

TAMING THE PEER TO PEER MONSTER USING SERVICE CONTROL

Michael Ben-Nun
P-Cube Inc.

Abstract

This document explains the increasing bandwidth and network capacity planning challenges peer-to-peer file exchange applications cause Internet Service Providers. It discusses how Service Control – the concept of statefully tracking network usage and enforcing advanced subscriber, application and destination differentiated policies – is key to resolving the peer-to-peer traffic issues within existing network infrastructure.

(e.g., Napster). Completely decentralized P2P has no central server (e.g., Gnutella) to provide search capabilities due to the fact that the clients search amongst themselves. Other variations of P2P provide application specific networks (e.g., KazaA) and some utilize an open standard (e.g., Gnutella and OpenNAP) to allow clients share all sorts of content. All of these applications allow individual users (conveniently shielded by the anonymity of the network) to share files over the Internet. These files often contain copyrighted materials (e.g., songs, movies, software, etc.) that no commercial content provider could legally afford to publish.

PEER-TO-PEER AFFECT ON NETWORK CONGESTION

The Evolution of Peer-to-Peer (P2P)

Understanding the relatively short history of P2P applications and its underlying technologies is critical to the comprehension as the impact it has on broadband IP networks. Internet based P2P is a relatively new technology, which allows for the creation of decentralized, dynamic, and anonymous logical networks for information exchange using the public Internet. In “traditional” client/server model a well-known source provides content and information to requesting clients, whereas in P2P, applications utilize various techniques to allow users to search and share content between themselves. There are several different P2P technologies and architectures that evolved from the most basic type – one that has a central “coordinating” server utilized for content searches between clients


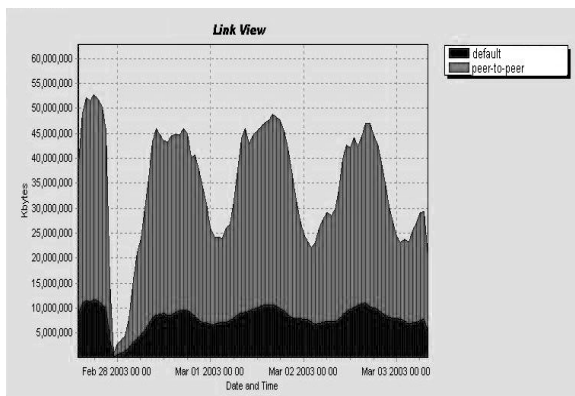
Due to this simple file sharing method, Napster, which is considered to be the first P2P application with mainstream appeal, was an immediate success among Internet users, especially those with high-speed Internet connections. A court ordered shutdown of the Napster service did little to decrease the amount of P2P file swapping activities, rather it can be argued that the added publicity probably achieved the opposite effect and the popularity of P2P applications has increased ever since. With new P2P clients and applications released to provide more functionality and ease of use, P2P traffic comprises a large part of Internet bandwidth usage. The popularity and use of different P2P clients is varied and can be determined by a variety of factors. Some clients are more popular in certain geographies (such as Winny which has wide spread acceptance in Japan), while others have a strong following among the “distributors” of specific types of material.

Peer-to-Peer Incurred Congestion

P2P clients, due to their numbers and intensive need for network bandwidth are causing significant network congestion. With less bandwidth left for other network traffic, this results in a reduction of the overall broadband experience for other subscribers on the network, and raises network capacity, planning, and management issues. Every IP network is built with assumptions about usage, which in turn is used to analyze and compute the necessary amount of network capacity and resources needed to support a given subscriber base. P2P applications are different from traditional client/server applications in the way that users run them and how the applications use the network. The table below provides a glimpse of some of the parameters used by service providers, their importance for planning the network,

and the influence P2P technologies have on these parameters. P2P applications are increasing in popularity and constitute a growing percentage of network traffic. These applications are so popular that a new term has been coined to describe the more avid users of these technologies. Often referred to as “bandwidth hogs” or “abusive subscribers,” these users are using their broadband network connections to generate a disproportional amount of network traffic and significantly contributing to network congestion.

The following charts, produced from analyzing the usage of a particular network, serving HSD cable subscribers uncovers the alarming truth: Approximately 70% of network bandwidth is being used by P2P applications.



CONTROLLING PEER-TO-PEER TRAFFIC: TECHNICAL REQUIREMENTS

With the growing amount of P2P traffic, there is a clear need to address the link congestion and bandwidth issues it creates. To solve the problem, service-providers must use a solution that is able to:

- (a) Identify, account and report on P2P usage.
- (b) Control the bandwidth these applications consume.

The following section provides detailed technical requirements that a solution must provide.

Technical Requirements

When attempting to identify and control P2P traffic, it is important to remember the underlying technical requirements from a proposed solution. Once the requirements are fully understood they could be used to evaluate possible solutions. The unique technical requirements that need to be addressed are:

IDENTIFY:

- Ability to classify traffic based on layer3-7 parameters: Peer-to-peer applications do not utilize well-known port numbers, and thus cannot be classified by simply looking at IP packet headers (IP addresses, TCP port-numbers, etc.). Rather, deep inspection of packets, including the identification of layer-7 patterns and sequences *must* be supported.
- Ability to maintain bi-directional flow state: In order to identify a particular flow of packets as peer-to-peer, carriers cannot inspect each packet within that flow to make the identification. The solution that performs proper identification of P2P traffic *must* ensure that once a particular flow (e.g. a TCP connection between two hosts) is identified as P2P, all packets on that flow are tracked, and treated as such. Of critical importance is the ability to tie between both directions (i.e. upstream & downstream) of a flow, since in many cases the initial identifying pattern resides in a packet sent from one host, yet the majority of traffic can flow in the other direction.
- Ability to provide quick turn-around for new P2P applications: As peer-to-peer applications constantly change, and

new ones emerge, the underlying protocols used to carry the peer-to-peer traffic change frequently. The solution *must* be quick to adapt to new protocols, and provide new identification mechanisms.

Note that the importance of the above-mentioned identification requirements increase in complexity and number with the growing speed of the development of new peer-to-peer applications/protocols. Even today, P2P applications use well-known ports, assigned to other network uses (such as port-80 for web-browsing), and they are constantly migrating to these port numbers in an attempt to masquerade as ‘traditional’ network activities and thereby avoid detection. Hence, simple analysis based on port-numbers leaves most of the P2P traffic unaccounted for, and will not truly address the problem.

CONTROL:

- Ability to control bandwidth at various isolation levels & granularities: To control the bandwidth impact of P2P applications it is necessary to provide a network control mechanism for different levels of isolation and control. The solution *must* provide the means to control bandwidth at “subscriber granularity”, whereby it limits the total amount of bandwidth each subscriber can consume. It *must* be able to control the bandwidth of particular flows, so as only the P2P identified traffic of a particular subscriber is limited, while the rest of that subscriber’s traffic is left unaffected.
- Ability to enforce time, destination and subscriber differentiated policies: To control the bandwidth congestion cause by P2P, and enforce various control policies, while maintaining the necessary flexibility to actually implement these on real-life subscribers, the solution *must* provide the

means to create differentiated enforcement schemes (or policies) based on time of day, destination and subscriber. Specifically, the ability to create different enforcement packages for different subscribers *must* be supported.

- Ability to maintain subscriber level quotas: In order to control P2P traffic in a persistent manner for each subscriber, the solution *must* provide the infrastructure to maintain a usage state for subscribers, and account for the total amount of P2P traffic over time. As an example, the ability to maintain the total amount of P2P traffic each subscriber has consumed on a daily/weekly/monthly basis, and apply different bandwidth quota based consumption restrictions based is key to moderating the use of the network.
- Note that while the issue of controlling and enforcing P2P bandwidth consumption is crucial for maintaining a congestion-free and predictable broadband network, it can cause customer expectation issues, as the current subscriber-base is unaccustomed to imposed limitations on its high-speed data access. Therefore the above flexibility is mandatory as service-providers create the policies best suited for their subscriber-base.
- Support high-speed network rates, and subscriber-capacities: As today's broadband networks are built to sustain significant traffic loads, the solution *must* support today's network interfaces and traffic rates. Typical broadband networks use Gigabit Ethernet and OC interfaces with high throughput. In addition, the solution *must* have the

capacity to support the total number of subscribers served by the network links, for both existing subscriber numbers today, and for forecasted growth.

APPROACHES TO CONTROLLING PEER-TO-PEER TRAFFIC

With the technical requirements in mind, the following section explores possible solutions to identifying and controlling peer-to-peer traffic.

Using Router/Switch QoS Mechanisms

Existing routers, switches or similar network devices contain various types of traffic classification and QoS mechanism, which could potentially be used to control P2P bandwidth.

However, as these devices were not designed to address these issues, they do not provide the following capabilities:

- They do not provide Layer 3-7 traffic classification. Nor do they maintain state across packets flows.
- They are not “subscriber-aware” and cannot provide subscriber differentiated enforcement

As a result, switches and routers do not provide the means by which the peer-to-peer traffic can be identified, and network usage policies be applied to it. Additionally, as the QoS mechanisms in switches and routers attempt to deal with link congestion and bandwidth distribution, they do not provide the necessary subscriber-differentiated policies, required to control the peer-to-peer traffic once identified.

Using DOCSIS 1.1

The DOCSIS 1.1 specifications, contains many features and capabilities to control bandwidth utilization, and offer differentiated services to subscribers. However, by itself the DOCSIS 1.1 specifications cannot fully address the issue of controlling P2P applications. This is due to the fact that DOCSIS 1.1 does not:

- Provide the mechanisms to classify traffic based on layer-7 capabilities, or maintain state for bi-directional network flows.
- Provide the required bandwidth control isolation and granularities. DOCSIS 1.1 provides the means to control traffic at a defined flow specification (typically a combination of layer3-4 parameters). However, as mentioned above, to fully control P2P bandwidth consumption, there is a need to implement various layer of bandwidth control, which the DOCSIS 1.1 specifications does not attempt to address.

As a result, while DOCSIS 1.1 is a potential key component in service differentiated high speed data networks, it does not provide the mechanisms to control the peer-to-peer abuse problem.

Using Service Control Platforms

A Service Control Platform is defined as a platform that is able maintain state for each network flow, classify it according to layer3-7 parameters, and implement various bandwidth shaping and control rules, based on the

classification of the traffic and the subscriber it is mapped to.

The following diagram depicts the internal operations of a service control platform.

On step (1), the platform classifies each packet received into a stateful, bi-directional flow.

On step (2), the platform performs dynamic stateful reconstruction of the application (layer-7) message exchange in the flow, and identifies the application used by each (peer-to-peer, web, mail, etc.)

On step (3), the platform maps each such flow into a particular subscriber. Typically there is a many-to-many relationship, in which many application-flows are mapped to many subscribers.

On step (4), once the traffic has been classified, identified and mapped, it is accounted for on a subscriber basis. Subscribers' state is updated according to the traffic they transmit or receive, and this impacts (along with the their assigned policies) the final bandwidth enforcement policy (5) applied.

On step (6), the selected policy is translated into packet level decisions, indicating how the actual implementation of the bandwidth restriction is performed.

Ultimately the total bandwidth consumed is reduced through control implemented in the service control platform and the overall network congestion is reduced to a level acceptable to the network provider.

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts



Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research



With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips



Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.