0 |                               Start chunk (32) |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                   |                               End chunk (32)  |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                   |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

where the first octet is the REQUEST message (0x08), followed by the start chunk and the end chunk describing the chunk range.

8.11. CANCEL

A CANCEL message (type 0x09) consists of a chunk specification for the chunks the requester no longer is interested in.

A CANCEL message using 32-bit chunk ranges as Chunk Addressing method:

```

      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
| 0 0 0 0 1 0 0 1 |                               Start chunk (32) |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                   |                               End chunk (32)  |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                   |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

where the first octet is the CANCEL message (0x09), followed by the start chunk and the end chunk describing the chunk range.

8.12. CHOKe and UNCHOKe

Both CHOKe and UNCHOKe messages (types 0x0a and 0x0b, respectively) carry no payload.

A CHOKe message:

```

0
0 1 2 3 4 5 6 7
+---+---+---+---+
|0 0 0 0 1 0 1 0|
+---+---+---+---+
```

where the first octet is the CHOKe message (0x0a).

An UNCHOKe message:

```

0
0 1 2 3 4 5 6 7
+---+---+---+---+
|0 0 0 0 1 0 1 1|
+---+---+---+---+
```

where the first octet is the UNCHOKe message (0x0b).

8.13. PEX_REQ, PEX_RESv4, PEX_RESv6 and PEX_REScert

A PEX_REQ (0x06) message has no payload. A PEX_RES (0x05) message consists of an IPv4 address in big endian format followed by a UDP port number in big endian format. A PEX_RESv6 (0x0c) message contains a 128-bit IPv6 address instead of an IPv4 one. If a PEX_REQ message does not originate from a private or link-local address [RFC1918][RFC4291], then the PEX_RES* messages sent in reply MUST NOT contain such addresses. This is to prevent leaking of internal addresses to external peers.

A PEX_REQ message:

```

0
0 1 2 3 4 5 6 7
+---+---+---+---+
|0 0 0 0 0 1 1 0|
+---+---+---+---+
```

where the first octet is the PEX_REQ message (0x06).

A PEX_RES message:

```

      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|0 0 0 0 0 1 0 1|               IPv4 Address (32)               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               |               Port (16)               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

where the first octet is the PEX_RES message (0x05), followed by the IPv4 address and the port number.

A PEX_RESv6 message:

```


      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|0 0 0 0 1 1 0 0|               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               |               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               |               IPv6 Address (128)               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               |               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               |               Port (16)               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

where the first octet is the PEX_RESv6 message (0x0c), followed by the IPv6 address and the port number.

A PEX_REScert (0x0d) message consists of a 16-bit integer in big endian specifying the size of the membership certificate that follows, see Section 13.2.1. This membership certificate states that peer P at time T is a member of swarm S and is a X.509v3 certificate [RFC5280] that is encoded using the ASN.1 distinguished encoding rules (DER) [CCITT.X208.1988]. The certificate MUST contain a "Subject Alternative Name" extension, marked as critical, of type uniformResourceIdentifier.

A PEX_REScert message:



where the first octet is the PEX_REScert message (0x0d), followed by the size of the membership certificate, and the membership certificate.


The URL contained in the name extension MUST follow the generic syntax for URLs [RFC3986], where its scheme component is "ppsp", the host in the authority component is the DNS name or IP address of peer P, the port in the authority component is the port of peer P, and the path contains the swarm identifier for swarm S, in hexadecimal form. In particular, the preferred form of the swarm identifier is `xxxyzz...`, where the 'x's, 'y's and 'z's are 2 hexadecimal digits of the 8-bit pieces of the identifier. The validity time of the certificate is set with `notBefore` UTCTime set to T and `notAfter` UTCTime set to T plus some expiry time defined by the issuer. An example URL:

`ppsp://192.168.0.1:6778/e5a12c7ad2d8fab33c699d1e198d66f79fa610c3`

8.14. KEEPALIVE

Keepalives do not have a message type on UDP. They are just simple datagrams consisting of the 4-byte channel ID of the destination only.

A keepalive datagram:



8.15. Detecting a Dead Peer

A guideline for declaring a peer dead consist of a 3 minute delay since that last packet has been received from that peer, and at least 3 datagrams were sent to that peer during the same period.

8.16. Flow and Congestion Control

Explicit flow control is not necessary in PPSPP-over-UDP. In the case of video-on-demand the receiver will request data explicitly from peers and is therefore in control of how much data is coming towards it. In the case of live streaming, where a push-model may be used, the amount of data incoming is limited to the bitrate, which the receiver must be able to process otherwise it cannot play the stream. Should, for any reason, the receiver get saturated with data that situation is perfectly detected by the congestion control.

PPSPP-over-UDP can support different congestion control algorithms. At present, it uses the LEDBAT congestion control algorithm [RFC6817]. LEDBAT is a delay-based congestion control algorithm that is used by millions of users everyday as part of the uTP transmission protocol of BitTorrent [LBT],[LCOMPL] and is suitable for P2P streaming [PPSPPERF].

8.17. Example of Operation

We present a small example of communication between a leecher and a seeder. The example presents the transmission of the file "Hello World!", which fits within a 1024 bytes chunk. For an easy understanding we use the message description names, as listed in Table 7, and the protocol option names as listed in Table 2, rather than the actual binary value.

To do the handshake the initiating peer sends a datagram that MUST start with an all 0-zeros channel ID (0x00000000), followed by a HANDSHAKE message, whose payload is a locally unused channel ID (0x00000001) and a list of protocol options.


```

      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
|0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0|
+-----+-----+-----+-----+-----+-----+-----+-----+
|  HANDSHAKE  |0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0|
+-----+-----+-----+-----+-----+-----+-----+-----+
|0 0 0 0 0 0 0 1|   Version   |0 0 0 0 0 0 0 1|  Min Version  |
+-----+-----+-----+-----+-----+-----+-----+-----+
|0 0 0 0 0 0 0 1|   Swarm ID   |0 0 0 0 0 0 0 0 0 0 0 0 1 0 1 0 0|
+-----+-----+-----+-----+-----+-----+-----+-----+
|0 1 0 0 0 1 1 1 1 0 1 0 0 0 0 0 0 0 0 1 0 0 1 1 1 1 1 0 0 1 1 0|
~.....~
|1 0 0 0 0 1 1 0 1 0 1 0 1 0 1 0 1 1 0 0 0 0 0 0 1 1 1 0 1 1|
+-----+-----+-----+-----+-----+-----+-----+-----+
|  Cont. Int.  |0 0 0 0 0 0 0 1| Mer.H.Tree F. |0 0 0 0 0 0 0 0|
+-----+-----+-----+-----+-----+-----+-----+-----+
|  Chunk Add.  |0 0 0 0 0 0 1 0|      End      |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

The protocol options are:

Version: 1

Minimum supported Version: 1

Swarm Identifier: A 20 bytes root hash (47a0...b03b) identifying the content.

Content Integrity Protection Method: Merkle Hash Tree.

Merkle Tree Hash Function: SHA1.

Chunk Addressing Method: 32-bit chunk ranges.

The receiving peer MAY respond, in which case the returned datagram MUST consist of the channel ID from the sender's HANDSHAKE message (0x000000001), a HANDSHAKE message, whose payload is a locally unused channel ID (0x000000008) and a list of protocol options, followed by any other messages it wants to send.

```

      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
|0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 1|
+-----+-----+-----+-----+-----+-----+-----+-----+
|  HANDSHAKE  |0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0|
+-----+-----+-----+-----+-----+-----+-----+-----+
|0 0 0 0 1 0 0 0|   Version   |0 0 0 0 0 0 0 0 1|   Cont. Int.  |
+-----+-----+-----+-----+-----+-----+-----+-----+
|0 0 0 0 0 0 0 1| Mer.H.Tree F. |0 0 0 0 0 0 0 0 0|   Chunk Add.  |
+-----+-----+-----+-----+-----+-----+-----+-----+
|0 0 0 0 0 0 1 0|       End      |       HAVE      |0 0 0 0 0 0 0 0 0|
+-----+-----+-----+-----+-----+-----+-----+-----+
|0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0|
+-----+-----+-----+-----+-----+-----+-----+-----+
|0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0|
+-----+-----+-----+-----+-----+-----+-----+-----+

```

With the protocol options the receiving peer agrees on speaking protocol version 1, on using the Merkle Hash Tree as Content Integrity Protection Method, SHA1 hash as Merkle Tree Hash Function, and 32-bit chunk ranges as Chunk Addressing Method. Furthermore, it sends a HAVE message within the same datagram, announcing that it has locally available the first chunk of content.

At this point, the initiator knows that the peer really responds; for that purpose channel IDs MUST be random enough to prevent easy guessing. So, the third datagram of a handshake MAY already contain some heavy payload. To minimize the number of initialization round trips, the first two datagrams MAY also contain some minor payload, e.g. the HAVE message.

The initiating peer MAY send a request for the chunks of content it wants to retrieve from the receiving peer, e.g. the first chunk announced during the handshake. It always precedes the message with the channel ID of the peer it is communicating with (e.g. 0x000000008 in our example), as described in Section 3.11. Furthermore, it MAY add additional messages such as a PEX_REQ.

0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9
+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-
	0	0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	1
+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-
	REQUEST									0	0	0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	
+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-
	0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	
+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-
	0	0	0	0	0	0	0	0		PEX_REQ																													
+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-

When receiving the third datagram, both peers have the proof they really talk to each other; the three-way handshake is complete. The receiving peer responds to the request by sending a DATA message containing the requested content.

0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9
+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-
	0	0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0
+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-
+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-
	0	0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0
+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-
	0	0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0
+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-
	0	0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	0	0		0	0	0	0	0	0	0	1	0		0	0	1	1	1	0	1	0	0
+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-
	0	1	0	0	0	0	0	1	1		0	0	0	0	0	0	1	0	1		1	0	1	1	1	1	1	1	0		1	1	0	1	1	0	1	1	
+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-
	0	1	0	0	0	1	0	0	0		0	1	0	0	1	0	0	0	0		1	1	0	0	1	0	1	0	1		1	0	1	1	0	1	1	0	0
+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-
~										~										~										~									
	0	1	1	0	1	1	0	0	0		1	0	0	1	0	0	0	0	0		1	0	0	0	0	0	1	0	0		0	0	0	0	0	0	1	0	1
+	-																																						

The DATA message consists of:

The 32-bit chunk range: 0,0 (the first chunk).

The timestamp value: 0004e94180b7db44

The Data message: 48656c6c6f20776f7226c6421 (the "Hello world!" file)

Note that the above datagram does not include the INTEGRITY message, as the entire content can fit into a single message, hence the

initiating peer is able to verify it against the root hash. Also, in this example the peer does not respond to the PEX_REQ as it does not know any third peer participating in the swarm.

Upon receiving the requested data, the initiating peer responds with an acknowledgement message for the first chunk, containing a one way delay sample (100ms). Furthermore it also adds a HAVE message for the chunk.

```

      0              1              2              3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
|0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0|
+-----+-----+-----+-----+-----+-----+-----+-----+
|      ACK      |0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0|
+-----+-----+-----+-----+-----+-----+-----+-----+
|0 0 0 0 0 0 0 0 0|0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0|
+-----+-----+-----+-----+-----+-----+-----+-----+
|0 0 0 0 0 0 0 0 0|0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0|
+-----+-----+-----+-----+-----+-----+-----+-----+
|0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0|
+-----+-----+-----+-----+-----+-----+-----+-----+
|0 1 1 0 0 1 0 0|      HAVE      |0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0|
+-----+-----+-----+-----+-----+-----+-----+-----+
|0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0|
+-----+-----+-----+-----+-----+-----+-----+-----+
|0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0|
+-----+-----+-----+-----+-----+-----+-----+-----+

```

At this point the initiating peer has successfully retrieved the entire file. It then explicitly closes the connection by sending a HANDSHAKE message that contains an all 0-zeros channel ID.

```

      0              1              2              3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
|0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0|
+-----+-----+-----+-----+-----+-----+-----+-----+
|  HANDSHAKE  |0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0|
+-----+-----+-----+-----+-----+-----+-----+-----+
|0 0 0 0 0 0 0 0 0|      End      |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

9. Extensibility

9.1. Chunk Picking Algorithms

Chunk (or piece) picking entirely depends on the receiving peer. The sender peer is made aware of preferred chunks by the means of REQUEST messages. In some (live) scenarios it may be beneficial to allow the sender to ignore those hints and send unrequested data.

The chunk picking algorithm is external to the PPSPP protocol and will generally be a pluggable policy that uses the mechanisms provided by PPSPP. The algorithm will handle the choices made by the user consuming the content, such as seeking, switching audio tracks or subtitles. Example policies for P2P streaming can be found in [BITOS], and [EPLIVEPERF].

9.2. Reciprocity Algorithms

The role of reciprocity algorithms in peer-to-peer systems is to promote client contribution and prevent freeriding. A peer is said to be freeriding if it only downloads content but never uploads to others. Examples of reciprocity algorithms are tit-for-tat as used in BitTorrent [TIT4TAT] and Give-to-Get [GIVE2GET]. In PPSPP, reciprocity enforcement is the sole responsibility of the sender peer.

10. Acknowledgements

Arno Bakker, Riccardo Petrocco and Victor Grishchenko are partially supported by the P2P-Next project (<http://www.p2p-next.org/>), a research project supported by the European Community under its 7th Framework Programme (grant agreement no. 216217). The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the P2P-Next project or the European Commission.

The PPSPP protocol was designed by Victor Grishchenko at Technische Universiteit Delft. The authors would like to thank the following people for their contributions to this draft: the chairs (Martin Stiernerling, Yunfei Zhang, Stefano Previdi, Ning Zong) and members of the IETF PPSP working group, and Mihai Capota, Raul Jimenez, Flutra Osmani, Johan Pouwelse, and Raynor Vliedendhart.

11. IANA Considerations

IANA is to create the new registries defined below for the extensibility of the protocol. For all registries, assignments consist of a name and its associated value. Also for all registries, the "Unassigned" ranges designated are governed by the policy 'IETF Review' as described in [RFC5226].

11.1. PPSP Peer Protocol Message Type Registry

Registry name is "PPSP Peer Protocol Message Type Registry". Values are integers in the range 0-255, with initial assignments and reservations given in Table 7.

11.2. PPSP Peer Protocol Option Registry

Registry name is "PPSP Peer Protocol Option Registry". Values are integers in the range 0-255, with initial assignments and reservations given in Table 2.

11.3. PPSP Peer Protocol Version Number Registry

Registry name is "PPSP Peer Protocol Version Number Registry". Values are integers in the range 0-255, with initial assignments and reservations given in Table 3.

11.4. PPSP Peer Protocol Content Integrity Protection Method Registry

Registry name is "PPSP Peer Protocol Content Integrity Protection Method Registry". Values are integers in the range 0-255, with initial assignments and reservations given in Table 4.

11.5. PPSP Peer Protocol Merkle Hash Tree Function Registry

Registry name is "PPSP Peer Protocol Merkle Hash Tree Function Registry". Values are integers in the range 0-255, with initial assignments and reservations given in Table 5.

11.6. PPSP Peer Protocol Chunk Addressing Method Registry

Registry name is "PPSP Peer Protocol Chunk Addressing Method Registry". Values are integers in the range 0-255, with initial assignments and reservations given in Table 6.

12. Manageability Considerations

This section presents operations and management considerations

following the checklist in [RFC5706], Appendix A.

In this section "PPSPP client" is defined as a PPSPP peer acting on behalf of an end user which may not yet have a copy of the content, and "PPSPP server" as a PPSPP peer that provides the initial copies of the content to the swarm on behalf of a content provider.

12.1. Operations

12.1.1. Installation and Initial Setup

A content provider wishing to use PPSPP to distribute content should setup at least one PPSPP server. PPSPP servers need to have access to either some static content or to some live audio/video sources. To provide flexibility for implementors, this configuration process is not standardized. The output of this process will be a list of metadata records, one for each swarm. A metadata record consists of the swarm ID, the chunk size used, the chunk addressing method used, the content integrity protection method used, and the Merkle hash tree function used (if applicable). If automatic content size detection (see Section 5.6) is not used, the content length is also part of the metadata record for static content. Note the swarm ID already contains the Live Signature Algorithm used, in case of a live stream.

In addition, a content provider should setup a tracking facility for the content by configuring, for example, a PPSP tracker [I-D.ietf-ppsp-base-tracker-protocol] or a Distributed Hash Table. The output of the latter process is a list of transport addresses for the tracking facility.

The list of metadata records of available content, and transport address for the tracking facility, can be distributed to users in various ways. Typically, they will be published on a Web site as links. When a user clicks such a link the PPSPP client is launched, either as a standalone application or by invoking the browser's internal PPSPP protocol handler, as exemplified in Section 2. The clients use the tracking facility to obtain the transport address of the PPSPP server(s) and other peers from the swarm, executing the peer protocol to retrieve and redistribute the content. The format of the PPSPP URLs should be defined in an extension document. The default protocol options should be exploited to keep the URLs small.

The minimal information a tracking facility must return when queried for a list of peers for a swarm is as follows. Assuming the communication between tracking facility and requestor is protected, the facility must at least return for each peer in the list its IP address, transport protocol identifier (i.e., UDP), and transport

protocol port number.

12.1.2. Requirements on Other Protocols and Functional Components

When using the PPSP tracker protocol, PPSPP requires a specific behavior from this protocol for security reasons, as detailed in Section 13.2.

12.1.3. Migration Path

This document does not detail a migration path since there is no previous standard protocol providing similar functionality.

12.1.4. Impact on Network Operation

PPSPP is a peer-to-peer protocol that takes advantage of the fact that content is available from multiple sources to improve robustness, scalability and performance. At the same time, poor choices in determining which exact sources to use can lead to bad experience for the end user and high costs for network operators. Hence, PPSPP can benefit from the ALTO protocol to steer peer selection, as described in Section 3.10.1.

12.1.5. Verifying Correct Operation

PPSPP is operating correctly when all peers obtain the desired content on time. Therefore the PPSPP client is the ideal location to verify the protocol's correct operation. However, it is not feasible to mandate logging the behavior of PPSPP peers in all implementations and deployments, for example, due to privacy reasons. There are two alternative options:

- o Monitoring the PPSPP servers initially providing the content, using standard metrics such as bandwidth usage, peer connections and activity, can help identify trouble, see next section and [RFC2564].
- o The PPSP tracker protocol may be used to gather information about all peers in a swarm, to obtain a global view of operation, according to [I-D.ietf-ppsp-problem-statement] (requirement PPSP.TP.REQ-3).

Basic operation of the protocol can be easily verified when a tracker and swarm ID are known by starting a PPSPP download. Deep packet inspection for DATA and ACK messages help to establish that actual content transfer is happening and that the chunk availability signaling and integrity checking are working.

12.1.6. Configuration

Table 8 shows the PPSPP parameters, their defaults and where the parameter is defined. For parameters that have no default, the table row contains the word "var" and refers to the section discussing the considerations to make when choosing a value.

Name	Default	Definition
Chunk Size	var, 1024 bytes recommended	Section 8.1
Static Content Integrity Protection Method	1 (Merkle Hash Tree)	Section 7.5
Live Content Integrity Protection Method	3 (Unified Merkle Tree)	Section 7.5
Merkle Hash Tree Function	0 (SHA1)	Section 7.6
Live Signature Algorithm	13 (ECDSAP256SHA256)	Section 7.7
Chunk Addressing Method	2 (32-bit chunk ranges)	Section 7.8
Live Discard Window	var	Section 6.2, Section 7.9
NCHUNKS_PER_SIG	var	Section 6.1.2.1
Dead peer detection	No reply in 3 minutes + 3 datagrams	Section 8.15

Table 8: PPSPP Defaults

12.2. Management Considerations

The management considerations for PPSPP are very similar to other protocols that are used for large-scale content distribution, in particular HTTP. How does one manage large numbers of servers? How does one push new content out to a server farm and allows staged releases? How to detect faults and how to measure servers and end-user performance? As standard solutions to these challenges are still being developed, this section cannot provide a definitive recommendation on how PPSPP should be managed. Hence, it describes the standard solutions available at this time, and assumes a future extension document will provide more complete guidelines.

12.2.1. Management Interoperability and Information

As just stated, PPSPP servers providing initial copies of the content are akin to WWW and FTP servers. They can also be deployed in large numbers and thus can benefit from standard management facilities. PPSPP servers may therefore implement an SNMP management interface based on the APPLICATION-MIB [RFC2564], where the file object can be used to report on swarms.

What is missing is the ability to remove or rate limit specific PPSPP swarms on a server. This corresponds to removing or limit specific virtual servers on a Web server. In other words, as multiple pieces of content (swarms, virtual WWW servers) are multiplexed onto a single server process, more fine-grained management of that process is required. This functionality is currently missing.

Logging is an important functionality for PPSPP servers and, depending on the deployment, PPSPP clients. Logging should be done via syslog [RFC5424].

12.2.2. Fault Management

The facilities for verifying correct operation and server management (just discussed) appear sufficient for PPSPP fault monitoring. This can be supplemented with host resource [RFC2790] and UDP/IP network monitoring [RFC4113], as PPSPP server failures can generally be attributed directly to conditions on the host or network.

Since PPSPP has been designed to work in a hostile environment, many benign faults will be handled by the mechanisms used for managing attacks. For example, when a malfunctioning peer starts sending the wrong chunks, this is detected by the content integrity protection mechanism and another source is sought.

12.2.3. Configuration Management

Large-scale deployments may benefit from a standard way of replicating a new piece of content on a set of initial PPSPP servers. This functionality may need to include controlled releasing, such that content becomes available only at a specific point in time (e.g. the release of a movie trailer). This functionality could be provided via NETCONF [RFC6241], to enable atomic configuration updates over a set of servers. Uploading the new content could be one configuration change, making the content available for download by the public another.

12.2.4. Accounting Management

Content providers may offer PPSPP hosting for different customers and will want to bill these customers, for example, based on bandwidth usage. This situation is a common accounting scenario, similar to billing per virtual server for Web servers. PPSPP can therefore benefit from general standardization efforts in this area [RFC2975] when they come to fruition.

12.2.5. Performance Management

Depending on the deployment scenarios, the application performance measurement facilities of [RFC3729] and associated [RFC4150] can be used with PPSPP.

In addition, when the PPSPP tracker protocol is used, it provides a built-in, application-level, performance measurement infrastructure for different metrics. See [I-D.ietf-ppsp-problem-statement] (requirement PPSP.TP.REQ-3).

12.2.6. Security Management

Malicious peers should ideally be locked out long-term. This is primarily for performance reasons, as the protocol is robust against attacks (see next section). Section 13.7 describes a procedure for long-term exclusion.

13. Security Considerations

As any other network protocol, the PPSPP faces a common set of security challenges. An implementation must consider the possibility of buffer overruns, DoS attacks and manipulation (i.e. reflection attacks). Any guarantee of privacy seems unlikely, as the user is exposing its IP address to the peers. A probable exception is the case of the user being hidden behind a public NAT or proxy. This section discusses the protocol's security considerations in detail.

13.1. Security of the Handshake Procedure

Borrowing from the analysis in [RFC5971], the PPSPP peer protocol may be attacked with 3 types of denial-of-service attacks:

1. DOS amplification attack: attackers try to use a PPSPP peer to generate more traffic to a victim.
2. DOS flood attack: attackers try to deny service to other peers by allocating lots of state at a PPSPP peer.

3. Disrupt service to an individual peer: attackers send bogus e.g. REQUEST and HAVE messages appearing to come from victim peer A to the peers B1..Bn serving that peer. This causes A to receive chunks it did not request or to not receive the chunks it requested.

The basic scheme to protect against these attacks is the use of a secure handshake procedure. In the UDP encapsulation the handshake procedure is secured by the use of randomly chosen channel IDs as follows. The channel IDs must be generated following the requirements in [RFC4960] (Sec. 5.1.3).

When UDP is used, all datagrams carrying PPSPP messages are prefixed with a 4-byte channel ID. These channel IDs are random numbers, established during the handshake phase as follows. Peer A initiates an exchange with peer B by sending a datagram containing a HANDSHAKE message prefixed with the channel ID consisting of all 0s. Peer A's HANDSHAKE contains a randomly chosen channel ID, chanA:

A->B: chan0 + HANDSHAKE(chanA) + ...

When peer B receives this datagram, it creates some state for peer A, that at least contains the channel ID chanA. Next, peer B sends a response to A, consisting of a datagram containing a HANDSHAKE message prefixed with the chanA channel ID. Peer B's HANDSHAKE contains a randomly chosen channel ID, chanB.

B->A: chanA + HANDSHAKE(chanB) + ...

Peer A now knows that peer B really responds, as it echoed chanA. So the next datagram that A sends may already contain heavy payload, i.e., a chunk. This next datagram to B will be prefixed with the chanB channel ID. When B receives this datagram, both peers have the proof they are really talking to each other, the three-way handshake is complete. In other words, the randomly chosen channel IDs act as tags (cf. [RFC4960] (Sec. 5.1)).

A->B: chanB + HAVE + DATA + ...

13.1.1. Protection Against Attack 1

In short, PPSPP does a so-called return routability check before heavy payload is sent. This means that attack 1 is fended off: PPSPP does not send back much more data than it received, unless it knows it is talking to a live peer. Attackers sending a spoofed HANDSHAKE to B pretending to be A now need to intercept the message from B to A to get B to send heavy payload, and ensure that that heavy payload goes to the victim, something assumed too hard to be a practical

attack.

Note the rule is that no heavy payload may be sent until the third datagram. This has implications for PPSPP implementations that use chunk addressing schemes that are verbose. If a PPSPP implementation uses large bitmaps to convey chunk availability these may not be sent by peer B in the second datagram.

13.1.2. Protection Against Attack 2

On receiving the first datagram peer B will record some state about peer A. At present this state consists of the chanA channel ID, and the results of processing the other messages in the first datagram. In particular, if A included some HAVE messages, B may add a chunk availability map to A's state. In addition, B may request some chunks from A in the second datagram, and B will maintain state about these outgoing requests.

So presently, PPSPP is somewhat vulnerable to attack 2. An attacker could send many datagrams with HANDSHAKES and HAVEs and thus allocate state at the PPSPP peer. Therefore peer A **MUST** respond immediately to the second datagram, if it is still interested in peer B.

The reason for using this slightly vulnerable three-way handshake instead of the safer handshake procedure of SCTP [RFC4960] (Sec. 5.1) is quicker response time for the user. In the SCTP procedure, peer A and B cannot request chunks until datagrams 3 and 4 respectively, as opposed to 2 and 1 in the proposed procedure. This means that the user has to wait shorter in PPSPP between starting the video stream and seeing the first images.

13.1.3. Protection Against Attack 3

In general, channel IDs serve to authenticate a peer. Hence, to attack, a malicious peer T would need to be able to eavesdrop on conversations between victim A and a benign peer B to obtain the channel ID B assigned to A, chanB. Furthermore, attacker T would need to be able to spoof e.g. REQUEST and HAVE messages from A to cause B to send heavy DATA messages to A, or prevent B from sending them, respectively.

The capability to eavesdrop is not common, so the protection afforded by channel IDs will be sufficient in most cases. If not, point-to-point encryption of traffic should be used, see below.

13.2. Secure Peer Address Exchange

As described in Section 3.10, a peer A can send Peer-Exchange messages PEX_RES to a peer B, which contain the IP address and port of other peers that are supposedly also in the current swarm. The strength of this mechanism is that it allows decentralized tracking: after an initial bootstrap no central tracker is needed anymore. The vulnerability of this mechanism (and DHTs) is that malicious peers can use it for an Amplification attack.

In particular, a malicious peer T could send PEX_RES messages to well-behaved peer A with addresses of peers B1,B2,...,BN and on receipt, peer A could send a HANDSHAKE to all these peers. So in the worst case, a single datagram results in N datagrams. The actual damage depends on A's behavior. E.g. when A already has sufficient connections it may not connect to the offered ones at all, but if it is a fresh peer it may connect to all directly.

In addition, PEX can be used in Eclipse attacks [ECLIPSE] where malicious peers try to isolate a particular peer such that it only interacts with malicious peers. Let us distinguish two specific attacks:

E1. Malicious peers try to eclipse the single injector in live streaming.

E2. Malicious peers try to eclipse a specific consumer peer.

Attack E1 has the most impact on the system as it would disrupt all peers.

13.2.1. Protection against the Amplification Attack

If peer addresses are relatively stable, strong protection against the attack can be provided by using public key cryptography and certification. In particular, a PEX_REScert message will carry swarm-membership certificates rather than IP address and port. A membership certificate for peer B states that peer B at address (ipB,portB) is part of swarm S at time T and is cryptographically signed. The receiver A can check the cert for a valid signature, the right swarm and liveliness and only then consider contacting B. These swarm-membership certificates correspond to signed node descriptors in secure decentralized peer sampling services [SPS].

Several designs are possible for the security environment for these membership certificates. That is, there are different designs possible for who signs the membership certificates and how public keys are distributed. As an example, we describe a design where the

PPSP tracker acts as certification authority.

13.2.2. Example: Tracker as Certification Authority

A peer A wanting to join swarm S sends a certificate request message to a tracker X for that swarm. Upon receipt, the tracker creates a membership certificate from the request with swarm ID S, a timestamp T and the external IP and port it received the message from, signed with the tracker's private key. This certificate is returned to A.

Peer A then includes this certificate when it sends a PEX_REScert to peer B. Receiver B verifies it against the tracker public key. This tracker public key should be part of the swarm's metadata, which B received from a trusted source. Subsequently, peer B can send the member certificate of A to other peers in PEX_REScert messages.

Peer A can send the certification request when it first contacts the tracker, or at a later time. Furthermore, the responses the tracker sends could contain membership certificates instead of plain addresses, such that they can be gossiped securely as well.

We assume the tracker is protected against attacks and does a return routability check. The latter ensures that malicious peers cannot obtain a certificate for a random host, just for hosts where they can eavesdrop on incoming traffic.

The load generated on the tracker depends on churn and the lifetime of a certificate. Certificates can be fairly long lived, given that the main goal of the membership certificates is to prevent that malicious peer T can cause good peer A to contact *random* hosts. The freshness of the timestamp just adds extra protection in addition to achieving that goal. It protects against malicious hosts causing a good peer A to contact hosts that previously participated in the swarm.

The membership certificate mechanism itself can be used for a kind of amplification attack against good peers. Malicious peer T can cause peer A to spend some CPU to verify the signatures on the membership certificates that T sends. To counter this, A SHOULD check a few of the certificates sent and discard the rest if they are defective.

The same membership certificates described above can be registered in a Distributed Hash Table that has been secured against the well-known DHT specific attacks [SECDHTS].

Note that this scheme does not work for peers behind a symmetric Network Address Translator, but neither does normal tracker registration.

13.2.3. Protection Against Eclipse Attacks

Before we can discuss Eclipse attacks we first need to establish the security properties of the central tracker. A tracker is vulnerable to Amplification attacks too. A malicious peer T could register a victim B with the tracker, and many peers joining the swarm will contact B. Trackers can also be used in Eclipse attacks. If many malicious peers register themselves at the tracker, the percentage of bad peers in the returned address list may become high. Leaving the protection of the tracker to the PPSP tracker protocol specification, we assume for the following discussion that it returns a true random sample of the actual swarm membership (achieved via Sybil attack protection). This means that if 50% of the peers is bad, you'll still get 50% good addresses from the tracker.

Attack E1 on PEX can be fended off by letting live injectors disable PEX. Or at least, let live injectors ensure that part of their connections are to peers whose addresses came from the trusted tracker.

The same measures defend against attack E2 on PEX. They can also be employed dynamically. When the current set of peers B that peer A is connected to doesn't provide good quality of service, A can contact the tracker to find new candidates.

13.3. Support for Closed Swarms (PPSP.SEC.REQ-1)

The Closed Swarms [CLOSED] and Enhanced Closed Swarms [ECS] mechanisms provide swarm-level access control. The basic idea is that a peer cannot download from another peer unless it shows a Proof-of-Access. Enhanced Closed Swarms improve on the original Closed Swarms by adding on-the-wire encryption against man-in-the-middle attacks and more flexible access control rules.

The exact mapping of ECS to PPSP is defined in [I-D.gabrijelcic-ppsp-ecs].

13.4. Confidentiality of Streamed Content (PPSP.SEC.REQ-2+3)

No extra mechanism is needed to support confidentiality in PPSP. A content publisher wishing confidentiality should just distribute content in cyphertext / DRM-ed format. In that case it is assumed a higher layer handles key management out-of-band. Alternatively, pure point-to-point encryption of content and traffic can be provided by the proposed Closed Swarms access control mechanism, or by DTLS [RFC6347] or IPsec [RFC4301].

13.5. Strength of the Hash Function for Merkle Hash Trees

Implementations MUST support SHA1 as the hash function for content integrity protection via Merkle Hash trees. SHA1 is preferred over stronger hash functions for two reasons. First, it reduces on-the-wire overhead. Second, few implementations need the extra strength of other functions because the function is used in a hash tree. In particular, if attackers manage to find a collision for a hash it can replace just one chunk, so the impact is limited. If fixed sized chunks are used, the collision has to be of the same size as the original chunk. For hashes higher up in the hash tree, a collision must be a concatenation of two hashes. In sum, finding collisions that fit with the hash tree are generally harder to find than regular SHA1 collisions, which are, at the time of writing, still hard to find.

13.6. Limit Potential Damage and Resource Exhaustion by Bad or Broken Peers (PPSP.SEC.REQ-4+6)

In this section an analysis is given of the potential damage a malicious peer can do with each message in the protocol, and how it is prevented by the protocol (implementation).

13.6.1. HANDSHAKE

- o Secured against DoS amplification attacks as described in Section 13.1.
- o Threat HS.1: An Eclipse attack where peers T1..Tn fill all connection slots of A by initiating the connection to A.

Solution: Peer A must not let other peers fill all its available connection slots, i.e., A must initiate connections itself too, to prevent isolation.

13.6.2. HAVE

- o Threat HAVE.1: Malicious peer T can claim to have content which it hasn't. Subsequently T won't respond to requests.

Solution: peer A will consider T to be a slow peer and not ask it again.

- o Threat HAVE.2: Malicious peer T can claim not to have content. Hence it won't contribute.

Solution: Peer and chunk selection algorithms external to the protocol will implement fairness and provide sharing incentives.

13.6.3. DATA

- o Threat DATA.1: peer T sending bogus chunks.

Solution: The content integrity protection schemes defend against this.

- o Threat DATA.2: peer T sends peer A unrequested chunks.

To protect against this threat we need network-level DoS prevention.

13.6.4. ACK

- o Threat ACK.1: peer T acknowledges wrong chunks.

Solution: peer A will detect inconsistencies with the data it sent to T.

- o Threat ACK.2: peer T modifies timestamp in ACK to peer A used for time-based congestion control.

Solution: In theory, by decreasing the timestamp peer T could fake there is no congestion when in fact there is, causing A to send more data than it should. [RFC6817] does not list this as a security consideration. Possibly this attack can be detected by the large resulting asymmetry between round-trip time and measured one-way delay.

13.6.5. INTEGRITY and SIGNED_INTEGRITY

- o Threat INTEGRITY.1: An amplification attack where peer T sends bogus INTEGRITY or SIGNED_INTEGRITY messages, causing peer A to check hashes or signatures, thus spending CPU unnecessarily.

Solution: If the hashes/signatures don't check out A will stop asking T because of the atomic datagram principle and the content integrity protection. Subsequent unsolicited traffic from T will be ignored.

- o Threat INTEGRITY.2: An attack where peer T sends old SIGNED_INTEGRITY messages in the Unified Merkle Tree scheme, trying to make peer A tune in at a past point in the live stream.

Solution: The timestamp in the SIGNED_INTEGRITY message protects against such replays. Subsequent traffic from T will be ignored.

13.6.6. REQUEST

- o Threat REQUEST.1: peer T could request lots from A, leaving A without resources for others.

Solution: A limit is imposed on the upload capacity a single peer can consume, for example, by using an upload bandwidth scheduler that takes into account the need of multiple peers. A natural upper limit of this upload quatum is the bitrate of the content, taking into account that this may be variable.

13.6.7. CANCEL

- o Threat CANCEL.1: peer T sends CANCEL messages for content it never requested to peer A.

Solution: peer A will detect the inconsistency of the messages and ignore them. Note that CANCEL messages may be received unexpectedly when a transport is used where REQUEST messages may be lost or reordered with respect to the subsequent CANCELs.

13.6.8. CHOKe

- o Threat CHOKe.1: peer T sends REQUEST messages after peer A sent B a CHOKe message.

Solution: peer A will just discard the unwanted REQUESTs and resend the CHOKe, assuming it got lost.

13.6.9. UNCHOKe

- o Threat UNCHOKe.1: peer T sends an UNCHOKe message to peer A without having sent a CHOKe message before.

Solution: peer A can easily detect this violation of protocol state, and ignore it. Note this can also happen due to loss of a CHOKe message sent by a benign peer.

- o Threat UNCHOKe.2: peer T sends an UNCHOKe message to peer A, but subsequently does not respond to its REQUESTs.

Solution: peer A will consider T to be a slow peer and not ask it again.

13.6.10. PEX_RES

- o Secured against amplification and Eclipse attacks as described in Section 13.2.

13.6.11. Unsolicited Messages in General

- o Threat: peer T could send a spoofed PEX_REQ or REQUEST from peer B to peer A, causing A to send a PEX_RES/DATA to B.

Solution: the message from peer T won't be accepted unless T does a handshake first, in which case the reply goes to T, not victim B.

13.7. Exclude Bad or Broken Peers (PPSP.SEC.REQ-5)

A receiving peer can detect malicious or faulty senders as just described, which it can then subsequently ignore. However, excluding such a bad peer from the system completely is complex. Random monitoring by trusted peers that would blacklist bad peers as described in [DETMAL] is one option. This mechanism does require extra capacity to run such trusted peers, which must be indistinguishable from regular peers, and requires a solution for the timely distribution of this blacklist to peers in a scalable manner.

14. References

14.1. Normative References

[CCITT.X208.1988]

International International Telephone and Telegraph Consultative Committee, "Specification of Abstract Syntax Notation One (ASN.1)", CCITT Recommendation X.208, November 1988.

[FIPS180-3]

Information Technology Laboratory, National Institute of Standards and Technology, "Federal Information Processing Standards: Secure Hash Standard (SHS)", Publication 180-3, Oct 2008.

[IANADNSSECALGNUM]

IANA, "Domain Name System Security (DNSSEC) Algorithm Numbers",
<<http://www.iana.org/assignments/dns-sec-alg-numbers>>.

[RFC1918] Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G., and

- E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3110] Eastlake, D., "RSA/SHA-1 SIGs and RSA KEYS in the Domain Name System (DNS)", RFC 3110, May 2001.
- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, January 2005.
- [RFC4034] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "Resource Records for the DNS Security Extensions", RFC 4034, March 2005.
- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, February 2006.
- [RFC5280] Cooper, D., Santesson, S., Farrell, S., Boeyen, S., Housley, R., and W. Polk, "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile", RFC 5280, May 2008.
- [RFC5905] Mills, D., Martin, J., Burbank, J., and W. Kasch, "Network Time Protocol Version 4: Protocol and Algorithms Specification", RFC 5905, June 2010.
- [RFC6605] Hoffman, P. and W. Wijngaards, "Elliptic Curve Digital Signature Algorithm (DSA) for DNSSEC", RFC 6605, April 2012.

14.2. Informative References

- [ABMRKL] Bakker, A., "Merkle hash torrent extension", BitTorrent Enhancement Proposal 30, Mar 2009, <http://bittorrent.org/beps/bep_0030.html>.
- [BINMAP] Grishchenko, V. and J. Pouwelse, "Binmaps: hybridizing bitmaps and binary trees", Technical Report PDS-2011-005, Parallel and Distributed Systems Group, Fac. of Electrical Engineering, Mathematics, and Computer Science, Delft University of Technology, The Netherlands, Apr 2009.
- [BITOS] Vlavianos, A., Iliofotou, M., Mathieu, F., and M. Faloutsos, "BiToS: Enhancing BitTorrent for Supporting

Streaming Applications", IEEE INFOCOM Global Internet Symposium Barcelona, Spain, Apr 2006.

[BITTORRENT]

Cohen, B., "The BitTorrent Protocol Specification", BitTorrent Enhancement Proposal 3, Feb 2008, <http://bittorrent.org/beps/bep_0003.html>.

[CLOSED]

Borch, N., Mitchell, K., Arntzen, I., and D. Gabrijelcic, "Access Control to BitTorrent Swarms Using Closed Swarms", ACM workshop on Advanced Video Streaming Techniques for Peer-to-Peer Networks and Social Networking (AVSTP2P '10), Florence, Italy, Oct 2010, <<http://doi.acm.org/10.1145/1877891.1877898>>.

[DETMAL]

Shetty, S., Galdames, P., Tavanapong, W., and Ying. Cai, "Detecting Malicious Peers in Overlay Multicast Streaming", IEEE Conference on Local Computer Networks (LCN'06). Tampa, FL, USA, Nov 2006.

[ECLIPSE]

Sit, E. and R. Morris, "Security Considerations for Peer-to-Peer Distributed Hash Tables", IPTPS '01: Revised Papers from the First International Workshop on Peer-to-Peer Systems pp. 261-269, Springer-Verlag, 2002.

[ECS]

Jovanovikj, V., Gabrijelcic, D., and T. Klobucar, "Access Control in BitTorrent P2P Networks Using the Enhanced Closed Swarms Protocol", International Conference on Emerging Security Information, Systems and Technologies (SECURWARE 2011), pp. 97-102, Nice, France, Aug 2011.

[EPLIVEPERF]

Bonald, T., Massoulie, L., Mathieu, F., Perino, D., and A. Twigg, "Epidemic Live Streaming: Optimal Performance Trade-offs", Proceedings of the 2008 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems Annapolis, MD, USA, Jun 2008.

[GIVE2GET]

Mol, J., Pouwelse, J., Meulpolder, M., Epema, D., and H. Sips, "Give-to-Get: Free-riding Resilient Video-on-demand in P2P Systems", Proceedings Multimedia Computing and Networking conference (Proceedings of SPIE Vol. 6818) San Jose, California, USA, Jan 2008.

[HAC01]

Menezes, A., van Oorschot, P., and S. Vanstone, "Handbook of Applied Cryptography", CRC Press, (Fifth Printing,

August 2001), Oct 1996.

- [I-D.gabrijelcic-ppsp-ecs]
Gabrijelcic, D., "Enhanced Closed Swarm protocol",
draft-ppsp-gabrijelcic-ecs (work in progress),
November 2012.
- [I-D.ietf-alto-protocol]
Alimi, R., Penno, R., and Y. Yang, "ALTO Protocol",
draft-ietf-alto-protocol-16 (work in progress), May 2013.
- [I-D.ietf-ppsp-base-tracker-protocol]
Cruz, R., Nunes, M., Yingjie, G., Xia, J., Taveira, J.,
and D. Lingli, "PPSP Tracker Protocol-Base Protocol (PPSP-
TP/1.0)", draft-ietf-ppsp-base-tracker-protocol-00 (work
in progress), February 2013.
- [I-D.ietf-ppsp-problem-statement]
Zhang, Y. and N. Zong, "Problem Statement and Requirements
of Peer-to-Peer Streaming Protocol (PPSP)",
draft-ietf-ppsp-problem-statement-15 (work in progress),
May 2013.
- [JIM11] Jimenez, R., Osmani, F., and B. Knutsson, "Sub-Second
Lookups on a Large-Scale Kademlia-Based Overlay", IEEE
International Conference on Peer-to-Peer
Computing (P2P'11), Kyoto, Japan, Aug 2011.
- [LBT] Rossi, D., Testa, C., Valenti, S., and L. Muscariello,
"LEDBAT: the new BitTorrent congestion control protocol",
Computer Communications and Networks (ICCCN), Zurich,
Switzerland, Aug 2010.
- [LCOMPL] Testa, C. and D. Rossi, "On the impact of uTP on
BitTorrent completion time", IEEE International Conference
on Peer-to-Peer Computing (P2P'11), Kyoto, Japan,
Aug 2011.
- [MERKLE] Merkle, R., "Secrecy, Authentication, and Public Key
Systems", Ph.D. thesis Dept. of Electrical Engineering,
Stanford University, CA, USA, pp 40-45, 1979.
- [POLLIVE] Dhungel, P., Hei, Xiaojun., Ross, K., and N. Saxena,
"Pollution in P2P Live Video Streaming", International
Journal of Computer Networks & Communications
(IJCNC) Vol.1, No.2, Jul 2009.
- [PPSPPERF]

- Petrocco, R., Pouwelse, J., and D. Epema, "Performance analysis of the Libswift P2P streaming protocol", IEEE International Conference on Peer-to-Peer Computing (P2P'12), Tarragona, Spain, Sept 2012.
- [RFC2564] Kalbfleisch, C., Krupczak, C., Presuhn, R., and J. Saperia, "Application Management MIB", RFC 2564, May 1999.
- [RFC2790] Waldbusser, S. and P. Grillo, "Host Resources MIB", RFC 2790, March 2000.
- [RFC2975] Aboba, B., Arkko, J., and D. Harrington, "Introduction to Accounting Management", RFC 2975, October 2000.
- [RFC3365] Schiller, J., "Strong Security Requirements for Internet Engineering Task Force Standard Protocols", BCP 61, RFC 3365, August 2002.
- [RFC3729] Waldbusser, S., "Application Performance Measurement MIB", RFC 3729, March 2004.
- [RFC4113] Fenner, B. and J. Flick, "Management Information Base for the User Datagram Protocol (UDP)", RFC 4113, June 2005.
- [RFC4150] Dietz, R. and R. Cole, "Transport Performance Metrics MIB", RFC 4150, August 2005.
- [RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, December 2005.
- [RFC4821] Mathis, M. and J. Heffner, "Packetization Layer Path MTU Discovery", RFC 4821, March 2007.
- [RFC4960] Stewart, R., "Stream Control Transmission Protocol", RFC 4960, September 2007.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.
- [RFC5389] Rosenberg, J., Mahy, R., Matthews, P., and D. Wing, "Session Traversal Utilities for NAT (STUN)", RFC 5389, October 2008.
- [RFC5424] Gerhards, R., "The Syslog Protocol", RFC 5424, March 2009.
- [RFC5706] Harrington, D., "Guidelines for Considering Operations and Management of New Protocols and Protocol Extensions",

RFC 5706, November 2009.

- [RFC5971] Schulzrinne, H. and R. Hancock, "GIST: General Internet Signalling Transport", RFC 5971, October 2010.
- [RFC6241] Enns, R., Bjorklund, M., Schoenwaelder, J., and A. Bierman, "Network Configuration Protocol (NETCONF)", RFC 6241, June 2011.
- [RFC6347] Rescorla, E. and N. Modadugu, "Datagram Transport Layer Security Version 1.2", RFC 6347, January 2012.
- [RFC6709] Carpenter, B., Aboba, B., and S. Cheshire, "Design Considerations for Protocol Extensions", RFC 6709, September 2012.
- [RFC6817] Shalunov, S., Hazel, G., Iyengar, J., and M. Kuehlewind, "Low Extra Delay Background Transport (LEDBAT)", RFC 6817, December 2012.
- [SECDHTS] Urdaneta, G., Pierre, G., and M. van Steen, "A Survey of DHT Security Techniques", ACM Computing Surveys vol. 43(2), Jun 2011.
- [SIGMCAST] Wong, C. and S. Lam, "Digital Signatures for Flows and Multicasts", IEEE/ACM Transactions on Networking 7(4), pp. 502-513, 1999.
- [SPS] Jesi, G., Montresor, A., and M. van Steen, "Secure Peer Sampling", Computer Networks vol. 54(12), pp. 2086-2098, Elsevier, Aug 2010.
- [SWIFTIMPL] Grishchenko, V., Paananen, J., Pronchenkov, A., Bakker, A., and R. Petrocco, "Swift reference implementation", 2012, <<https://svn.tribler.org/libswift/>>.
- [TIT4TAT] Cohen, B., "Incentives Build Robustness in BitTorrent", 1st Workshop on Economics of Peer-to-Peer Systems, Berkeley, CA, USA, Jun 2003.

Appendix A. Revision History

- 00 2011-12-19 Initial version.
- 01 2012-01-30 Minor text revision:
- * Changed heading to "A. Bakker"
 - * Changed title to *Peer* Protocol, and abbreviation PPSPP.
 - * Replaced swift with PPSPP.
 - * Removed Sec. 6.4. "HTTP (as PPSP)".
 - * Renamed Sec. 8.4. to "Chunk Picking Algorithms".
 - * Resolved Ticket #3: Removed sentence about random set of peers.
 - * Resolved Ticket #6: Added clarification to "Chunk Picking Algorithms" section.
 - * Resolved Ticket #11: Added Sec. 3.12 on Storage Independence
 - * Resolved Ticket #14: Added clarification to "Automatic Size Detection" section.
 - * Resolved Ticket #15: Operation section now states it shows example behaviour for a specific set of policies and schemes.
 - * Resolved Ticket #30: Explained why multiple REQUESTs in one datagram.
 - * Resolved Ticket #31: Renamed PEX_ADD message to PEX_RES.
 - * Resolved Ticket #32: Renamed Sec 3.8. to "Keep Alive Signaling", and updated explanation.
 - * Resolved Ticket #33: Explained NAT hole punching via only PPSPP messages.
 - * Resolved Ticket #34: Added section about limited overhead of the Merkle hash tree scheme.
- 02 2012-04-17 Major revision
- * Allow different chunk addressing and content integrity protection schemes (ticket #13):

- * Added chunk ID, chunk specification, chunk addressing scheme, etc. to terminology.
- * Created new Sections 4 and 5 discussing chunk addressing and content integrity protection schemes, respectively and moved relevant sections on bin numbering and Merkle hash trees there.
- * Renamed Section 4 to "Merkle Hash Trees and The Automatic Detection of Content Size".
- * Reformulated automatic size detection in terms of nodes, not bins.
- * Extended HANDSHAKE message to carry protocol options and created Section 8 on Protocol options. VERSION and MSGTYPE_RCVD messages replaced with protocol options.
- * Renamed HASH message to INTEGRITY.
- * Renamed HINT to REQUEST.
- * Added description of chunk addressing via (start,end) ranges.
- * Resolved Ticket #26: Extended "Security Considerations" with section on the handshake procedure.
- * Resolved Ticket #17: Defined recently as "in last 60 seconds" in PEX.
- * Resolved Ticket #20: Extended "Security Considerations" with design to make Peer Address Exchange more secure.
- * Resolved Ticket #38+39 / PPSP.SEC.REQ-2+3: Extended "Security Considerations" with a section on confidentiality of content.
- * Resolved Ticket #40+42 / PPSP.SEC.REQ-4+6: Extended "Security Considerations" with a per-message analysis of threats and how PPSP is protected from them.
- * Progressed Ticket #41 / PPSP.SEC.REQ-5: Extended "Security Considerations" with a section on possible ways of excluding bad or broken peers from the system.
- * Moved Rationale to Appendix.
- * Resolved Ticket #43: Updated Live Streaming section to include "Sign All" content authentication, and reference to

[SIGMCAST] following discussion with Fabio Picconi.

- * Resolved Ticket #12: Added a CANCEL message to cancel REQUESTs for the same data that were sent to multiple peers at the same time in time-critical situations.

-03 2012-10-22 Major revision

- * Updated Abstract and Introduction, removing download case.
- * Resolved Ticket #4: Added explicit CHOKE/UNCHOKE messages.
- * Removed directory lists unused in streaming.
- * Resolved Ticket #22, #23, #28: Failure behaviour, error codes and dealing with peer crashes.
- * Resolved Ticket #13: Chunk ranges are the default chunk addressing scheme that all peers MUST support.
- * Added a section on compatibility between chunk addressing schemes.
- * Expanded the explanation of Unified Merkle Trees as a method for content integrity protection for live streams.
- * Added a section on forgetting chunks in live streaming.
- * Added "End" option to protocol options and corrected bugs in UDP encapsulation, following Karl Knutsson's comments.
- * Added SHA-2 support for Merkle Hash functions.
- * Added content integrity protection methods for live streaming to the relevant protocol option.
- * Added a Live Signature Algorithm protocol option.
- * Resolved Ticket #24+27: The choice for UDP + LEDBAT as transport has now been reflected in the draft. TCP and RTP encapsulations have been removed.
- * Superfluous parts of Section 10 on extensibility have been removed.
- * Removed appendix with Rationale.

- * Resolved Ticket #21+25: PPSPP currently uses LEDBAT and the DATA and ACK messages now contain the time fields it requires. Should other congestion control algorithms be supported in the future, a protocol option will be added.

-04 2012-11-07 Minor revision

- * Corrected typos.
- * Added empty protocol option list when HANDSHAKE is used for explicitly closing a channel in the UDP encapsulation.
- * Corrected definition of a range chunk specification to be a single (start,end) pair. To send multiple disjunct ranges multiple messages should be used.
- * Clarified that in a range chunk specification the end is inclusive. I.e., [start,end] not [start,end)
- * Added PEX_REScert message to carry a membership certificate. Renamed PEX_RES to PEX_RESv4.
- * Added a guideline about private and link-local addresses in PEX_RES messages.
- * Defined the format of the public key that is used as swarm ID in live streaming.
- * Clarified that a HANDSHAKE message must be the first message in a datagram.
- * Clarified sending INTEGRITY messages ahead in a separate datagram if not all necessary hashes that still need to be sent and the chunk fit into a single datagram. Defined an order for the INTEGRITY messages.
- * Clarified rare case of sending multiple DATA messages in one datagram.
- * Clarified UDP datagrams carrying PPSPP should adhere to the network's MTU to avoid IP fragmentation.
- * Defined value for version protocol option.
- * Added small clarifications and corrected typos.
- * Extended versioning scheme to Min/max versioning scheme defined in [RFC6709], Section 4.1, following Riccardo

Bernardini's suggestion.

- * Processed comments on unclear phrasing from Riccardo Bernardini.
- * Added a guideline on when to declare a peer dead.
- * Made sure all essential references are listed as Normative references following RFC3967.

-05 2013-01-23 Minor revision

- * Corrected category to Standards Track.
- * Clarified that swarm identifier is a required protocol option in an initiating HANDSHAKE in the UDP encapsulation.
- * Added IANA considerations and tablisted name spaces for registry definition.

-06 2013-02-11 Minor revision

- * Updated "Overall Operation" to have more context (HTML5 video).
- * Clarified wording on PEX_REQ.
- * Clarified wording on SIGNED_INTEGRITY.
- * Added a reference on how ALTO can be used with PPSP.
- * Added Manageability Consideration section following RFC5706.
- * Clarified that implementations SHOULD implement the "Unified Merkle Tree" content integrity protection method for live, and MAY implement "Sign All".
- * Made SHA1 hash function mandatory-to-implement as Merkle Tree Hash function and explained the security considerations.
- * Made RSA/SHA1 mandatory-to-implement as Live Signature Algorithm for integrity protection while live streaming.
- * Clarified that implementations MUST implement addressing via 32-bit chunk ranges.
- * Made LEDBAT an Informational reference to prevent a so-called "down ref".

- * Updated reference to PPSP problem statement and requirements document.
- * Used kibibyte unit in formal sections.

-07 2013-06-19 Revision following AD Review

Quoting the AD review by Martin Stiernerling: ***High-level issues:

1) Merkle Hash Trees I have found the document very confusing on whether Merkle Hash Trees (MHTs) and the for the MHT required bin numbering scheme are now optional or mandatory. Parts of the draft make the impression that either of them or both or optional (mainly in the beginning of the document), while Section 5 and later Sections are relying heavily on MHTs. My naive reading of the current draft is that you could rely on start-end ranges for chunk addressing and MHTs for content protection. However, I do know that this combination is not working. If MHTs are really optional, including the bin numbering, the document should really state this and make clear what the operations of the protocol are with the mandatory to implement (MTI) mechanisms. The MHT, bins, and all the protocol handling should go in an appendix. There is a call to make for the WG: I do know that MHTs were considered by some as burden and they have called for a leaner way, i.e., the start-end ranges. The call for the leaner way has been implemented in the document but not fully.

- + The text now states that MHTs SHOULD be used unless in benign environments and are mandatory-to-implement. It also states that only start-end chunk range is mandatory-to-implement, and bins are optional.

2) LEDBAT as congestion control vs. PPSP The PPSP peer protocol is intended for the Standards Track and relies in a normative manner on LEDBAT (RFC 6817). LEDBAT as such is an ****experimental**** delay-based congestion control algorithm. A Standards Track protocol cannot normatively rely on an Experimental congestion control mechanism (or RFC in general). There are ways out of this situation: i) Do not use ledbat: this would call for another congestion control mechanism to be described in the PPSP draft. ii) Work on an 'upgrade' of the LEDBAT specification to Standards Track: Possible, but a very long way. iii) Agree on having PPSP also as Experimental protocol. I'm currently leaning towards option iii), but this is my pure personal opinion as an

individual in the IETF.

- + A new paragraph has been added to Section 8.16 describing the widespread use of LEDBAT in current P2P systems. Hence, aim is a DOWNREF procedure.

3) No formal protocol message definition Section 7 and more specific Section 8 describe the protocol syntax of the protocol options and the messages, though Section 8 is talking about UDP encapsulation. Section 7 is hard to digest if someone should implement the options, see also later, but Section 8 is almost impossible to understand by somebody who has not been involved in the PPSP working group. See also further down for a more detailed review of the sections. To give an example out of Section 8.4: This section describes the HANDSHAKE message and gives examples how such a HANDSHAKE message could look like. But no formal definition of the message is given leaving a number of things unclear, such as what the local channel number and what's the remote channel number is. This is implicitly defined, but that is not a good way of writing Standards Track drafts.

- + We added the usual bit-based ASCII art representations.

4) Implicit use of default values There are a number of places all over the draft where default values are defined. Many of those default values are used when there are no values explicitly signaled, e.g., the default chunk size of 1 Kbyte in Section 8.4 or Section Section 7.5. with the default for the Content Integrity Protection Method. I have the feeling that the protocol and the surroundings (e.g., what comes in via the 'tracker') are over-optimized, e.g., always providing the Content Integrity Protection Method as part of the Protocol options will not waste more than 2 bytes in a HANDSHAKE message. Further, I do not see the need to define a default chunk size in the base protocol specification, as this default can look very different, depending on who is deploying the protocol and in what context. This calls for a more dynamic way of handling the system chunk size, either as part of an external mechanisms (e.g. via the tracker) or in the HANDSHAKE message.

- + Removed implicit defaults from protocol options. Chunk size is part of the content's metadata and thus configurable. The default 1KiB has been turned into a recommendation.

5) Concept of channels The concept of channels is good but it is introduced too late in the draft, namely in Section 8.3, and it is introduced with very few words. Why isn't this introduced as part of Section 2 or Section 3, also in the relationship to the used transport protocol? I.e., the intention is to keep only one transport 'connection' between two distinct peers and to allow to run multiple swarm instances at the same time over the same transport. And how do swarms and channels correlate?

+ Concept now introduced in Section 3 with a figure.

***Technicals:

- Section 2.1, 2nd paragraph, about the tracker: I haven't seen a single place where the interaction with a tracker is discussed or where the tracker less operation is discussed in contrast. It is further unclear what type of information is really required from a tracker. A tracker (or a resource directory) would need to provide more than IP address & port, e.g., the used transport protocol for the protocol exchange (given that other transports are allowed), used chunk size, chunk addressing scheme, etc

+ Interaction with tracking facilities in general is discussed in the Operations and Management section, Section 12.1.1. This also discusses swarm metadata and information required from tracking facility. Decentralized tracking in PPSP is discussed in Section 3.10.1.

- Section 2.3, the 1st paragraph, 'close-channel': This has been the first time where I stumbled over the channel without knowing the concept.

+ Rephrased.

- Section 3.1: ordering of messages The 1st sentence implies that ordering of messages in a datagram matters a lot. This is outlined later in the document, but I would add this as part of 3., i.e., the messages are processed in the strict order or something along this line.

+ Phrase added.

- Section 3.1, 1st paragraph, options to include I would not say anything about 'SHOULD include options' here, as this is anyhow described in Section 8.

+ Phrase removed.

- Section 3.1, 2nd paragraph: "Datagrams exchanged MAY also contain some minor payload, e.g. HAVE messages to indicate the current progress of a peer or a REQUEST (see Section 3.7)." to be added, just to make it clear IMHO: ", but MUST NOT include any DATA message".

+ Added.

- Section 3.2, 2nd paragraph: "In particular, whenever a receiving peer has successfully checked the integrity of a chunk or interval of chunks it MUST send a HAVE message to all peers it wants to interact with in the near future." This looks like a place where a lot of traffic can be send out of a peer, i.e., whenever a chunk arrives a HAVE message must be sent. I don't believe that this should be mandated by the protocol specification, but there should guidance on when to send this, e.g., peers might be also able to wait for a short period of time to gather more chunks to be reported in HAVE. Or should in this case a single UDP datagram contain multiple HAVES?

+ Clarified that this is indeed controlled by a policy outside the peer protocol that can decide to piggyback onto other traffic or wait till multile chunks are verified.

- Section 3.4 on ACKs This section looks pretty weak, as ACKs may be sent but on the other hand MUST be sent if ledbat is used. I would simply say: - ACK MUST be sent if an unreliable transport protocol is used - ACK MAY be sent if a reliable transport protocol is used - keep clarification about ledbat.

+ DONE.

- Section 3.5: Give text where INTEGRITY is described at least for the MTI scheme.

+ DONE.

- Section 3.7, 2nd paragraph - all 'MAY' are actually not right here. Please remove or replace them with lower letters if appropriate. - It is not clear what the 'sequentially' means exactly. Is it in the received order?

- + Rephrased MAYs. "Sequentially" replaced with "received order".
- Section 3.8: Please replace 'MAY' by can, as those are not normative behaviors but more the fact that peers can, for instance, request urgent data.
- + DONE.
- Section 3.9 Same comment as for the Section 3.8 just above this comment.
- + DONE.
- Section 3.9 waiting for responses OLD " When peer B receives a CHOKe message from A it MUST NOT send new REQUEST messages and SHOULD NOT expect answers to any outstanding ones." NEW " When peer B receives a CHOKe message from A it MUST NOT send new REQUEST messages and it cannot expect answers to any outstanding ones, as the transfer of chunks is choked."
- + DONE.
- Section 3.10.2 This whole section about PEX hole punching reads very, very experimental. The STUN method is ok, but PEX isn't. First of all, the safe behavior for a peer when it receives unsolicited PEX messages, is to discard those messages. Second, this unsolicited PEX messages trigger some behavior which may open an attack vector. The best way, but this needs more discussion, is to include to some token in the messages that are exchanged in order to make avoid any blind attacks here. However, this will need more and detailed discussions of the purpose of this.
- + We moved parts of the security analysis of PEX up, such that all mechanisms are explained in the main text, and the analysis of what attacks there are and how these mechanisms prevent them is in the Sec. Considerations section.
- + The section about hole punching was removed, lacking a reference to the experiments we conducted with this exact variant of the mechanism.
- Section 3.11 I don't see the 'MUST send keep-alive' as a mandatory requirement, as peers might have good reasons not to send any keep alive. Why not saying 'A peer can send a

keep-alive' and it 'MUST use the simple datagram...' as already described. Though there is also no really need to say MUST.

- + Now Section 3.12. Rephrased and clarified the reason and consequences of sending keep-alive msgs.

- Section 4 The syntax definition for each of the chunk addressing schemes is missing. This is not suitable for any specification that aims at interoperable implementations.

- + We added the usual bit-based ASCII art representations.

- Section 4.3.2 PPSPP peers MUST use the ACK message if an unreliable transport protocol is used.

- + DONE.

- Section 4.4 Has been tested in an implementation? I would like to understand the need for such a section, as in my understanding a peer implementation should chose one scheme and support this and there shouldn't be the need to convert between the different schemes.

- + Yes, the reference implementation translates from chunk ranges on the wire to bins internally. However, for simplicity we now state that all peers in a swarm MUST use the same method and the compatibility section has been removed.

- Section 5 This reads that MHTs are mandatory to implement while the document makes the impression that MHTs are optional.

- + Rephrased, see High-level issues.

- Section 5.3 " so each datagram SHOULD be processed separately and a loss of one datagram MUST NOT disrupt the flow" The MUST NOT is not a protocol specification requirement, but more an informative part saying that a lost message shouldn't impact the protocol machinery, but it can impact the overall operation. What is the flow here in that sentence?

- + Rephrased.

- Section 5.6.2. An illustrative example explaining how the automatic size detection works is required here.
- + Added a paragraph with an example that follows the figure used during the explanation. A state diagram could also be added, but might be a bit redundant.
- Section 6.1, 4th paragraph: Where do I find the 1 byte algorithm field in the swarm ID? The swarm ID is not really defined in a single place.
- + Expanded. Added a formal definition.
- Section 7.3 The described min/max versioning relies on the fact that there are major and minor version numbers. I cannot find any major and minor version number scheme in the draft.
- + Actually, it does not. There is a single unstructured version number.
- Section 7.4, Length field It is not clear what the 'Length' field is referring to. Further, it is not clear of the swarm IDs are concatenated in one swarm ID option, of each swarm ID must be placed in a separate swarm ID option.
- + Clarified.
- Section 7.6 MHTs are mandatory to support though MHTs are optional?
- + Clarified.
- Section 7.7 'key size ... derived from the swarm ID'. This relates to my high level comment no 4. on the use of implicit information. Either it is clearly specified how this information is derived or there is a protocol field/information about the size.
- + Key size derivation procedure added to description of SIGNED_INTEGRITY in UDP encapsulation.
- Section 7.8 I would recommend to say that the default MUST be supported, but the peer must always signal what method it is supporting or at least using.
- + Corrected, see High-level issues 4.)

- Section 7.10 I have not understood how the 'Lenght' field relates to the message bitmap and how long the message bitmap can grow. The figure looks like a maximum of 16 bits?

+ Clarified.

- Section 8 I do not see the value of the text in the preface of Section 8. I would say that this text should say what is mandatory and what's not, i.e., MUST use UDP and MUST use LEDBAT. Potentially saying that future protocol versions can also run over other transport protocols.

+ Adjusted.

- Section 8.1 about Maximum Transfer Unit (MTU) The text is discussing that a Ethernet can carry 1500 bytes. This is true, but the Ethernet payload is not the normative MTU across all of the Internet. For IPv6 the min MTU is 1280 bytes and for IPv4 it is 576 bytes, though for IPv4 it can be theoretically much lower at 64 bytes. It would move the definition of the default chunk size to a recommendation with text saying that this size has a high likelihood to travel end-to-end in the Internet without any fragmentation. Fragmentation might increase the loss of complete chunks, as one lost fragment will cause the loss of a complete chunk. One way of getting an informed decision on whether chunks can travel in their size is to use the Don't Fragment (DF) bit in IPv4 and also to watch for ICMP error messages. However, ICMP error messages are not a reliable indication, but they can be some indication.

+ 1 KiB chunk size has been made a recommendation.

+ Added a small paragraph discussing the optional integration of MTU path discovery.

- Section 8.1 Definition of the default chunk size There is no need to define a default chunk size, if the chunk size would be always signaled per swarm. This is another default/implicit value places that is unnecessary.

+ The chunk size is always part of the content's metadata.

- Section 8.3: see also my comment no 3. The concept of channels is introduced very late and with few words. A figure to explain the concept will help a lot and also more formal text on what a channel is and how they are identified. Also what the init channel is.

- + Concept now introduced in Section 3.11.
- Section 8 in general: There is no formal definition of the messages, just bit pattern examples.
- + We added the usual bit-based ASCII art representations.
- Section 8.4 (as example for the other Sections in 8.x): i) What is the '(CHANNEL' paramter? Is it actually a parameter? ii) it is implicit that the first channel no (0000000) is the remote peer's channel and that the second channel no (00000011) is the local peer's channel, right? This isn't clear from the text, but my guess.
- + We added the usual bit-based ASCII art representations.
- Section 8.5 Can HAVE messages multiple bin specs in one message or do I have to make a HAVE message for each bin?
- + Clarified.
- Section 8.6 What is the formal defintion of a DATA message? That's completely missing or I have not understood it.
- + We added the usual bit-based ASCII art representations.
- Section 8.7 looks just underspecified, especially as this is the link to LEDBAT.
- + Implementors will unfortunately need to read the full LEDBAT specification.
- Section 8.11 How are the chunks specified here? The formal syntax definition or reference to one is missing.
- + We added the usual bit-based ASCII art representations.
- Section 8.13 I'm lost on this section, as I haven't fully understood the concept of the PEX in this document. Especially not why there is the PEX_REScert.
- + We moved parts of the security analysis of PEX up into 3.10, such that all mechanisms are explained in the main text, and the analysis of what attacks there are and how these mechanisms prevent them is in the Sec. Considerations section.

- Section 11 The RFC required for protocol extensions of a standards track protocol looks odd. This must be at least IETF Review or Standards Action.

+ Policy changed to "IETF Review" and the section was extended with information about data types and required information.

***Editorials:

- Abstract (and probably also other places), 1st sentence of, PPSP is not a transport protocol, just a protocol

+ DONE.

- Section 1.1, 4th paragraph: I would remove the reference to rmcatt, as it is not yet clear what the outcome of the rmcatt wg will be

+ DONE.

- Section 1.3, on page 8, about seeding/leeching: I would break it in to sub-bullets.

+ DONE.

- Section 2.1 and following: These are examples, isn't it? If so, this should be mentioned or clarified.

+ DONE. All subsections now labeled "Example:".

- Section 2.1: What is the PPSP Url?

+ Reformulated in terms of "Imagine there is a PPSP URL".

- Section 2.3, the 1st paragraph, detection of dead peers: It would be good to say where this detection is described in the remainder of the draft. Just for completeness.

+ DONE. Dead peer detection is now a separate section and referenced here.

- Section 2.2, the very last paragraph, 'Peer A MAY also': This 'MAY' is not useful here. I would just write 'Peer A can also', as there is nothing normative described here.

+ DONE.

- Section 3.2, last paragraph: What is the latter confinement? This is not clear to me.
- + Rephrased.
- Section 3.9, last sentence I am not sure to what the reference to Section 3.7 is pointing in this respect.
- + Rephrased.
- Section 3.10.1 about PEX messages The text says 'PPSPP optionally features...'. I have not understood if this optionally refers to mandatory to implement but optionally to use, or if the PEX messages are optionally to implement.
- + Made it clear that is OPTIONAL and not mandatory-to-implement.
- Section 3.12 I'm not sure what this section is telling exactly. Isn't just saying that PPSPP as such does not care how chunks are stored locally, as this is implementation dependent?
- + Yes. Removed.
- Section 4.2, page 15, 1st paragraph: OLD 'A PPSPP peer MAY support' NEW 'The support for this scheme is OPTIONAL'
- + DONE, for byte ranges as well.
- Section 6.1.1 This section is not describing sign-all, but rather a justification why it may still work. This doesn't help at all.
- Section 7, 1st paragraph Why is there a reference to RFC 2132?
- + Removed, just similarity in format.
- Section 7 in general i) It is common to give bit positions in the figures where the syntax of options is described. This allows to count how many bits are used for a protocol field more easily and also way more reliable. ii) Please add also Figure labels to the syntax definitions of the options. This makes it easier to reference them later on if needed.
- Section 8.1 1 kibibyte is 1 kbyte?

- + Mentioned base 1024 in Terminology. Changed to 1024 bytes where appropriate.
- Section 8.2, last paragraph i) "All messages are idempotent" in what respect? ii) "or recognizable as duplicates" but how are they recognized as duplicates?
- + Idempotent means that processing a message twice does not lead to a different state than processing them once. Resent handshakes can be recognized as duplicates because a peer already recorded the first connection attempt in its state. Updated text.
- Section 8.5, last sentence in brackets: What is this last sentence about?
- + Was explanation of the on-the-wire bytes shown.
- Section 8.13 " If sender of the PEX_REQ message does not have a private or link-local address, then the PEX_RES* messages MUST NOT contain such addresses [RFC1918][RFC4291]." What is this text saying? Do not include what you do not have anyway?
- + Rephrased. It tries to say that internal addresses must not be leaked to external peers.
- Section 8.14 There is no single place where all the constants are collected and also documented what the default values or the recommended values. For instance in this Section 8.14 where the dead peer time out is set to 3 minutes and also the number of datagrams that should have sent. I would make a section or subsection to discuss dead peers and how they are detected and just link to the keep-alive mechanism in Section 8.14.
- + The Section 12.1.6 section was rewritten for this in the Ops & Mgmt part.
- Section 11 This section needs to be overhauled once the document is ready for the IESG. The section is not wrong but can be improved to help IANA.
- + The section was extended with information about data types and required information.

-08 2013-08-8 Continued Revision following AD Review

Please see the -07 entry for our responses to the comments.

Added ECDSAP256SHA256 and ECDSAP384SHA384 as mandatory-to-implement live signature algorithms, as they provide small swarm IDs.

Added line that a peer SHOULD NOT send HAVEs to peers that already have the complete content (e.g. in video-on-demand scenarios).

In response to a remark at WG meeting at IETF 87 we added a paragraph on OPTIONAL MTU discovery using PPSPP messages to Section 8.1.

Authors' Addresses

Arno Bakker
Vrije Universiteit Amsterdam
De Boelelaan 1081
Amsterdam, 1081HV
The Netherlands

Phone:
Email: arno@cs.vu.nl

Riccardo Petrocco
Technische Universiteit Delft
Mekelweg 4
Delft, 2628CD
The Netherlands

Phone:
Email: r.petrocco@gmail.com

Victor Grishchenko
Technische Universiteit Delft
Mekelweg 4
Delft, 2628CD
The Netherlands

Phone:
Email: victor.grishchenko@gmail.com

