Controlling Peer-to-Peer Applications

Palo Alto Networks
232 E. Java Dr.
Sunnyvale CA 92089
www.paloaltonetworks.com

Table of Contents

Executive Summary

Palo Alto Networks provides enterprises with visibility into and control over applications traversing the network irrespective of port, protocol, SSL encryption or evasive tactic used. With the knowledge of the application identity in hand, administrators can then use that data to implement granular security policies.

As highlighted in this paper, P2P applications are just one example of the type of applications that are identified and can be controlled by Palo Alto Networks. The visibility and control outlined in this paper can be applied to more than 1,000 applications across 25 categories including email, web mail, business applications, networking and more. Visit the Applipedia at for a complete breakdown of the categories and their respective applications.

Introduction

Peer-to-Peer (P2P) technology is not new, in fact, some of the earliest cases of P2P use were in Usenet and news server systems with the use case being the distribution of news articles. P2P applications are designed to leverage shared resources, CPU cycles and bandwidth, across an ad hoc network. The advantage of a P2P network is that it distributes the load across a network and decentralizes command and control, rather than focusing it on a small number of centralized servers. For many years, P2P was used quietly in the technical community and in fact can be a very useful tool for an IT department or anyone who needs to deal with moving large files around.

The Internet boom and the release of Napster brought P2P squarely into the music sharing application business and more recently it has expanded into video sharing and distribution. While Napster has been shut down, many more P2P offspring have risen, garnering an extremely bad reputation brought on not by it's distribution efficiency but by what was being distributed - copyrighted music files and other materials.

The Dark Side of P2P

In addition to being a common source of pirated music and movies, P2P applications have been at the heart of some high profile examples of inadvertent sharing of proprietary or confidential data. Examples include employee records, blueprints for President Obama's helicopter, the location of a former first lady's safe house and health care records. The FTC recently took steps to tell more than 100 organizations that proprietary data including health-related information, financial records and drivers' license and Social Security numbers was found on P2P networks.

In every case, the data sharing was a result of the presence of P2P applications on the user's computers. Whether these data exposure examples were purposeful or not is unknown – what is known is that many of the readily available P2P applications are confusing and can be improperly configured, resulting in the unintentional sharing of folders as highlighted in a report assembled by the US Patent and Trademark Office.

Not to be left out in the cold, malware creators and thieves have joined the fray. Knowing that P2P can easily evade today's firewalls and that it is pervasively deployed by inexperienced users, hackers have begun using P2P as a means of BOT distribution and file collections. The Mariposa botnet is the most recent example. Mariposa spreads itself across nine different P2P networks including: Ares, Bearshare, Direct Connect, eMule, iMesh, Kazaa, Gnutella, BitTorrent, (via LimeWire client), and Shareaza. Essentially, for each P2P network, there is a Mariposa foldershare feeding the bot executable. In addition to P2P applications, MSN Instant Messaging is also used as a spreader.

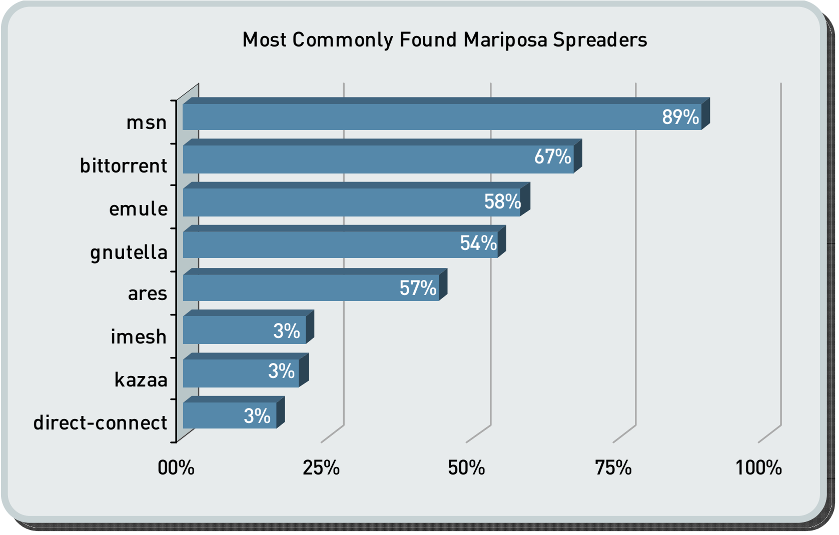

Figure 1: Most commonly detected mariposa spreaders detected.

A more detailed analysis of the data shows the following statistics:

- 312 (86%) of the organizations had at least one of the P2P applications used by Mariposa.
- An average of three of the nine P2P applications were found in each organization.
- Total bandwidth consumed by the nine P2P applications that spread Mariposa was 17.3 terabytes or an average of 55 gigabytes per organization.
- Session consumption by nine P2P applications was 555 million or 1.8 million sessions per organization average.

P2P Applications For Commercial Use?

What is important to point out is that P2P was not developed to be used maliciously – it can be a valuable tool in the hands of an experienced user. And as the trend towards web-based media continues, P2P technology can be used to deliver a service or improve productivity. An analysis of more than 1,000 applications within Applipedia shows 102 applications are based on peer-to-peer technology. As a reminder, only 40 of them are file sharing applications.

Category	Subcategory	Count	Examples
Collaboration (32)	voip-video	23	Camfrog, AIM Audio/video
	instant-messaging	9	IP Messanger, Gtalk File Transfer
General Internet (41)	file-sharing	40	BitTorrent, Xunlei, Winny
	internet-utility	1	Trinoo
Business Systems (2)	storage-backup	1	NDMP
	office-programs	1	MS Groove
Media (21)	audio-streaming	2	Simplify, Octoshape
	photo-video	19	PPLive, TVants, PPStream
Networking (6)	encrypted-tunnel	3	Hamachi
	infrastructure	3	UPNP, JXTA
Total		102	

There are many other applications where P2P is used in a beneficial manner, in fact, one of the most common mechanisms to download updated Linux binaries is via BitTorrent. Given the fact that the Internet is an established piece of most corporate networks, P2P technology will continue to be used on corporate networks and rolled into commercial software solutions. As a result, corporate IT departments will be doubly challenged in their efforts to control P2P file sharing applications--blocking the “bad” P2P applications (and their owners) while allowing the “good” ones. Best practices need to be implemented—both in terms of policies and technology—to control the use of P2P and its associated risks.

Controlling P2P Applications

As a tool that leverages distributed resources, P2P applications need to be connected to a network of other users. This enables them to evade detection through various means. They will port hop, continually searching for a port that the firewall allows traffic over to get connected. They are known to emulate or use HTTP as another means of passing through the firewall as if it is web browsing traffic, undetected. And some P2P applications are known to use encryption, yet another means of evading detection. P2P applications clearly provide real business benefits to enterprises but they also bring certain risks. In order to mitigate those risks, it is recommended that a multi-faceted approach to controlling them be implemented.

If application controls are going to be implemented and enforced, they should be part of the overarching corporate security policy, whether it is P2P or otherwise. And as part of the process of implementing an application control policy, IT should make a concerted effort to learn about the P2P applications. Proactively installing them in a lab environment to see how they act can be very educational. Another form of education is peer discussions, asking other IT professionals for their input is an invaluable source of information. Yet a third source of information would be the P2P focused web sites, message boards and blogs. Ignoring the issue or acknowledging it exists but doing nothing about it until “something happens” can be a career limiting tact.

Employee Controls

Most companies have some type of application usage policy, outlining which applications are allowed, and which are prohibited. It is assumed that every employee understands the contents of this policy and the ramifications of not complying with it. While employee policies are a key piece to the P2P control puzzle but in many cases, they represent a number of challenges.

- Given the increasing number of “bad” applications, how will an employee know which applications are allowed and which are banned?
- How is the list of unapproved applications updated, and who ensures employees know the list has changed?
- What constitutes a policy violation?
- What are the ramifications of policy violations – firing or a reprimand?

Documented employee policies are a key piece to the P2P control puzzle but will remain largely ineffective as a stand alone control mechanism.

Desktop Controls

Desktop controls present IT departments with significant challenges. Careful consideration should be applied to the granularity of the desktop controls and the impact on employee productivity. As with employee policies, desktop controls are a key piece to the gaining the upper hand on the growth of P2P applications in the enterprise and if used alone, will be ineffective for several reasons.

The drastic step of desktop lock down to keep users from installing their own applications is a task that is easier said than done.

- Laptops connecting remotely, Internet downloads, USB drives and email are all means of installing applications that may or may not be approved.
- Removing administrative rights completely has proven to be difficult to implement and, in some cases, limits end-user capabilities.
- USB drives are now capable of running an application so in effect, the P2P application could be accessed after the network admission was granted.

Desktop controls can complement the documented employee policies as a means to gain the upper hand on P2P applications.

Network Controls

At the network level, what is needed is a means to identify the P2P applications and block or control them. By implementing network level controls, IT is able to minimize the possibility that internal files are shared and network congestion associated with large file transfers is alleviated. Several possible control mechanisms can be used at the network level, each of which carry certain drawbacks that reduce their effectiveness.

- Stateful firewalls can be used as a first line of defense, providing coarse filtering of traffic and segmenting the network into different, password protected zones. One drawback to Stateful firewalls is that they use protocol and port to identify and control what gets in and out of the network. This port-centric design is relatively ineffective when faced with applications such as P2P that hop from port to port until they find an open connection to the network.
- Adding IPS to a firewall deployment enhances the network threat prevention capability by looking at a subset of traffic and blocking known threats or bad applications. IPS offerings lack the breadth of applications and the performance required to look at all traffic across all ports and as such, cannot be considered a full solution.
- IPS technologies are typically designed to look only at a partial set of traffic to avoid impeding performance, and as such, would be unable to cover the breadth of applications needed by today’s enterprises. And finally, managing a FW+IPS combination is usually a cumbersome task, requiring different management interfaces pointed at separate policy tables. Simply put, the current bolt-on solutions do not have the accuracy, policy, or performance to solve today’s application visibility and control requirements.
- Proxy solutions are another means of traffic control but here too, they look at a limited set of applications or protocols and as such only see a partial set of the traffic that needs to be monitored. So a P2P application will merely see a port blocked by a proxy and hop over to the next one that is open. By design, proxies need to mimic the application they are trying to control so they struggle with updates to existing applications as well as development of proxies for new applications. A final issue that plagues proxy solutions is throughput performance brought on by how the proxy terminates the application, and then forwards it on to its destination.

The challenge with any of these network controls is that they do not have the ability to identify P2P applications; look only at a portion of the traffic and suffer from performance issues.

Using the Palo Alto Networks Firewalls to Control P2P

Palo Alto Networks avoids many of the issues that existing network control solutions suffer from by delivering a high performance firewall that takes an application-centric approach to traffic classification, accurately identifying all applications traversing the network, irrespective of port, protocol, SSL encryption or evasive tactic employed. By addressing security evasion tactics commonly used in many of today's new applications with a new traffic classification technology called App-ID^{TM}, Palo Alto Networks can help IT regain control over P2P applications at the network level, complementing any existing desktop and employee control mechanisms.

App-ID^{TM} accurately identifies more than 1,000 applications, including more than 40 P2P applications, flowing in and out of the network. All traffic flowing through the Palo Alto Networks firewalls is processed by App-ID^{TM} and the identity of the application, P2P or otherwise, is displayed in the management interface using their common application names. The application identity is then used as the basis of all firewall security policies including access control, user permissions, threat prevention and more.

Identifying P2P Applications

Palo Alto Networks' App-ID uses multiple traffic classification mechanisms, operating in concert, to determine exactly what applications are traversing the network. App-ID looks at traffic flowing across all ports, not just a subset of known used ports, thereby increasing the odds of detecting the application. By taking an application-centric approach to traffic classification, Palo Alto Networks addresses security evasion tactics used by P2P applications, such as the use of non-standard ports, dynamically changing ports and protocols, emulating other applications, and tunneling to bypass existing firewalls. As a result, App-ID identifies more than 40 different P2P applications, which translates hundreds of P2P clients. By identifying the P2P network, as opposed to the clients, broader coverage is achieved in the effort to control P2P usage.

As traffic flows across the network, App-ID establishes the application session and maintains session state, while application identification mechanisms are methodically applied to accurately classify the application. App-ID uses as many as four traffic classification mechanisms to identify traffic. The specific mechanisms that are used to identify P2P are Application Signatures, Decoders and Heuristics.

- **Application Signatures:** Context-based signatures are used first to look for unique application properties and related transaction characteristics to correctly identify the application regardless of the protocol and port being used. The signature also determines if the application is being used on its default port or it is using a non-standard port (for example, RDP across port 80 instead of port 3389, its standard port).
- **Application Protocol Decoding:** If needed, protocol decoders are then employed to determine whether the application is using a protocol as its normal transport (such as HTTP for web browsing applications), or if it is only using the protocol as an obfuscation technique to hide the real application (for example, Yahoo! Instant Messenger used across HTTP). Protocol decoders also help narrow the range of possible applications, providing valuable context when applying signatures and they identify files and other content that should be scanned for threats or sensitive data.

- **Heuristics:** In certain cases, evasive applications still cannot be detected even through advanced signature and protocol analysis. In those situations, it is necessary to apply additional heuristic, or behavioral analysis to identify certain applications such as peer-to-peer or VoIP applications that use proprietary encryption. Heuristic analysis is used as needed, with the other App-ID techniques discussed here, to provide visibility into applications that might otherwise elude positive identification.

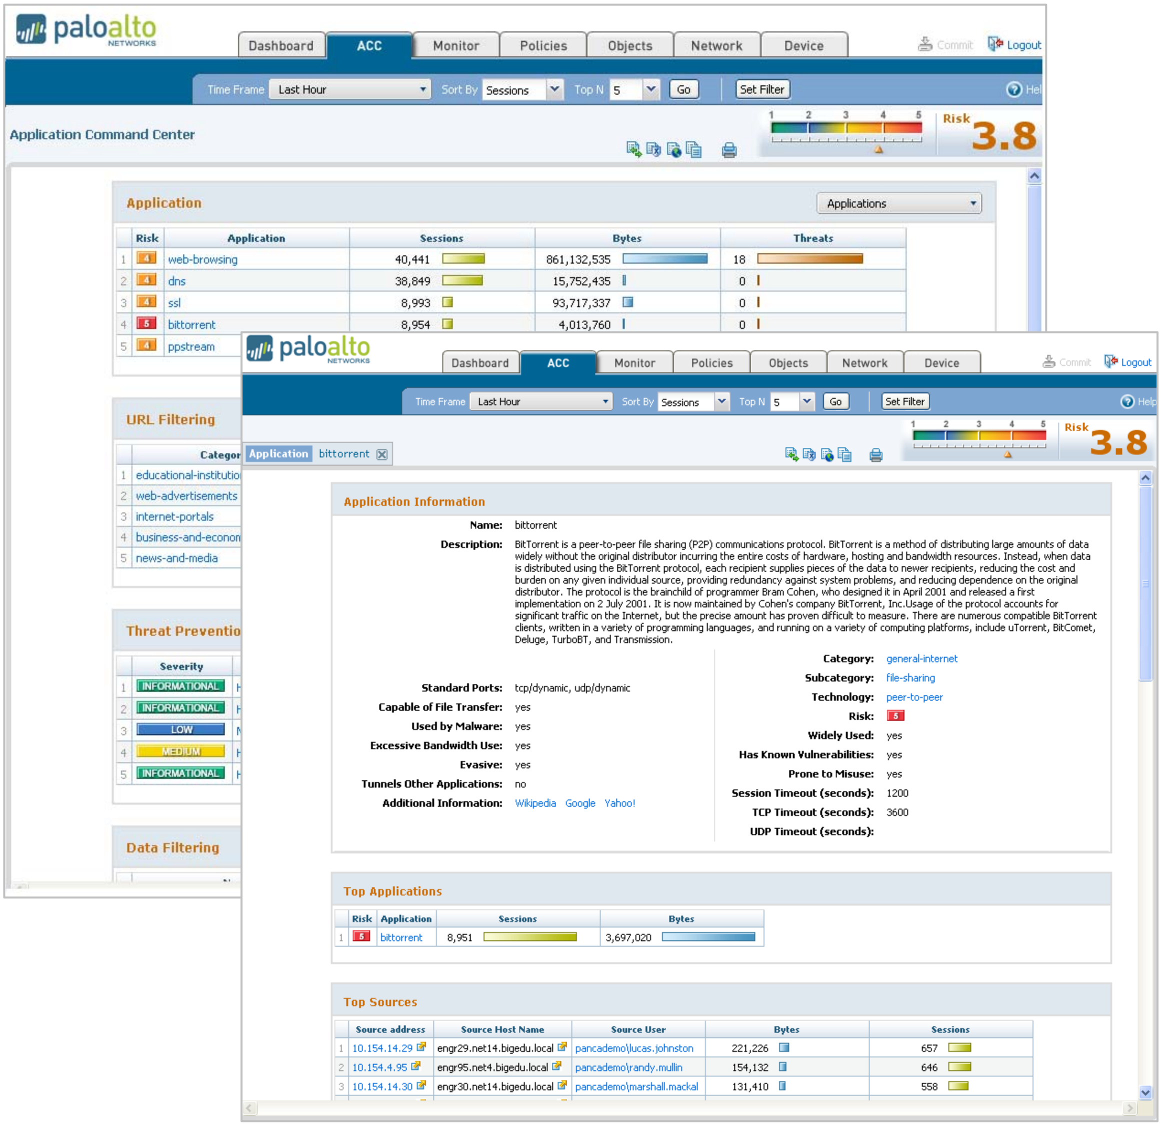

Time Frame: Last Hour · Sort By: Sessions · Top N: 5

Application Command Center

Application (Applications)

	Risk	Application	Sessions	Bytes	Threats
1	4	web-browsing	40,441	861,132,535	18
2	4	dns	38,849	15,752,435	0
3	4	ssl	8,993	93,717,337	0
4	5	bittorrent	8,954	4,013,760	0
5	4	ppstream			

URL Filtering

	Categor
1	educational-institutio
2	web-advertisements
3	internet-portals
4	business-and-econon
5	news-and-media

Threat Preventio

	Severity
1	INFORMATIONAL
2	INFORMATIONAL
3	LOW
4	MEDIUM
5	INFORMATIONAL

Data Filtering

Application: bittorrent

Application Information

Name: bittorrent

Description: BitTorrent is a peer-to-peer file sharing (P2P) communications protocol. BitTorrent is a method of distributing large amounts of data widely without the original distributor incurring the entire costs of hardware, hosting and bandwidth resources. Instead, when data is distributed using the BitTorrent protocol, each recipient supplies pieces of the data to newer recipients, reducing the cost and burden on any given individual source, providing redundancy against system problems, and reducing dependence on the original distributor. The protocol is the brainchild of programmer Bram Cohen, who designed it in April 2001 and released a first implementation on 2 July 2001. It is now maintained by Cohen's company BitTorrent, Inc.Usage of the protocol accounts for significant traffic on the Internet, but the precise amount has proven difficult to measure. There are numerous compatible BitTorrent clients, written in a variety of programming languages, and running on a variety of computing platforms, include uTorrent, BitComet, Deluge, TurboBT, and Transmission.

Standard Ports: tcp/dynamic, udp/dynamic
Capable of File Transfer: yes
Used by Malware: yes
Excessive Bandwidth Use: yes
Evasive: yes
Tunnels Other Applications: no
Additional Information: Wikipedia Google Yahoo!

Category: general-internet
Subcategory: file-sharing
Technology: peer-to-peer
Risk: 5
Widely Used: yes
Has Known Vulnerabilities: yes
Prone to Misuse: yes
Session Timeout (seconds): 1200
TCP Timeout (seconds): 3600
UDP Timeout (seconds):

Top Applications

	Risk	Application	Sessions	Bytes
1	5	bittorrent	8,951	3,697,020

Top Sources

	Source address	Source Host Name	Source User	Bytes	Sessions
1	10.154.14.29	engr29.net14.bigedu.local	pancademo\lucas.johnston	221,226	657
2	10.154.4.95	engr95.net4.bigedu.local	pancademo\randy.mullin	154,132	646
3	10.154.14.30	engr30.net14.bigedu.local	pancademo\marshall.mackal	131,410	558

Figure 1: *Application Command Center (background screen) provides a current view of application activity with drill down into specific applications such as BitTorrent for additional details (foreground screen).*

When App-ID identifies a P2P application, the administrator can quickly see exactly which of the P2P applications are running on the network. A few clicks of a mouse tells the administrator who's using the application, the bytes and sessions being consumed, the traffic source/destination and more.

Additional information to assist in making an informed decision on how to treat the application includes a definition of the application, its behavioral characteristics, its potential risks, and 3rd party sources of information. Within minutes, administrators can make a more informed decision on how to treat the application: deny, allow, scan, apply QoS, inspect for files/datapatterns.

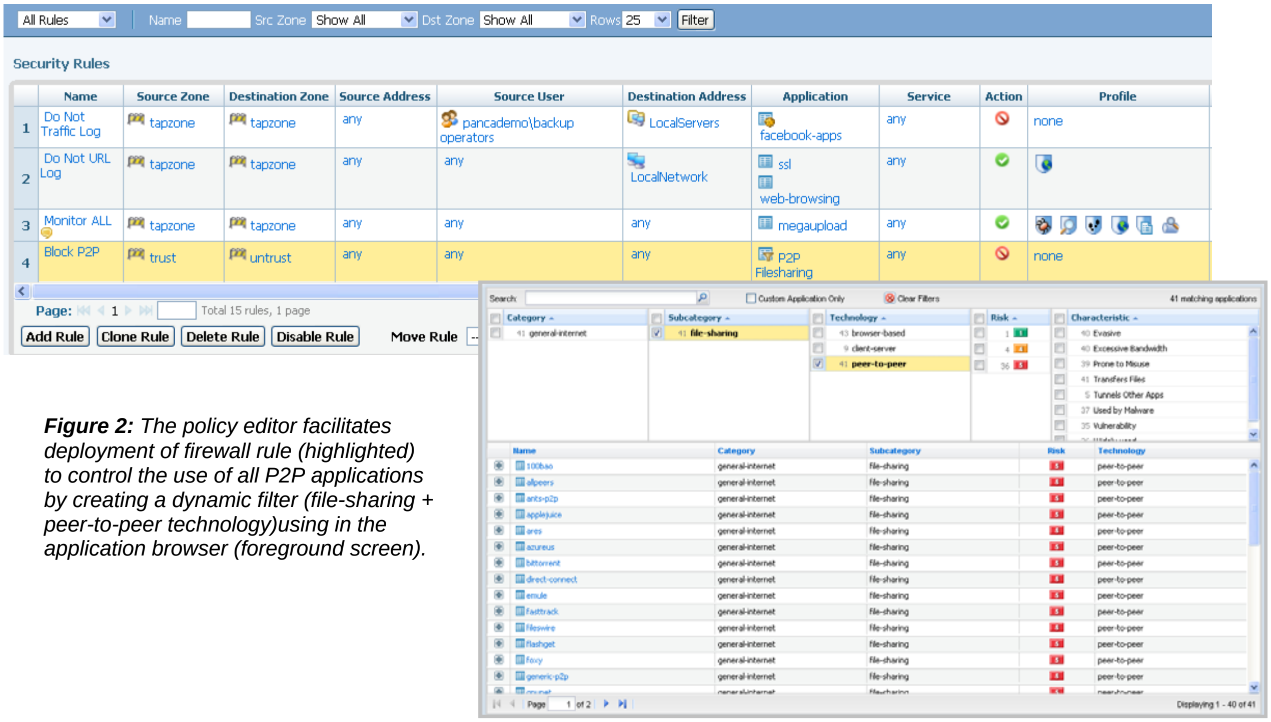

All Rules | Name | Src Zone Show All | Dst Zone Show All | Rows 25 | Filter

Security Rules

	Name	Source Zone	Destination Zone	Source Address	Source User	Destination Address	Application	Service	Action	Profile
1	Do Not Traffic Log	tapzone	tapzone	any	pancademo\backup operators	LocalServers	facebook-apps	any	⊘	none
2	Do Not URL Log	tapzone	tapzone	any	any	LocalNetwork	ssl web-browsing	any	✓	
3	Monitor ALL	tapzone	tapzone	any	any	any	megaupload	any	✓	
4	Block P2P	trust	untrust	any	any	any	P2P Filesharing	any	⊘	none

Page: 1 | Total 15 rules, 1 page

Search: | □ Custom Application Only | Clear Filters | 41 matching applications

Category	Subcategory	Technology	Risk	Characteristic
□ 41 general-internet	☑ 41 file-sharing	□ 43 browser-based	□ 1 🟩 1	□ 40 Evasive
		□ 9 client-server	□ 4 🟧 4	□ 40 Excessive Bandwidth
		☑ 41 peer-to-peer	□ 36 🟥 5	□ 39 Prone to Misuse
				□ 41 Transfers Files
				□ 5 Tunnels Other Apps
				□ 37 Used by Malware
				□ 35 Vulnerability

Name	Category	Subcategory	Risk	Technology
100bao	general-internet	file-sharing	5	peer-to-peer
allpeers	general-internet	file-sharing	4	peer-to-peer
ants-p2p	general-internet	file-sharing	5	peer-to-peer
applejuice	general-internet	file-sharing	5	peer-to-peer
ares	general-internet	file-sharing	4	peer-to-peer
azureus	general-internet	file-sharing	5	peer-to-peer
bittorrent	general-internet	file-sharing	5	peer-to-peer
direct-connect	general-internet	file-sharing	4	peer-to-peer
emule	general-internet	file-sharing	5	peer-to-peer
fasttrack	general-internet	file-sharing	5	peer-to-peer
fileswire	general-internet	file-sharing	4	peer-to-peer
flashget	general-internet	file-sharing	5	peer-to-peer
foxy	general-internet	file-sharing	5	peer-to-peer
generic-p2p	general-internet	file-sharing	4	peer-to-peer

Page 1 of 2 | Displaying 1 - 40 of 41

Figure 2: *The policy editor facilitates deployment of firewall rule (highlighted) to control the use of all P2P applications by creating a dynamic filter (file-sharing + peer-to-peer technology)using in the application browser (foreground screen).*

Applying Positive Control (Firewall) Policies to P2P Applications

Once the P2P application(s) have been identified, an administrator can use the rule-based editor to create a P2P application usage control policy for BitTorrent. Or, to be safer and implement a wider net, the policy can be established against the all 40 peer to peer applications using a filter (shown above). The advantages of using a filter in the policy editor is that it covers all current P2P applications as well as those that are added in the future. As new P2P applications are added by the Palo Alto Networks development team, they are automatically covered through the Palo Alto Networks dynamic update service.

In some cases, administrators will want to block all P2P applications, while in others, such as when a company uses BitTorrent to distribute software, they may wish to enforce specific rules to allow it - but only for key individuals. In this scenario, an application such as BitTorrent can be selected, and then the specific users allowed to use BitTorrent are selected based on their Active Directory information – in the example below, the Administrator group are allowed to use BitTorrent.

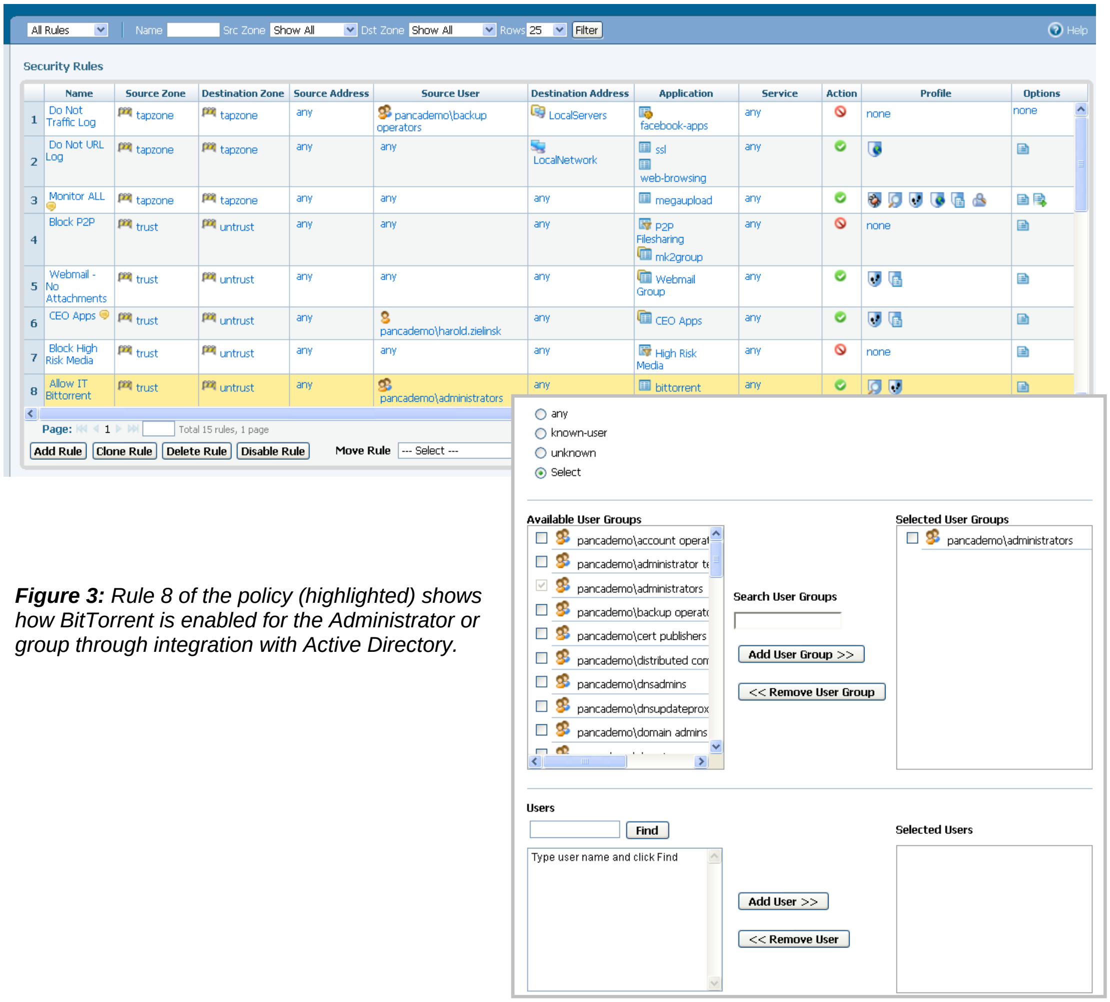

Security Rules

	Name	Source Zone	Destination Zone	Source Address	Source User	Destination Address	Application	Service	Action	Profile	Options
1	Do Not Traffic Log	tapzone	tapzone	any	pancademo\backup operators	LocalServers	facebook-apps	any	⊘	none	none
2	Do Not URL Log	tapzone	tapzone	any	any	LocalNetwork	ssl web-browsing	any	✓		
3	Monitor ALL	tapzone	tapzone	any	any	any	megaupload	any	✓		
4	Block P2P	trust	untrust	any	any	any	P2P Filesharing mk2group	any	⊘	none	
5	Webmail - No Attachments	trust	untrust	any	any	any	Webmail Group	any	✓		
6	CEO Apps	trust	untrust	any	pancademo\harold.zielinsk	any	CEO Apps	any	✓		
7	Block High Risk Media	trust	untrust	any	any	any	High Risk Media	any	⊘	none	
8	Allow IT Bittorrent	trust	untrust	any	pancademo\administrators	any	bittorrent	any	✓		

Page: 1 Total 15 rules, 1 page

○ any
○ known-user
○ unknown
● Select

Available User Groups

- □ pancademo\account operat
- □ pancademo\administrator te
- ☑ pancademo\administrators
- □ pancademo\backup operato
- □ pancademo\cert publishers
- □ pancademo\distributed com
- □ pancademo\dnsadmins
- □ pancademo\dnsupdateprox
- □ pancademo\domain admins

Search User Groups

Selected User Groups

- □ pancademo\administrators

Users

Type user name and click Find

Selected Users

Figure 3: *Rule 8 of the policy (highlighted) shows how BitTorrent is enabled for the Administrator or group through integration with Active Directory.*

Conclusion

By discarding the traditional traffic classification mechanisms of port and protocol, and taking an application centric approach, the Palo Alto Networks next-generation firewall is able to bring unparalleled application visibility and control back to the IT department. Whether the need is to control one of the application categories such as P2P or a more general application visibility and control requirement, the Palo Alto Networks firewall allows administrators to define traditional firewall policies to control their application traffic.