




Controlling Peer to Peer Bandwidth Consumption



As Peer-to-Peer (P2P) file exchange applications gain popularity, Internet service providers are faced with new challenges and opportunities to sustain and increase profitability from the broadband IP network. Due to the unique and aggressive usage of network resources by Peer-to-Peer technologies, network usage patterns are changing and provisioned capacity is no longer sufficient. Extensive use of Peer-to-Peer file exchange causes network congestion and performance deterioration, and ultimately leads to customer dissatisfaction and churn. This paper discusses the unique problems associated with the growing popularity of P2P applications and how they affect the IP network. It then discusses various remedies that the P-Cube Service Control Platform provides to address these problems. From detailed usage accounting of Peer-to-Peer traffic, to facilitating monitoring and billing based on this information, to the creation of smart control policies, the Service Control Platform's high level of data abstraction and presentation helps ISPs manage these business challenges while avoiding alienation of the customer base with policies that are perceived as intrusive or "unfair."



White Paper

A Brief History of Internet Peer-to-Peer Technologies

Internet Peer-to-Peer (P2P) is a relatively new technology, which allows the creation of decentralized, dynamic and anonymous logical networks for information exchange over the public Internet. As opposed to “traditional” client/server models in which a well-known source provides content and information to requesting clients, Peer-to-Peer applications utilize various techniques to allow users to search and share information and content between themselves. There are several different P2P technologies and architectures, such as those with a central server used to coordinate and conduct searches (e.g., Napster), or those that are completely decentralized in which no central node exists (e.g., Gnutella) and some different levels of interoperability, ranging from application specific P2P networks (e.g., iMesh) to those utilizing an open standard (e.g., Gnutella and OpenNAP). 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.)— materials no commercial content provider could legally afford to publish.

Due to this simple file sharing, Napster (commonly seen as the first P2P application with mainstream appeal) was an immediate success among Internet users, especially those with high-speed Internet connections. The famous court ordered shutdown of the Napster service did little to decrease the amount of P2P file swapping activities (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 clients and applications released on an ongoing basis. The popularity of P2P clients often changes: some are more popular in certain geographies (such as WinMX, which, due to its support for Unicode character sets has become very popular in Japan and other parts of Asia), while others have a strong following among the “distributors” of specific types of material. Following are some current popular clients and applications:



The Peer-to-Peer Challenge for Service Providers

While the popularity of P2P applications is causing significant challenges and issues for the rightful owners of the copyrighted material being freely distributed, they are also creating network capacity and subscriber expectation management problems for Internet service providers. Every IP network is built with assumptions about its usage, which are in turn used to analyze and compute the necessary amount of network capacity and resources needed to support a given subscriber base. This analysis is essential for service providers in their attempt to maintain a solid ROI model for their networks and to maintain some level of visibility into their future needs. As P2P applications are different from traditional client/server applications in the way they are used and in turn the way they use the network, they are, in many ways, changing the game for service providers trying to architect and maintain their networks. 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.

Parameter	Importance for network planning	Influence of traditional applications	Change caused by P2P applications
Upstream / Downstream Traffic Ratio	Networks are asymmetrical in nature: the amount of traffic that a network can sustain upstream (i.e., from the subscribers to the network), is different from the amount it can sustain in the opposite direction. The ratio required between these two directions is in direct correlation to the requirements of the applications using the network. Networks are built with a certain ratio in mind, which, if incorrect, may cause high rates of congestion and unutilized capacity.	A typical residential user uses the network for downstream applications. These applications (e-mail, web browsing, etc.) generate a larger amount of downstream traffic for a single upstream request, and service providers have come to rely on this ratio	P2P applications encourage users to share files, and a typical peer serves multi-megabytes of files. This causes a shift in the upstream/downstream ratio, and as a result congestion on the upstream link (due to individual users' increased uploading of files).
Time of Day and Percentage of Activity	Service providers typically assume an average duration of network use per subscriber per day, and (based on subscriber profiling) peak use periods. A service provider would typically be able to predict and account for network "rush hours" and less congested periods of network use. This subscriber profiling is based on assumptions that residential home users primarily use the network during weekends and at nighttime, and that telecommuters and small offices use it primarily during business hours. Sudden or sporadic changes in these patterns may cause congestion during certain hours that were not evident before.	The time of day and percentage of activity assumption for residential broadband subscribers are rooted in the premise that a typical residential user uses the network only when the subscriber is physically present and actively using the connection. Such is the case when web browsing, reading e-mails, etc.	As P2P applications are usually used to upload or download large, multi-megabyte files, they are typically left unattended for days at a time while the application constantly attempts to download a list of files on one hand, and serves multiple file requests of other peers on the other. This creates a neverending, high volume stream of network activity throughout the day. For example, a student's computer with a broadband connection can compete with telecommuters for vital network resources during business hours while the student is at school.
Traffic Destination and Peering points	The costs associated with serving each network packet and connection sometimes depends on the location of the peer of the subscriber. Carefully	Traditional uses of the data network are either mainly OnNET (email, nntp, web-	P2P traffic has increased the amount of traffic between home users in a significant way. Before the existence of

Parameter	Importance for network planning	Influence of traditional applications	Change caused by P2P applications
	crafted peering agreements with other network providers can help reduce the amount of traffic, and hence the cost of expensive transit connections. Furthermore local traffic (often referred to as OnNET) which does not leave the service provider's own backbone network, is significantly lower in cost than traffic leaving the provider's domain (OffNET).	proxies) or destined to a small number of content providers, and data sites.	this technology, two home users (whether they used the same or different providers) would almost never form a direct connection. P2P file exchange has significantly increased this interconnected traffic.
Estimated Traffic Volume	No matter the topology and architecture of the network, there is a finite amount of bandwidth available for all its users, and certain over-subscription assumptions are used when planning the capacity of the network	Traditional applications have a large "time-to-consume" factor: A small web-page can take several minutes to read, a single e-mail message might take a number of hours to process.	P2P applications are mainly used to share large binary files that have a much lower "attention-per-byte" ratio. A three-minute song is usually 3-5 megabytes. A 10-minute movie can be hundreds of megabytes long.

Where and how P2P network traffic hurts service providers depends on the service provider's network architecture and topology, as well as the subscriber base and its usage patterns. Some common examples of difficulties caused by the rise in P2P application usage include:

- ❖ Due to the physical attributes of the shared cable infrastructure, cable HSD providers are particularly limited in the amount of upstream network resources they have and need to go through a costly configuration process (fiber node splits) to expand the capacity. Since P2P applications cause a dramatic increase in upstream data, they pose both cost and maintenance challenges for cable HSD providers.
- ❖ Regional service providers are concerned by the amount of P2P traffic traversing expensive network access peering points (such as those connecting their own IP network to the Internet over an international link), whereas local file swapping (between subscribers on the same segment of their network) is of less concern.

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 (often unknowingly) using their broadband network connections to generate a disproportional amount of network traffic, significantly contributing to network congestion.

Service providers cannot afford to ignore P2P's increasing popularity (recent statistics indicate P2P accounts for approximately 60% of all Internet traffic), the changes these applications cause to network traffic patterns, and the effect these programs have on planning and infrastructure. As the usage of P2P applications increases, so does network congestion, decreasing performance for all users and applications. As the primary selling factor for broadband Internet access to date has been network speed (both in comparison to dial-up access and other broadband providers), service providers cannot afford slow or unpredictable network access. This is especially true as providers attempt to grow their customer base in the telecommuter and small business markets where predictable performance is essential.

Addressing the Challenge

Service providers must find ways to deal with these “abusive subscribers” and address the challenges posed by the aggressive nature of P2P applications. Simply adding additional network capacity is too costly and cumbersome to manage. IP networks are expensive to maintain and competition is fierce. Even with today’s allocated capacities, some providers find it difficult to maintain a solid margin and profit from subscribers. In addition, the “eat-all-you-can-get” nature of P2P technologies means any additional bandwidth allocated will quickly be consumed again, leaving strained service providers in the same situation.

Rather, service providers must:

- ❖ Find how and where P2P traffic is causing network congestion or increased expenses, and restrain its effect.
- ❖ Identify those subscribers who are consuming an unacceptably large amount of traffic (and are consequently an expensive loss leader).
- ❖ Use different subscription plans to compensate for the increased expense in carrying this network traffic.

While attempting to address the challenges posed by P2P applications, service providers must also ensure they do not **alienate** their subscribers by changing subscription plans and implementing overly severe or strict restrictions. Familiar with flat-fee, “all-you-can-use” Internet access, subscribers will need to be guided through a delicate product marketing/conditioning phase to accept new subscription models. As such, it is essential that any remedy deployed to address P2P is focused, easy to understand and not globally restrictive.

In order to address these challenges, service providers must use tools capable of:

- ❖ Identifying and classifying all P2P traffic, so that it can be accurately accounted for and controlled
- ❖ Enforcing policies that address the problems caused by the excessive P2P traffic without influencing other traffic or alienating subscribers

As simple as this may sound, these business goals are difficult to achieve with conventional means.

Identifying P2P Traffic

To understand the nature of P2P traffic on a network, and to control its usage, a service provider must first have a tool in place that can *identify* P2P-related packets and differentiate them from others.

However, many of the communication protocols used by P2P applications utilize different techniques than other communication protocols, making it extremely difficult to identify them using traditional techniques. Specifically, many P2P protocols do not use static, well known port numbers, but rather dynamically utilize available port numbers, including those typically reserved for other applications. An example of this is KazaA, which can use port 80 (reserved for http/web browsing) for its communication, to penetrate firewalls and network packet filters. This makes it impossible to identify, track or control P2P traffic by using simple port-based classification.

In addition, even as P2P applications’ popularity increases, so do the variety of communication protocols used. The mechanism used to identify and classify P2P activities must be capable of rapid and simple adaptation to the ongoing change inherent to this community and the applications they use.



Accounting for P2P Traffic

After P2P traffic is identified, it is essential that it be accounted for. This information is required for the identification of “bandwidth hogs” and is important information for service providers as they attempt to evaluate how P2P traffic affects the network. This information can also be used to fuel various usage- and consumption-based billing plans, to better align subscribers’ fees with their individual usage (and consequently their actual capacity and usage costs to the network).

In order to address scalability and reduce the cost of maintaining and managing this information, it is important that it be collected, accounted for and reported on at the appropriate level of granularity. Generating too much usage information can decrease network performance and cause an overload of usage data, making it difficult to generate useful, timely data. Overly coarse usage data, on the other hand, does not provide enough information to truly gain insight on network activities.

Smart Control Policies

It is critical that “smart policies” be implemented to control the effect P2P traffic has on network congestion and performance while limiting customer alienation or churn. Service providers should avoid aggressive policies, and instead opt for policies that reduce the strain of the P2P applications, condition and educate customers about tiered, consumption-based subscription plans, and provide new product offerings for power P2P users. These “smart policies” improve network performance while sidestepping the pitfalls of bad publicity and subscription termination. The table below provides a number of examples of policies that may be seen as too restrictive and carry the risk of potential subscriber alienation and some of the alternative “smart policies” that alleviate these issues.

Restrictive Policy	Potential Alienation or Subscriber Concern	Alternative “Smart” Policies
Block (or significantly limit) access to certain P2P applications	May aggravate subscribers using the blocked P2P applications, potentially causing bad publicity.	<ul style="list-style-type: none"> + De-prioritize P2P during congestion periods. + Throttle upstream (file upload) traffic, and do not limit downstream (file downloads). + Limit P2P access during certain periods of the day/week (e.g., business hours). + Limit P2P traffic traversing expensive peering points or transits. + Provide unrestricted subscription plans for an additional charge.
Subscriber byte caps (e.g., up to 1 gigabyte daily, followed by an additional surcharge or block of access)	May cause concern that application “unaware” byte caps do not provide the desired application isolation for the subscriber, as excessive P2P use might penalize other critical applications (e.g., VPN, email, etc.). This is especially true for cases in which one broadband connection is used by a number of different users.	<ul style="list-style-type: none"> + Provide a P2P quota, which once depleted will throttle P2P traffic, but will not affect other application uses. + Provide optional P2P “bandwidth-on-demand” for an additional charge to allow subscriber access to more bandwidth even if quota has been depleted.

Existing QoS, queuing and shaping mechanisms are insufficient to implement these “smart policies” as they do not provide the required level of control. Furthermore, as these devices cannot truly identify P2P traffic, they cannot isolate and control P2P traffic while distinguishing it from other applications such as web-browsing, VPN and email.

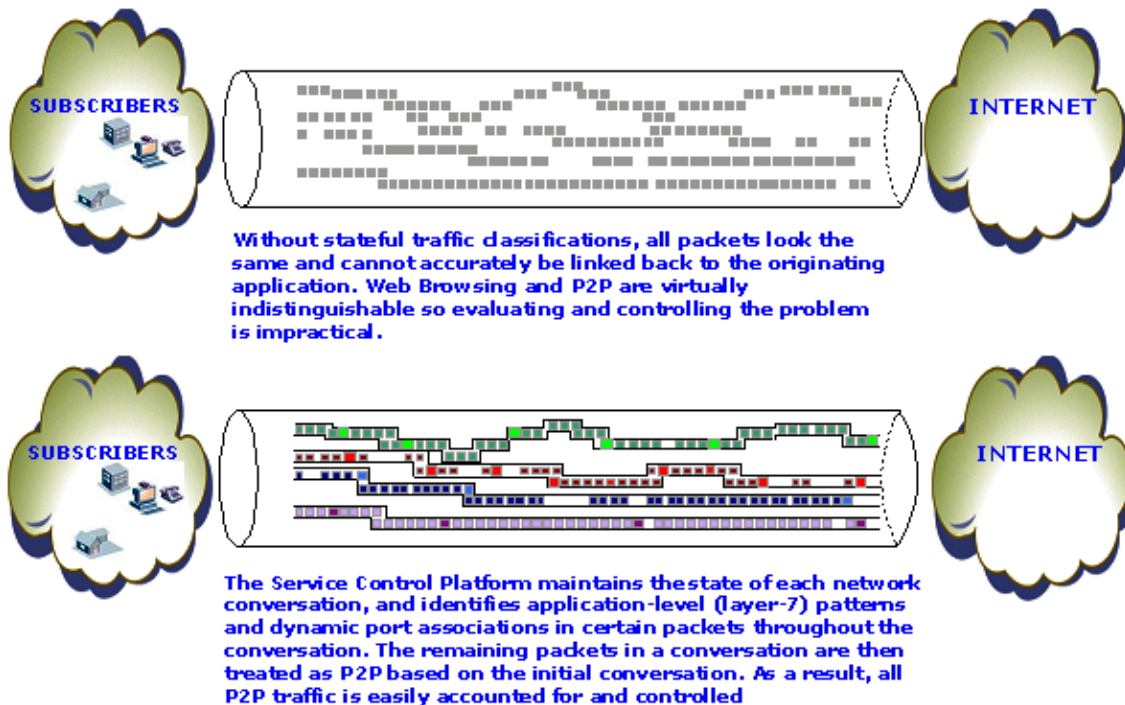


The Solution: P-Cube's Service Control Platform

P-Cube's Service Control Platform (SCP), a state-of-the-art, dedicated network device, provides the means and tools to detect and control P2P application usage and excessive bandwidth consumption. Using the platform's stateful traffic classification and analysis, the P-Cube Service Control Platform can accurately identify P2P traffic and identify "abusive subscribers". Furthermore the platform's advanced traffic management and control capabilities provide the means to contain and moderate this excessive bandwidth usage, while preventing subscriber alienation due to the implementation of overly restrictive or aggressive policies. All this due to the unique characteristics and architectural attributes of the SCP, especially designed to perform real-time traffic classification, accounting and control

Stateful Traffic Classification

P-Cube's Service Control Platform utilizes a revolutionary hardware architecture that maintains the state of each and every network conversation, while executing deep and detailed inspections of each and every data packet, up to the application layer (layer-7). As a result, the SCP can detect specific P2P application signatures usually witnessed during the initial message exchange between two network hosts and classify all traffic for that conversation as P2P.

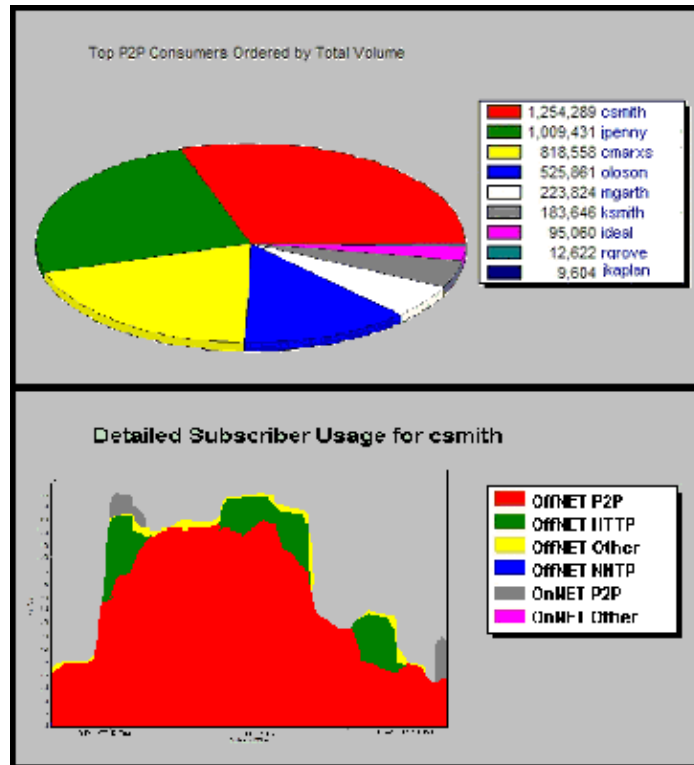


Subscriber Application Usage Analysis

The Service Control Platform generates usage statistics on each and every subscriber using the IP network for each application and protocol being used. This allows service providers to identify “abusive subscribers” in real-time.

The usage information generated by the SCP is used to produce simple to understand, detailed reports on network activities. Reports such as “Top Volume Consumers” (identifying the most active subscribers) and “Detailed Subscriber Usage” (detailing how the network was used by a particular subscriber) give unsurpassed insight into network activity and are key to gaining insight on which subscribers are abusing the network, how they are doing it, and what the best control policies to moderate their use are.

This information is critical for both understanding the usage pattern in the network and identifying those subscribers that are creating unacceptably high traffic volume.



Rapid P2P Application Identification Turnaround

The Service Control Platform’s protocol classification mechanisms are extremely flexible and programmable, facilitating rapid and reliable upgrades to classification code, in a fraction of the time it typically takes to upgrade a network device. This is important for the detection of P2P protocols, as these protocols change frequently with the release of new applications and new versions of existing ones.

Controlling P2P

The Service Control Platform's proactive traffic control and bandwidth management capabilities provide simple to use yet highly efficient means to relieve the strain caused by P2P traffic and "abusive subscribers," while enabling the creation of smart control policies to limit subscriber alienation. Following are some of the traffic control capabilities of the Service Control Platform, facilitating the creation of these "smart policies":

Aggregated Rate Limiting

Limit all P2P traffic to a certain percentage of the available bandwidth. While this does not provide fairness between subscribers, it can be used to eliminate the performance degradation caused by excessive P2P traffic to other network activities (e.g., VPN, browsing, streaming, etc.).

EXAMPLE:

Limit Total P2P traffic to 20% of link capacity.

This will leave 80% to other applications and uses.

Upstream Control

Limit upstream (file uploads) P2P traffic, while allowing downstream (file downloads) to continue uninterrupted. This can provide relief to over congested upstream links, while not disrupting file downloads (and thus posing less limitation on subscribers). The SCP ability to isolate and control specifically P2P uploads is essential to this policy as simply limiting upstream data would adversely influence all file transfer, uploads and downloads (by slowing down TCP connections).

EXAMPLE

Limit P2P file upload bandwidth to 24kbps per subscriber.

This will preserve the limited upstream resources while not restricting P2P file downloads (the "important" traffic for P2P users).

Destination Based Classification

Limit traffic that uses expensive or particularly congested links, peering points or transit connections (e.g., an expensive international link). This helps reduce the cost associated with serving P2P traffic on expensive connections.

EXAMPLE

Limit P2P traffic leaving the ISP domain through peering point **xyz** to 10mbps.
Limit each subscriber's OffNET P2P traffic to 64kbps.
No limit to OnNET P2P traffic.

This will reduce peering and transit costs. In addition, since OnNET and the cheaper OffNET traffic is less restricted, P2P applications would automatically shift downloads to use these connections, further reducing the cost associated with P2P traffic.

Time of Day Policies

The Service Control Platform can be configured to provide different limits on P2P usage during different periods of the day and week. This can be used to reduce the congestion caused by P2P during those hours when other mission critical and bandwidth sensitive applications (e.g., email, VPN, etc.) are used. This will also encourage P2P developers and users to automatically shift the usage of the network to different hours of the day and week.



EXAMPLE

Limit P2P traffic during business hours to 5% of link capacity

During business hours, telecommuters and business users receive high levels of quality of service since P2P traffic does not congest the network.

Subscriber Application Quotas

Rather than simply limiting the bandwidth a subscriber receives at any given moment, the Service Control Platform can enforce a byte cap for a certain period of time (e.g., quota per day), after which access can be either completely blocked or throttled to a minimum. This provides a better sense of fairness to subscribers, as they are not fighting over the momentarily available resources but are each given a fair "piece of the pie." Note that the platform can enforce these quotas at the application level, meaning that the byte cap can be performed specifically on P2P traffic. This is superior to performing byte caps at the subscriber level (e.g., a daily limit on all traffic for each subscriber), as it provides better application-based isolation for the subscriber, and gives the subscriber the piece of mind that no matter how many P2P bytes have been consumed, he or she will not lose access to other, critical applications (e.g., VPN, email, etc.)

EXAMPLE

Unlimited P2P traffic up to 1 gigabyte per day, after which P2P is throttled to 64kbps.

Subscriber receives a known quota of P2P traffic. Once exceeded, all other network activities are undisrupted, but P2P is limited.

Subscriber Dynamic Policies

Using the Service Control Platform's subscriber awareness (its ability to correlate each packet to a subscriber), dynamic policies can be implemented to allow subscribers to control their own accounts. As an example, the service provider can design a subscription package that provides unlimited P2P access for an additional charge or develop a "bandwidth-on-demand" system (in which a subscriber can buy additional P2P bandwidth as needed). This gives subscribers the opportunity to manage their own accounts and can potentially open additional revenue streams for the service provider.

EXAMPLE

Unlimited P2P traffic up to 1 gigabyte per day, after which P2P is throttled to 64kbps.

Optional "Unrestricted P2P" subscription package also exists, without quota on that traffic.

Subscribers receive a known quota of P2P traffic. Once exceeded, all other network activities are undisrupted, but P2P is limited. Subscribers can opt to select unrestricted package for unlimited P2P access.



Conclusion

The combination of Peer-to-Peer subscribers' aggressive use of network resources and the growing popularity of these applications is straining broadband networks. Service providers must address these users in a way that limits customer alienation, improves subscriber experience and opens new avenues for revenue growth.

In contrast to costly capacity upgrades and fiber node splits or cumbersome tools, P-Cube's Service Control Platform is an ideal, cost-effective platform for addressing the P2P problem as a business opportunity.

Key capabilities of the platform include:

- ❖ By tracking the state of each network connection, and performing application (layer-7) traffic analysis the platform is able to reliably and accurately classify traffic as P2P, so that a true understanding of the implication of peer to peer traffic can be built and focused control policies implemented
- ❖ Due to its highly programmable architecture, the Service Control Platform provides rapid turnaround of new protocols and applications keeping the solution up to date with the ever-changing world of peer-to-peer applications and protocols.
- ❖ P-Cube's Service Control Platform is the first custom build carrier grade network device, specifically designed to perform real-time traffic classification, usage reporting and traffic control. It provides unsurpassed performance, supporting gigabit line rates, even while executing the most complex policies.
- ❖ Acting as a transparent network overlay, the platform is simple to deploy, requiring minimal network reconfiguration and maintains the investment in existing network equipment and infrastructure.
- ❖ Using the advanced control capabilities of the P-Cube Service Control Platform, effect of P2P traffic can be controlled and limited on the network to increase customer satisfaction and open additional revenue streams.

Visit P-Cube on the web at www.p-cube.com

U.S. Corporate Headquarters:

P-Cube, Inc.
388 Oakmead Parkway
Sunnyvale, CA 94085
Tel: 408-720-7770
Fax: 408-720-7772
info@p-cube.com

UK:

P-Cube Limited
200 Brook Drive
Green Park
Reading RG2 6BU
United Kingdom
Tel: +44 (0) 118 949 7101

Israel:

P-Cube, Ltd.
85 Medinat Hayehudim St.
P.O.B. 12331
Herzliya 46766 Israel
Tel: +972 9-956-9220

