Network Working Group R. Penno
Internet Draft Juniper Networks
Intended status: Informational May 9, 2008
Expires: November 2008

P2P Status and Requirements
draft-ietf-penno-p2p-status-requirements-00

Status of this Memo

Copyright Notice

Abstract

The use of P2P protocols by end users is widespread. These protocols are meant to exchange, replicate, stream or download files with little human intervention, trying at the same time to minimize the download time of the files requested by any single peer. The ubiquity of such P2P networks has created a steep rise in subscriber bandwidth

consumption that is at odds with ISP's original capacity planning. In this document we will describe the status and requirements for some of the proposed solutions to tackle this problem.

Conventions used in this document

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 [1].

Table of Contents

1. Introduction

The use of P2P protocols by end users is widespread. These protocols are meant to exchange, replicate, stream or download files with little human intervention, trying at the same time to minimize the download time of the files requested by any single peer. This is done by opening several connections to multiple peers and downloading one or more chunks of the file from each one, selecting faster peers, amongst others.

The availability of large amounts of content hosted by distributed peers across the globe coupled with the fact that once downloads are scheduled the clients pretty much run on their own, created a change in data traffic patterns.

Now end users that utilize P2P clients in general not only consume more bandwidth but consume bandwidth throughout the whole day, and more specifically consume more upstream bandwidth since they share files, and more downstream bandwidth since P2P clients try to minimize download time.

Given that ISP networks were planned for different traffic patterns, this surge in bandwidth consumption created several problems such as congestion, delays, unfairness between end users and a spike in inter-ISP transit costs.

In this document we will discuss the status and requirements for solutions tackling this problem.

2. Terminology

Cache - A IP Host that holds a copy of an original content

End user - A IP host running P2P applications, e.g., P2P client

NSIS - Next Steps in Signaling

P2P - Peer to Peer

Peer - Any IP host that joins a P2P network to share or download content

Rendezvous Point - A host in a P2P network that collects and disseminates peer information.

Subscriber - Used interchangeably with end user. Although a broadband subscriber as a paying entity can encompass many end users

3. P2P Conflicting Interests

Any P2P solution needs to juggle three different sets of interests. End users want to maximize their bandwidth consumption to receive downloads faster and upload files to others peers with privacy. The ISPs want to manage and provide a fair share for each subscriber in an economic viable way.

Finally, more recently content distributors are also utilizing P2P as a means to economically distribute content or services in a cost effective redundant manner. For this reason the requirements of content owners or distributors need to be taken into account when finding a solution for P2P in ISP networks.

4. P2P clients and Network

When it comes to interactions between P2P clients and the network, there are a number of possible ways to deal with the problem such as:

- o ISPs can take unilateral actions such as throttle or tear down P2P or transport sessions as whole for each subscriber, or place caches strategically in the network,
- o P2P clients can also take actions unilaterally such as react to congestion or delay in the network,
- o P2P clients can make use of rendezvous points located within the ISP's network
- o Or finally, P2P clients and the network can exchange information through policies or signaling.

Policies could be divided in two classes: session dependent and independent. Session dependent policies can change from session to session and can be a result of the client signaling for bandwidth on a per session basis using an existing protocol such as NSIS.

Session Independent polices have longer time spans and are a result of network wide provisioning. For example, an ISP can have a list of networks that the clients could give preference when selecting peers. Such exchange of information could be done independent of the P2P protocol.

5. P2P Dynamic Nature

There are many issues that contribute to the dynamic nature of the problem:

- o The number of established and dynamic sessions to other peers and rendezvous points,
- o Bandwidth consumption that depends not only on the originating ISP, but terminating ISP and other peer's client configuration,
- o The different protocol implementations and variations or extension within the same protocol
- o The distributed nature of P2P protocols usually implies that there is no effective single point of control for the providers

o Flash crowds are common since new popular content can become available and spread quickly.

So, any protocol that signals session dependent policies needs to be quite scalable, and session independent policy should have longer time spans when compared with time taken to download a large file.

6. Caches

One of the mechanisms that is an option to ISPs and can be deployed independent of clients is that of strategically placed caches.

Caches from a HTTP perspective are well understood [8][9]. There are many studies on web workload, both on real networks of diverse architectures and simulations. Moreover, Web follows a client-server model where many clients try to access a limited number of servers.

In P2P the number of servers is not limited and could be any or all of the peers; while total number of files might be very large, the number of popular files is usually observed to be limited.

The number of studies today in the area of P2P caching in real diverse networks is limited and although it is a promising area and intuitively caches can definitely help mitigate the bandwidth issue, it remains to be seen the economic viability in terms actual hit ratio and bandwidth savings vs. necessary storage and price.

7. Rendezvous Points

Rendezvous points are hosts that store and disseminate information amongst P2P clients such as which clients are active, their current shared files, bandwidth consumption, amongst others and in some cases help P2P clients join a P2P network (although these two functionalities can be done by different hosts). Depending on the P2P network these rendezvous points have different functionalities and are called supernodes, servers, ultrapeers or trackers.

Although the specifics of Rendezvous Points are different for each P2P protocol, a protocol agnostic version could help alleviate P2P bandwidth related issues through a smart peer selection.

8. High Availability

It is also worth noting that P2P is a desirable technology for content owners and distributors because P2P provides both high availability and redundancy since content is distributed throughout the Internet as opposed to individual data centers or servers.

9. Security Considerations

There are many possible attacks on P2P networks. For the specific solutions discussed, if ISP and P2P clients are to exchange information in whatever manner the most relevant ones would be related to the channel between P2P client and network. A non-exhaustive list would be rendezvous point poisoning, rendezvous point spoofing and network policy spoofing or poisoning.

10. IANA Considerations

None at this time.

11. Conclusions

From a technical standpoint there are many possible ways to tackle the P2P problem. It remains to be seen in real networks which combination of combination of caches, P2P signaling or provisioning protocol or rendezvous points is going to be the solution.

From an IETF perspective if the goal is standardization, working on solutions that are protocol agnostic is probably the way to move forward given the history on the number of P2P protocols and variations. Interception use of caches, network provisioning, and others should probably be captured in a BCP.

12. Acknowledgments

None at this time

13. References

13.1. Normative References

[1] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[2] R. Hancock, et. al. Next Steps in Signaling (NSIS): Framework, RFC4080, June 2005.

13.2. Informative References

[3] Nathaniel Leibowitz, Aviv Bergman, Roy Ben-shaul, Aviv Shavit, "Are file swapping networks cacheable? Characterizing p2p traffic", In Proc. of the 7th Int. WWW Caching Workshop 2002

[4]   T. Karagiannis, A. Broido, K. Claffy, and M. Faloutsos. "Transport Layer Identification of P2P Traffic", In International Measurement Conference, October 2004.

[5]   Gnutella Protocol Specification, http://wiki.limewire.org/index.php?title=GDF

[6]   Y. Kulbak and D. Bickson. The emule protocol specification. Technical report TR-2005-03, the Hebrew University of Jerusalem, 2005

[7]   Bittorrent Protocol Specification v1.0, http://wiki.theory.org/BitTorrentSpecification

[8]   Rabinovich, M. and Spatschek, O. "Web Caching and Replication". Addison-Wesley Longman Publishing Co., Inc. 2002

[9]   Krishnamurthy, B. and Rexford, J. "Web Protocols and Practice: Http/1.1, Networking Protocols, Caching, and Traffic Measurement." Addison-Wesley Longman Publishing Co., Inc. 2001

[10]   Saleh, O. and Hefeeda, M. "Modeling and Caching of Peer-to-Peer Traffic." In Proceedings of the Proceedings of the 2006 IEEE international Conference on Network Protocols, November 2006

[11]   David Erman, Dragos Ilie, and Adrian Popescu. "Bittorrent session characteristics and models". In Demetres Kouvatsos, editor, Proceedings of the 3rd International Working Conference on Performance Modelling and Evaluation of Heterogeneous Networks, pages P30/1--P30/10, Ilkley, West Yorkshire, UK, July 2005.

[12]   Ilie, D., Erman, D., Popescu, A. and Nilsson, A. "Measurement and Analysis of Gnutella Signaling Traffic," IPSI-2004 Internet Conference (IPSI-2004), Stockholm, Sweden, September 2004.

Author's Addresses

Reinaldo Penno
Juniper Networks
1194 N Mathilda Avenue

Phone:
Email: rpenno@juniper.net

Acknowledgment

Funding for the RFC Editor function is currently provided by the Internet Society.