

Peer-to-Peer File Sharing

The Effects of File Sharing on a Service Provider's Network

An Industry White Paper

Copyright © July 2002, Sandvine Incorporated
www.sandvine.com

408 Albert Street
Waterloo, Ontario
Canada
N2L 3V3

Executive Summary

Peer-to-peer (P2P) file sharing has emerged as the primary siphon of Internet bandwidth. Beginning with the Napster phenomenon of the late 1990s, the popularity of P2P has dramatically increased the volume of data transferred between Internet users. As a result, a rump percentage of the global Internet subscriber base is consuming a disproportionate share of bandwidth - certainly more than the per-user amounts provisioned by service providers.

Recent studies suggest that file sharing activity accounts for up to 60% of the traffic on any given service provider network. While asymmetric bandwidth consumption is a legitimate concern on its own, the ad-hoc nature of P2P communication means that a large amount of data traffic is pushed off-network (P2P clients don't care where other P2P clients are located) - driving up network access (NAP) fees.

By inflating the financial pressure on service providers' already low margins, P2P is quickly undermining the business model for basic Internet access. In the face of surging off-network traffic, the traditional provider approach - managing network costs through over-subscription - is no longer sufficient.

Yet the enormous popularity of file sharing and the breadth of competing protocols makes blocking P2P traffic a practical impossibility. Service providers have begun to experiment with tiered pricing based on monthly bandwidth consumption, or capping the amount of bandwidth available to P2P applications, but these approaches can easily be positioned as punitive by Internet lobby groups and competitors, generating dissatisfaction amongst subscribers and potentially aggravating customer churn.

This document explores the technology and infrastructure behind peer-to-peer file sharing, and its implications for long-term service provider profitability.

Peer-to-Peer Technology

A Brief History

Peer-to-peer refers to any relationship in which multiple, autonomous devices interact as equals. A peer-to-peer network is a type of network in which workstations may act as clients (requesting data), servers (offering data) and/or servents (both a client and a server). P2P technology enables the sharing of computer resources and services, including information, files, processing cycles and storage by direct exchange between systems (without the use of central servers). P2P technology allows computers, along with their users, to tap idle resources that would otherwise remain unused on individual workstations.

Prior to the remarkable rise and fall of Napster, systems for sharing files and information between computers were exceedingly limited, and largely confined to the World Wide Web (WWW), Local Area Networks (LANs) and the exchange of files via a File Transfer Protocol (FTP) connection.

As the speed and pervasiveness of personal computers (PCs) increased, as well as the speed and pervasiveness of Internet connections, so did public demand for file sharing technologies. Napster popularized file sharing and became, almost overnight, the single most popular P2P application. In February of 2001, Napster boasted a daily average of 1.57 million simultaneous users

(<http://www.cnn.com/2001/TECH/internet/06/28/napster.usage>, June 28 2001).

P2P has since emerged as the dominant component of bandwidth used by residential Internet subscribers. The evolution from Napster to KaZaA, Gnutella, Morpheus and others has dramatically increased the amount of data transferred across service provider networks. Many home PCs are now being used as P2P data servers twenty-four hours a day, seven days a week.

Styles of P2P

There are three basic styles of P2P file sharing:

- The One-to-One relationship, typically a transfer of files from PC to PC
- The more advanced One-to-Many relationship used by Napster, which enables a single host to communicate and share files with multiple nodes. Examples include mail servers connected to multiple mail clients and HTTP servers communicating with browsers.
- The Many-to-Many relationship used by Gnutella protocol clients like BearShare and Morpheus, which enables highly automated resource sharing among multiple nodes

Peer-to-Peer Frameworks

Centralized Framework

First generation P2P (e.g. Napster) utilizes the server-client network structure. The centralized server acts as a sort of “traffic cop,” as shown in Figure 1.




Figure 1: Centralized Peer-to-Peer Network

The central server maintains directories of shared files stored on each node. Each time a client logs on or off the network, the directory is updated. In this model, all control and search messages are sent to a central server. The central server then cross-references the client’s search request with its directory database and displays any matches to the requesting client. Once informed about a match, the client contacts the peer directly and downloads the requested file. The actual file is never stored on the central server.

The centralized P2P framework provides the highest performance when it comes to locating files. Every individual peer in the network must be registered, which ensures that all searches are comprehensive and execute quickly and efficiently.

Decentralized Framework

Second generation P2P (e.g. Gnutella protocol) uses a distributed model where there is no central server and every node has equal status. Each node acts as a servent, or a ‘peer’, and operates as both a client and a server to the network.

As is evident in Figure 2, every node within the framework tries to maintain some number of connections (typical range is 4-8) to other nodes at all times. This set of connected nodes carries the network traffic, which is essentially made up of queries, replies to those queries, and various control messages that facilitate the discovery of other nodes.




Figure 2: Decentralized Peer-to-Peer Network

In order to share files using the Gnutella protocol, the user requires a networked computer, ("Node X"), equipped with a Gnutella software program. Node X initiates a query by forwarding a request to another computer on the Gnutella network ("Node Y"). Node Y then forwards the query to everyone that it is connected to.

Although the span of this network is potentially infinite, it remains limited by the 'time-to-live' (TTL) constraint. Time-to-live refers to the layers of nodes that the request message will reach. Query messages are sent with a time-to-live field (typical range is 4-6) that is decremented and then forwarded by each node to all other connected nodes. If, after decrementing the TTL field, the TTL field is found to be zero, the message will not be forwarded along any further connections. Each node that receives the query will then reply (included in the reply is the file name, size, etc.) and all replies are forwarded back to the origin of the query, Node X, via Node Y. Node X can now open a direct connection with one of the replying nodes ("Node Z"), and download the file. Files are transferred directly, without the intervention of intermediate nodes (the download is executed using a protocol similar to HTTP version 1 protocol). This is the approach utilized by Gnutella protocol applications such as BearShare, Limewire, Gnucleus, and Morpheus.

A decentralized framework does not rely on a central server, and is therefore more robust than its centralized counterpart. Disadvantages to the decentralized model appear in the form of prolonged search times. An outgoing search request may need to travel through thousands of users before any results are identified.

Controlled Decentralized Framework

Third generation P2P (e.g. FastTrack, KaZaA, Grokster, Groove, and current Gnutella clients) employs a hybrid of the central-server and fully decentralized frameworks. Within this hybrid model, certain nodes in the network are elected 'super-nodes' or 'ultrapeers' and act as traffic cops for the other nodes. The super-nodes change dynamically as bandwidth and the network topology change. A client-node keeps only a small number of connections open and each of those connections is to a super-node. This has an effect of making the network scale, by reducing the number of nodes involved in message handling and routing, as well as by reducing the actual volume of traffic among them. It is because

of these super-nodes, which also act as search hubs, that the speed at which queries are answered within the controlled framework is comparable to a centralized network model. An example of this type of network is shown below in Figure 3.

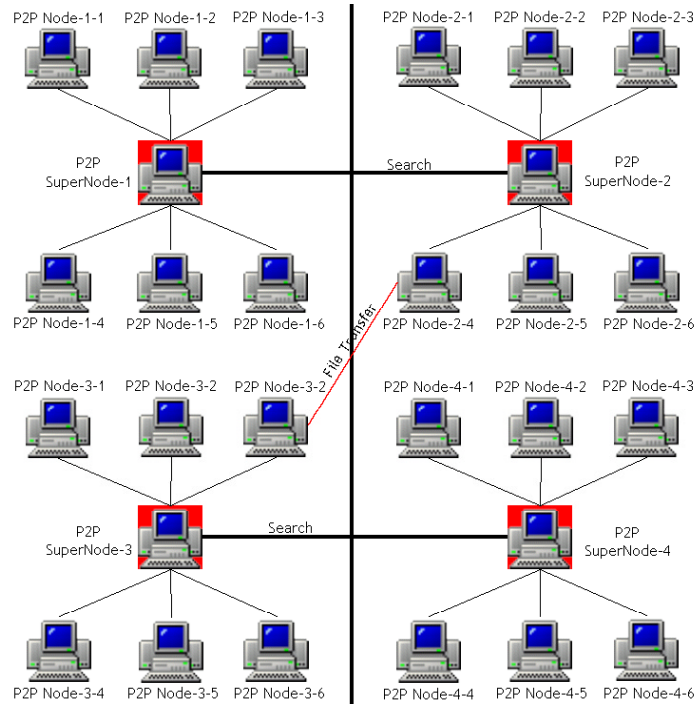


Figure 3: Controlled Decentralization Peer-to-Peer Network

In the controlled decentralization framework, each node forwards a list of its shared files to its super-node ("Node Y"). Search requests are directed to the appropriate Node Y, which will then forward the request to other super-nodes. When a match is found, the requesting node, or Node X, connects directly to the node with the match, Node Z, and downloads the file.

Peer-to-Peer Applications

Direct Exchange of Services

Peer-to-peer networks enable the sharing of services by direct exchange between nodes. Services include cache storage, disk storage, information, and files. It was through the efforts of Napster that this category of P2P application was first brought into public attention.

Grid Computing

Grid computing, also known as collaborative computing, is a form of P2P computing in which unused CPU cycles are channeled towards a common purpose. Grid computing became a popular topic when the [SETI@Home](http://www.setiathome.org) project was launched on May 17, 1999 (<http://www.berkeley.edu>). [SETI@Home](http://www.setiathome.org) is a screen saver application that harnesses the unused CPU cycles of hundreds of thousands of volunteers' computers to analyze the results of the search for extra-terrestrial intelligence. Grid computing is commonly found in science, biotech and financial environments, where there is a need for intense computer processing.

Distributed Information Infrastructure

Distributed information infrastructure is a method of P2P that brings all of the information assets and resources of an organization together to form a "Virtual Organization." A virtual organization may consist of multiple companies, or multiple branches, uniting as one to strive towards a common goal. Many firms within the healthcare industry, along with the scientific research and development sectors, use this type of P2P application to manage, distribute and retrieve important data and information. Distributed information infrastructure offers an effective, and efficient, way to span geographical and organizational boundaries.

Peer-to-Peer & Service Providers

Profitability of Service Providers

A service provider carries various costs that can be allocated to individual subscribers. One of the most significant of these costs is the service provider's Internet transit charge. Internet transit charges are a substantial variable cost; the more a subscriber uses the service, the more it costs the service provider. The competitive market for Internet access demands that customers be given unlimited access to the Internet. However, service providers purchase their bandwidth from an Inter-Exchange Carrier (IXC) based on the total bandwidth used. This undermines the profitability of offering flat rate Internet access service.

Direct Exchange of Services - The Contributing Factor

According to the Cooperative Association for Internet Data Analysis (CAIDA), service provider network traffic is dominated by peer-to-peer file sharing applications and WWW protocols. Figure 4 displays a breakdown of Internet traffic over a service provider's network in a given week (<http://www.caida.org>, Mar 13, 2002).

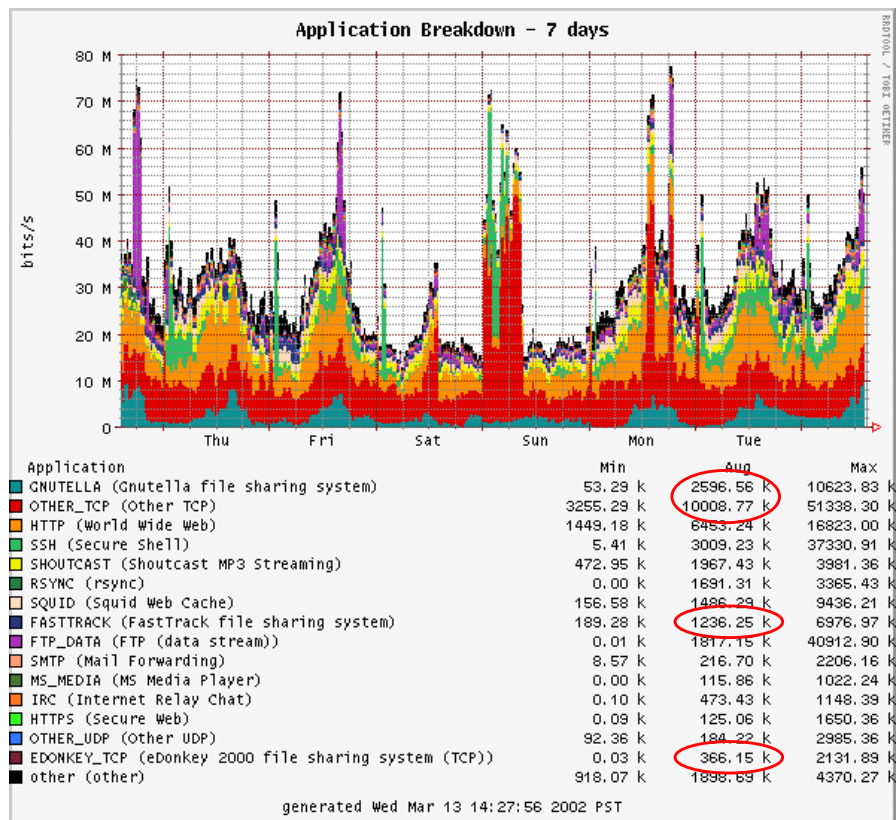


Figure 4: Internet Traffic Breakdown

P2P applications generate two types of network traffic:

- Network overhead traffic (searches, keep-alives)
- Data traffic (file transfers)

P2P network traffic consumes a large portion of bandwidth, and as P2P application usage continues to increase, so do service providers' Internet transit charges. At the height of Napster's popularity, Indiana University banned all P2P file swapping applications after discovering that the protocol was responsible for 50 percent of their network traffic (<http://www.redherring.com/insider/2000/1116/tech-oddcouple111600.html>, Nov 16, 2000).

P2P continues to evolve through the continuous development of new file swapping applications. Together, the FastTrack and Gnutella protocols currently boast an outstanding 2.9 million simultaneous users (www.slyck.com, July 2, 2002). Figure 5 illustrates the array of P2P applications currently available to online users.



Figure 5: Common P2P Applications

Peer-to-Peer Network

On start-up, a P2P application will connect to a number of other P2P nodes which can be anywhere on the Internet. Because there is no correlation to the underlying IP network structure and cost model, the closest P2P peer/node is rarely located on the same network. As a result, a very low percentage of P2P nodes within a service provider's network will ever connect to one another. The organization of a typical P2P network is displayed in Figure 6.




Figure 6: Typical P2P Network Overview

File sharing applications are comprised of two main bandwidth consumers:

- The connection management of the client to the network
- The downloading of files, which is simply the transfer of files from one P2P host to another, anywhere on the Internet.

The connection component of P2P traffic is comprised of a number of connections to different P2P hosts; anywhere in the network or the Internet. Each connection uses a number of messages intended to keep the connection alive over a period of time and insure that file searches are quickly resolved. This component of P2P networking is commonly referred to as 'protocol chatter'. Figure 7 illustrates the two main forms of P2P network traffic.



Figure 7: P2P Network Traffic

A popular misconception with P2P networking is that file transfers account for the vast majority of the bandwidth consumed. In reality, a very large percentage of bandwidth is required for protocol chatter. In some protocols, P2P chatter generates approximately 50 to 150 Kbps of traffic continuously per P2P host PC.

This leaves service providers with few options to reduce their Internet transit charges. Possible solutions include switching to tiered bandwidth services, or capping the amount of bandwidth available to P2P applications. However, this is likely to cause dissatisfaction among the subscriber base.

Conclusion

Given the current - and growing - popularity of P2P file sharing, service providers must adopt methods to manage the impact of P2P traffic on their networks. In order to enhance profitability and reinstate a prosperous business model, service providers must address both the issue of P2P file sharing *and* successfully manage the protocol chatter that consumes such a large percentage of network bandwidth.

Appendix

Glossary

- Client: a node that requests a service of a server, using some kind of network protocol and accepts the server's responses.
- Server: a node that provides a service for other clients that are connected to it via a network.
- Servent: a device that acts as both a client and a server simultaneously.
- Client-Server: A two-node relationship characterized by fixed capabilities where the client and server are confined to rigidly defined roles.
- Node/Peer: a device that is capable of both initiating communications and accepting communications initiated elsewhere.
- Peer-to-Peer: any relationship in which multiple, autonomous devices interact as equals.
- Time-to-Live (TTL): refers to the number of layers of peers that a request message will reach, therefore limiting the span of a peer-to-peer network.
- Network Topology: the pattern of interconnection between nodes in a network.
- Protocol Chatter: a number of messages that are intended to keep a network connection alive over a period of time, such as queries, query hits, pings and pongs.
- Bandwidth: the transmission capacity of an electronic line such as a communications network, computer bus or computer channel. It is expressed in bits per second, bytes per second or in Hertz (cycles per second).
- Message: the entity in which information is transmitted over the network. Sometimes the word "packet" or "descriptor" is used with the same meaning.

References

- Slyck, www.slyck.com
- RFC Gnutella, rfc-gnutella.sourceforge.net
- CAIDA, www.caida.org
- Peer-to-Peer Central, www.peertopeercentral.com
- OpenP2P, www.openp2p.com