

Meeting the Challenge of Today's Evasive P2P Traffic

Service Provider Strategies for Managing P2P Filesharing

An Industry White Paper

Copyright © September 2004, Sandvine Incorporated
www.sandvine.com

408 Albert Street
Waterloo, Ontario
Canada
N2L 3V3

Copyright © September 2004, Sandvine Incorporated

Executive Summary

As its popularity has risen, Peer-to-peer (P2P) technology has transformed broadband Internet business models worldwide. Accounting for up to 60% of total traffic on a typical residential network, P2P congestion consumes bandwidth at a disproportional rate, swelling transit costs and opening a network up to a deluge of non-revenue-generating traffic.

And while these escalating costs eat away at service providers' bottom-lines, subscribers are becoming more elastic to price and quality of service than ever before. Reducing subscriber churn by maximizing each subscriber's experience is becoming a necessity in the increasingly competitive broadband market.

P2P technology has become a driver of broadband adoption; for the most part however, service providers have been unable to completely harness this revenue-generating opportunity due to the high level of associated cost. The goal is to effectively manage P2P traffic and its costs while continuing to maximize the subscriber experience.

In order for broadband service providers to reach this goal, a variety of solutions have been proposed:

1. Acquire More Bandwidth
2. Block P2P Traffic
3. Utilize Network Caching
4. Implement Bandwidth Caps
5. Shape P2P Traffic
6. Introduce Stateful Policy Management

This document examines these solutions, evaluating each in light of the pressing need to simultaneously reduce costs and preserve the subscriber experience.

The Problem with P2P

A General Overview of P2P

P2P ('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.

As the speed and frequency of personal computers (PCs) increased, as well as the speed and frequency of Internet connections, so did public demand for file sharing technologies. P2P has since emerged as the dominant component of bandwidth used by residential Internet subscribers. The evolution from Napster to Gnutella to KaZaA; and now eDonkey and WinMX in Europe, has dramatically increased the amount of data transferred across service provider networks.

The Implications of P2P for Service Providers

The dominating use of P2P carries with it significant ramifications for service providers; on the one hand, P2P is a major driving force in the adoption of broadband. On the other, increased levels of P2P traffic result in a wide range of disadvantages for service providers and subscribers. The disadvantages of unrestrained P2P usage can be grouped into two categories: additional costs for service providers, and degradations in the subscriber experience.

Additional Costs for Service Providers

Costs are a primary concern to service providers. Below are a few of the many costs associated with unrestrained P2P traffic.

- Costly bandwidth consumed - on a typical service provider network, over 60% of total bandwidth is used by P2P traffic. This traffic is comprised of "protocol-chatter" as well as the transmission of the shared files themselves.

- Additional network transit costs - P2P traffic connects in an ad-hoc fashion; subscribers are as likely to download a file from halfway around the world as they are to download it from their neighbor.
- Over-subscription business model undermined - a common business model among service providers, over-subscription is unworkable when a large proportion of subscribers are consuming bandwidth 24-hours a day, 7 days a week.
- Loss of brand equity - in today's competitive broadband industry, a congested service provider network translates into churn as subscribers "vote with their feet".

Degradations in the Subscriber Experience

Acquiring and maintaining a strong customer base is a must in today's competitive broadband market. Below are a few of the negative effects that P2P has on a service provider's subscribers.

- Slower transfer speeds - as bandwidth is consumed by P2P traffic, less is available for other uses
- Unequal distribution of bandwidth - a disproportionate number of subscribers use up a large proportion of the total bandwidth with P2P; this traffic crowds out other legitimate traffic, decreasing the performance of:
 - "Time-critical" applications (streaming video, audio, VoIP)
 - Internet browsing (HTTP)
 - E-mail (SMTP)
- Overwhelmed network infrastructure - a congested network is more likely to fail, as flow-based equipment struggles to perform their routine tasks. When a strained network inevitably fails, subscribers suffer service outages as the service provider scrambles to (temporarily) fix the problem.

Any truly successful solution to the P2P problem must address both the service provider perspective of costs and the subscriber perspective of quality of service. The disadvantages of P2P traffic listed above are considerable, but not impossible to mitigate. The remainder of this document will examine the proposed strategies for dealing with the problems of P2P.

Proposed Service Provider Solutions

In order to effectively manage the P2P phenomenon, a variety of solutions have been proposed:

1. Acquire More Bandwidth
2. Block P2P Traffic
3. Utilize Network Caching
4. Implement Bandwidth Caps
5. Shape P2P Traffic
6. Introduce Stateful Policy Management

Each of these proposed solutions are evaluated below.

1. Acquire More Bandwidth

When a network is regularly overwhelmed with traffic, a common approach is to obtain more bandwidth by purchasing it from a larger provider and upgrading the existing infrastructure to handle the increase. To a certain extent, this is logical: if the present amount of bandwidth is not enough to handle traffic volumes, then additional bandwidth is required. And if the service provider is in a growth phase, then a solution that facilitates that growth is appropriate.

However, while acquiring more bandwidth and building up infrastructure does provide more bandwidth, it does nothing to mitigate the problems associated with P2P. In fact, the increased amount of bandwidth actually encourages bandwidth abuses, as the offending subscribers have increased resources to consume; the more that is provided, the more is consumed, while the associated costs of P2P only increase.

That said, this solution could delay substandard network performance; initially, subscribers may notice faster download rates and more consistent service. But as P2P usage continues to expand, this costly additional bandwidth will be increasingly consumed by P2P traffic. In the end, *upgrading* the network is not a replacement for *optimizing* the network.

This solution also comes with a prohibitively high price tag, as service providers are forced to invest significant amounts of money in the purchased bandwidth, the new equipment, and the many hours spent integrating it all with the existing system. Since this solution also leaves the problems associated with P2P untouched, the cost-to-benefit ratio becomes far too high for it to be a realistic way to manage P2P traffic.

The Verdict

Acquiring more bandwidth and upgrading equipment is an exceedingly costly way to address the P2P problem. The subscriber experience may initially improve, but since nothing is done to limit bandwidth abusers, it will eventually worsen for the same reason it did before - unmanaged P2P traffic. Far from reducing costs, this solution multiplies them many times over.

Because of the high-cost, low-benefit nature of this solution, it fails to fulfill the criteria of reducing costs and improving the subscriber experience.

Acquiring more bandwidth as part of a growth strategy is a necessary step; acquiring it to manage P2P is untenable.

2. Block P2P Traffic

P2P blocking refers to the practice of blocking ports at the Network Access Point (NAP) that are commonly used by the most popular P2P networks. The aim of P2P blocking is to reduce bandwidth usage by arresting all P2P traffic, and in so doing, completely avoid the typical costs of P2P usage.

In practical terms, P2P blocking can be very problematic: many clients allow users to select a desired port or assign ports dynamically, with the express purpose of circumventing standard P2P blocking attempts.

These technical difficulties in blocking P2P traffic can be overcome, but there is a far more important concern: P2P file sharing is a significant driver for broadband Internet

adoption. Blocking all P2P traffic is certain to lead to customer dissatisfaction and aggravate customer churn. In fact, some service providers are beginning to tout their high-speed services as “P2P Friendly”, leveraging their P2P position into a powerful marketing tool, capturing the interest - and wallets - of frustrated subscribers.

The Verdict

Costs are reduced when P2P traffic is blocked, but so are revenues and brand equity; network performance is increased, but so is subscriber churn. These tradeoffs are simply not good enough. It is clear that this short-term solution brazenly ignores marketplace indicators.

While this solution reduces the costs of P2P to zero, it also dramatically harms the subscriber experience, and so falls well short of the established criteria.

As P2P continues to thrive, service providers block P2P traffic at their own peril.

3. Utilize Network Caching

A P2P network cache allows a service provider to maintain a cache, or store, of the most frequently accessed P2P files. When a subscriber performs a file search, this centralized store of popular P2P files is accessed first. If the file is present in the store, the file can be retrieved directly from it, thus reducing incoming traffic over the service provider’s NAP and reducing the bandwidth consumed by downloads. The cache is automatically seeded; when a user requests a file that the cache does not have, it makes the connection to the source of the content and retrieves the file, simultaneously storing it on its local drive and sending it to the requesting user.

This solution minimizes some of the downstream bandwidth and transit costs associated with P2P traffic since traffic is kept on the local network as much as possible. It also leaves the subscriber experience unharmed.

Despite these advantages, network caching remains a controversial approach. The foremost concern for service providers is the legality of such a solution; the provider would no longer be merely providing the technology, but free copyrighted content as

well. In the eyes of active media lobby groups and increasingly irate music organizations (the RIAA, for example), this would push the service provider into copyright violation. From lobby group harassment, to civil lawsuits, to state legal action, network caching carries with it a disreputable host of agitations and problems that could negatively impact the public image of a company worldwide.

Another significant problem of network caching is that it does nothing to alleviate upstream congestion. P2P users from other networks still consume bandwidth as they download files from the service provider's subscribers - network activity that provides no value to either the provider or the subscriber.

The Verdict

On the surface, networking caching seems to be a workable solution, allowing for reduced downstream costs and an unharmed subscriber solution. Lurking beneath the surface however, are a host of legal issues and a mass of bad PR. Until some of the legal questions are definitively answered, a service provider choosing network caching exposes itself to a range of serious risks.

While network caching superficially satisfies the solution criteria, legal issues and negative PR effects prevent it from being feasible.

4. Implement Bandwidth Caps

A bandwidth cap relates to the aggregate total of a subscriber's bandwidth activity over a certain period of time. This cap, usually stated in gigabytes (GB), actually serves as the base breakpoint in a tiered pricing scheme.

Tiered pricing schemes allow service providers to offer varying levels of service; for the casual Internet user, the lowest level of service is likely suitable, providing them with enough monthly bandwidth to surf the Internet and use email. The high-traffic P2P user would require the highest tier, and would pay a higher rate for it. By charging different prices for each tier of service, service providers recoup some of the additional costs that are incurred by the heavy-traffic users.

Although bandwidth caps discourage most subscribers from using more than the prescribed cap of ‘n’ GB per month, they offer limited reductions in bandwidth per second, and thus may not affect a service provider’s monthly bandwidth charges. Upstream traffic is left untouched, and so continues to congest the network and incur unnecessary costs.

In addition to this, bandwidth caps cannot deal with P2P traffic specifically, as they only deal with the aggregate bandwidth activity of a subscriber. These caps lack the granularity to mitigate P2P as a separate phenomena, and must use broad-level restrictions to accomplish their goal - a far from ideal method. In essence, this solution merely sets the maximum boundary for P2P costs, but does nothing to mitigate the cost triggers themselves.

Perhaps the most important objection to bandwidth caps is that subscribers do not understand “bits and bytes” - they only understand what they want to do online, and are confused and frustrated when their service is limited by something that doesn’t make sense to them. And an end user running a P2P client does not generally have direct control over how much is uploaded from their system, and thus runs the very real risk of using up all of their available bits by unknowingly serving external P2P users. This results in either penalty charges or cut off service, two punishments, which, in the subscriber’s mind, don’t seem to fit the crime. Subscriber frustration is likely to aggravate customer churn, especially if other service providers are not implementing similar caps.

Today’s broadband market is usually characterized as an “all-you-can-eat” industry, with subscribers paying a flat rate for unlimited use. In light of this, adopting bandwidth caps as a P2P mitigation strategy is to risk losing one’s subscriber base.

The Verdict

Bandwidth caps are a heavy-handed and imprecise approach to the P2P problem. While they do reduce some of the overall costs of P2P, they do it by infringing on what subscribers see as their right to unlimited Internet access. Subscriber churn increases and the long-term profitability of the provider is put in jeopardy.

While this solution meets the criteria of reducing associated P2P costs, it drastically harms the subscriber experience, and so fails the second part of the criteria outlined at the beginning of this document.

In an “all-you-can-eat” broadband world, bandwidth rationing leaves subscribers frustrated and unsatisfied.

5. Shape P2P Traffic

Shaping refers to the practice of processing, buffering, and prioritizing all traffic traveling through the NAP. This allows a service provider to give priority to non-P2P traffic, leaving whatever bandwidth is left over for P2P. Each individual data packet that arrives at the NAP is examined and classified based on an identification key found in the packet. Based on the priority of each category of traffic, the packets are then entered into a queue and transmitted. In a P2P-shaping context, P2P packets are sent last, consuming whatever bandwidth is left over after all the higher priority traffic has been sent.

Shaping certainly has its advantages; a service provider can gain a degree of control over their network, and the ability to prioritize their traffic to suit their subscriber base and cost concerns is a useful tool. And associated P2P costs can be reduced in a way that avoids the sizeable pitfalls of completely blocking P2P traffic.

However, because shaping relies on accurately identifying packets as P2P, it is susceptible to a range of evasion tactics implemented by P2P developers. The P2P development community has shown itself to be very resistant to shaping techniques in the past, and has developed several tactics for hiding the true identity of packets. These include:

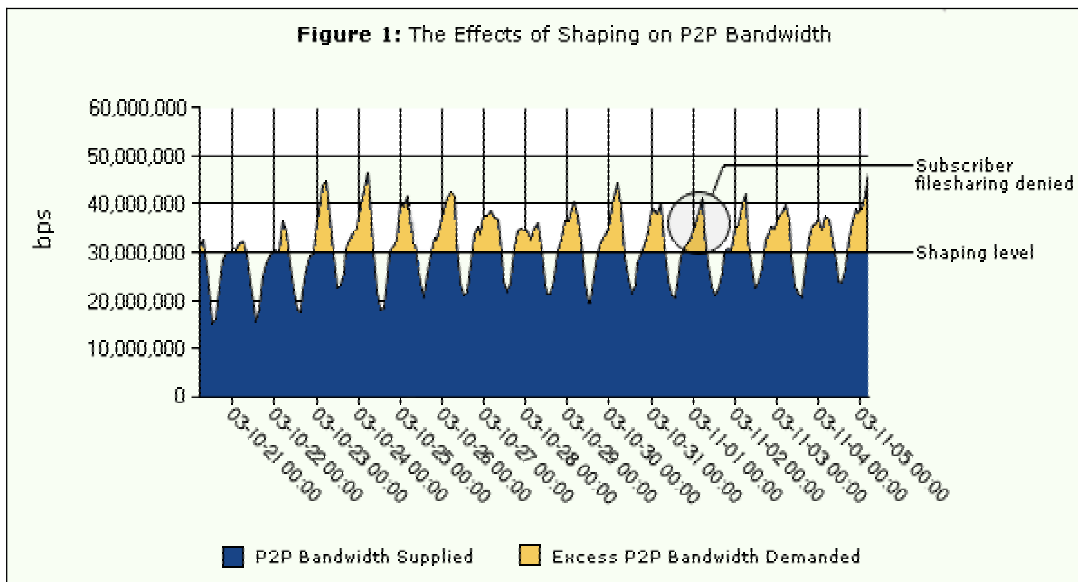
- Variable offsets of the packet identifier - the identifier is not always in the same place in the packet all the time
- Mathematical transpose of packet information - a mathematical operation is used to mask the true nature of the packet
- Identifier concatenates information across connections
- Identifier concatenates information across multiple packets - the identity of a packet can only be known by examining several packets (i.e. in context).
- Identifier artificially split by fragmenting packets.

Each of these strategies must be countered in order for shaping to remain effective at managing P2P traffic.

Even if shaping could identify P2P packets accurately 100% of the time, there is a more significant disadvantage to address: shaping dramatically harms the experience of all subscribers on the network - not just the P2P users.

Shaping requires that each data packet be inspected to determine which are P2P. This introduces significant network latency and slower processing times for all traffic, as each packet is queued, inspected, and then queued again for transmittal. Shaping therefore indirectly penalizes non-P2P users for the P2P traffic that is on the network; every user on the network suffers these lags. This damages the subscriber experience immensely. High network latency erodes the performance of services that require real-time performance; low quality VoIP, jittery streaming video, and online gaming hang ups leave subscribers scratching their heads and looking for alternatives.

If shaping punishes non-P2P users, then P2P users are given an even more severe penalty. P2P performance suffers dramatically under a shaping scheme, the foreseeable outcome of giving P2P traffic the lowest priority. As mentioned numerous times, P2P is driving broadband Internet adoption, and subscribers expect to get the speeds they paid for. When faced with dial-up performance on a supposed high-speed subscription, users start asking themselves why they are paying \$x more for the same performance. And so in response, service providers must ask themselves: are the gains made through choking P2P worth the long-term costs of losing previously loyal customers?



As illustrated by Figure 1, shaping results in subscribers being unable to use P2P during peak times. On the graph above, the shaping level at 30,000,000 bps is the upper limit of P2P allowed. The orange regions represent subscribers who wanted to use P2P but could not due to the shaping restriction. These regions represent the frustrations and confusion of denied subscribers; these regions represent reasons for switching service providers.

The misclassification of packets is another subscriber-oriented issue to consider. If a shaping device misclassifies some packets, they may not receive the priority they require, or may not be sent at all. This requires the shaping device to be able to identify all types of traffic accurately. At best, subscribers suffer some inconvenience when their packets are misclassified; at worst, key packets aren't sent at all, and the subscriber is left confused and frustrated.

The Verdict

Shaping is an interesting approach to the P2P problem. It provides service providers with some useful tools to control their network, and it can be used to minimize the effects of P2P. Unfortunately, it also severely damages the experience of all subscribers. P2P traffic is indeed minimized, but at the very high price of slower processing times for each and every subscriber, and abysmal performance for P2P users themselves.

Because it abjectly fails the subscriber-friendly aspect of the decision criteria, shaping is not an effective P2P managing solution - despite the reductions in cost that it offers.

6. Introduce Stateful Policy Management

Stateful Policy Management is a technique that manages P2P traffic both on the downstream and on the upstream. On the downstream, a redirecting agent reroutes P2P traffic along the least-cost network path, while P2P session management manages the upstream bandwidth by controlling the number of P2P connections with external networks.

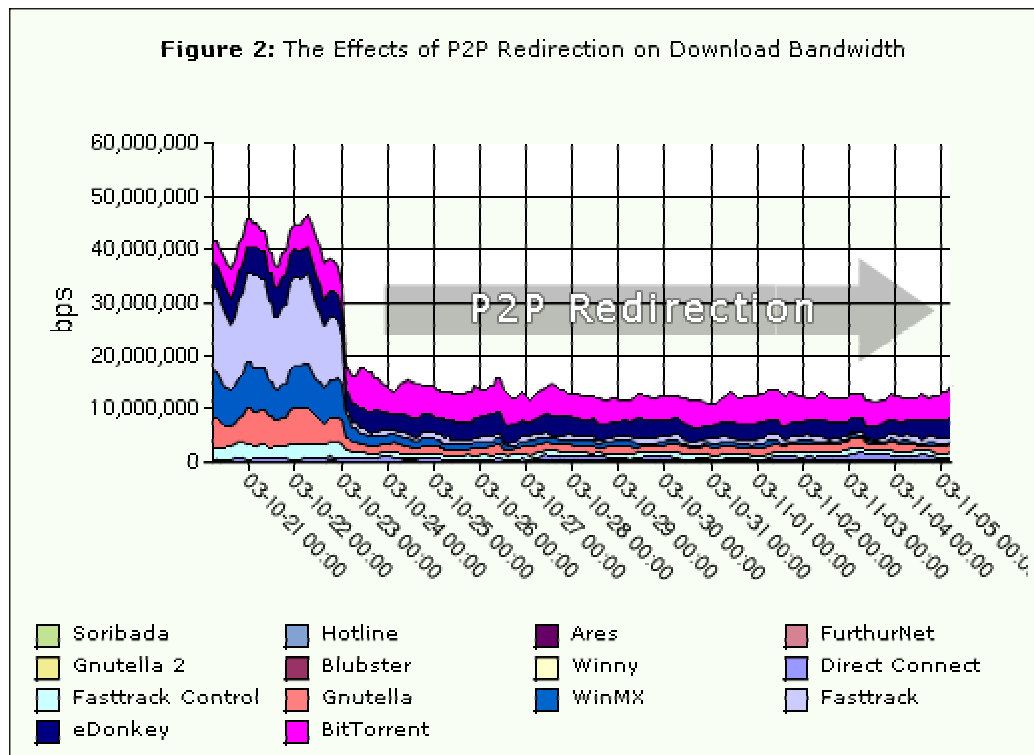
Stateful Policy Management is accomplished through the use of stateful, deep-packet inspection. The term “stateful” refers to an application and connection where information must be kept track of over time, as specific information pertinent to the communication is important for both the identification of the protocol and the subsequent actions and operations to be taken on the protocol. Stateful inspection is crucial to intelligently classifying and routing traffic, because it provides an overall understanding of the conversation between two clients. Packets are also deeply inspected and classified, which identifies the nature of the packet flow and the control plane information.

This Layer-7 application awareness facilitates the redirection of the P2P traffic flow. The redirecting agent keeps track of all control plane traffic, which typically consists of queries, send requests, termination requests, and lists of shared content. Because it is essentially ‘listening in’ on the P2P conversations, the redirector can step in and facilitate a transfer among local subscribers, rather than allowing the P2P protocol connect randomly to a client on an external network.

When a subscriber attempts to download a file, the redirecting agent first checks to see if any other subscriber has the requested file. If there is such a subscriber on the network, the redirector facilitates the connection of the requester and the subscriber who has the file. The file is shared among the service provider’s own subscribers, avoiding transit costs and congestion at the NAP point. If another subscriber is not sharing the file, then the file request gets sent out to an external network as per usual. This process is completely transparent; the subscriber is entirely unaware that it is occurring, and there is no delay introduced or any other performance aberrations.

Importantly, the total number of actual downloads is completely unaffected; Stateful Policy Management does not depend on limiting subscriber usage, but rather on the intelligent rerouting of the P2P traffic itself. And the cost savings are substantial.

Downstream traffic volume across the transit links is significantly decreased as redirection routes files along the least-cost network path (see Figure 2). Network congestion at the NAP point is also minimized, freeing up bandwidth for other uses.



Redirection minimizes the amount of P2P download bandwidth considerably without any impact on the subscriber experience. Unlike shaping, which can leave subscribers without P2P capability during peak hours, redirection does not infringe on any subscriber’s ability to use P2P whenever desired.

Acting as the facilitator of P2P conversations, the redirecting agent can also limit the number of P2P connections made to external users through the use of session management. Session management is a capability that allows service providers to control exactly how many P2P sessions are allowed with external users, and subsequently, how much upstream bandwidth is consumed by P2P traffic. If the maximum number of sessions is reached, the redirecting agent turns away requests from external users, forwarding them back out to external networks. And because it is customizable, session management allows service providers to save anywhere from 0% up to 100% of all upstream traffic - whatever rate of savings is desired.

From a cost perspective, Stateful Policy Management reduces downstream bandwidth consumption, transit costs, and up to 100% of all P2P upstream bandwidth. Out of each of the solutions examined thus far, Stateful Policy Management provides the most cost

savings. It is also the only solution to reduce the three types of costs: upstream bandwidth, downstream bandwidth, and transit fees.

Aside from the cost savings, another significant advantage of Stateful Policy Management is that non-P2P traffic is unaffected; unlike traffic shaping, where all traffic is processed, buffered, retained, and transmitted by the router, Stateful Policy Management only impacts P2P traffic.

By far the most important aspect of Stateful Policy Management is that the process is completely subscriber friendly. Featuring no delay time for any subscribers and no decreases in the number of downloads or download performance for P2P users, Stateful Policy Management manages P2P completely transparently; subscribers have no indication of what is happening - they are only aware of the improved network performance. Their choice is not limited, their performance is enhanced, and all of the pitfalls of having a subscriber-negative P2P policy are avoided - all while the service provider banks tangible savings.

Stateful Policy Management accomplishes the seemingly disparate goals of reducing P2P costs while improving the subscriber experience.

The Verdict

By intelligently inspecting traffic, service providers gain an understanding of what kinds of traffic are traveling on their networks. Once P2P traffic has been identified, it can be redirected along the least-cost path of the network and managed at the session level, decreasing both upstream and downstream bandwidth consumption, transit costs, and most importantly, subscriber frustration.

Stateful Policy Management is the only solution examined that transparently reduces costs while at the same time maximizing the subscriber experience.

Conclusion

Subscribers continue to cite P2P as a primary reason for their broadband connection, and are also increasingly sensitive to anything that adversely impacts their online experience. With a wide range of service providers to choose from, subscribers are becoming more and more willing to change who provides their Internet connection; in short, Internet access is becoming a commodity based on performance, price, and stability. In light of this, managing P2P effectively, and in a subscriber-friendly way, has become a necessity for service providers.

A variety of solutions that attempt to mitigate the costs of P2P and ensure that subscribers have the best possible experience have been proposed: acquiring more bandwidth, blocking P2P traffic, utilizing network caching, implementing bandwidth caps, shaping P2P traffic, and introducing Stateful Policy Management. Out of these, introducing Stateful Policy Management is the only one that fulfills the dual roles of reducing costs and protecting the subscriber experience. In fact, Stateful Policy Management features higher cost savings than the other options, and actually improves network performance - a powerful combination of abilities that can allow a service provider to differentiate their services from others, while realizing significant improvements to their bottom-lines.

As P2P continues to rise in popularity, service providers must find an effective way to manage it - or suffer the consequences of overwhelming costs, an unstable subscriber base, and poor brand integrity.

Solutions	Reduce Costs?	Subscriber Friendly?	Improve Brand Integrity?
Acquire more Bandwidth	✗	✓	✓
Block P2P Traffic	✓	✗	✗
Utilize Network Cache	✓	✓	✗
Implement Bandwidth Cap	✓	✗	✗
Shape P2P Traffic	✓	✗	✗
Introduce Stateful Policy Management	✓	✓	✓