The Application Usage and Threat Report

An Analysis of Application Usage and Related Threats within the Enterprise

10th Edition, February 2013

Palo Alto Networks
3300 Olcott Street
Santa Clara, CA 94089
www.paloaltonetworks.com

Table of Contents

Executive Summary

Since 2008, Palo Alto Networks has published trends and analysis in application usage across enterprise networks in its bi-annual *Application Usage and Risk Report*. This version of the report marks an evolution of sorts – it now includes threat activity, specifically malware and exploits, across the applications observed and therefore, the name of the report has been changed.

The *Application Usage and Threat Report (1st Edition, January 2013)* from Palo Alto Networks provides a global view into enterprise application usage and the associated threats by summarizing network traffic assessments conducted in 3,056 organizations worldwide between May 2012 and December 2012.

This report edition will be the first report of its kind to discuss application usage patterns and the specific type of threat they may or may not introduce. The application and threat patterns discussed within this report dispel the position that social networking, filesharing and video applications are the most common threat vectors, while reaffirming that internal applications are highly prized targets. Rather than use more obvious, commercially available applications, attackers are masking their activities through custom or encrypted applications.

Key findings include:

Applications commonly viewed as top threat sources are, in fact, not.

- 339 social networking, video, and filesharing applications represent 20% of the bandwidth but displayed only 0.4% of the threat logs.
- Exploits, not malware logs, were more commonly detected in social networking by a ratio of 49:1.
- Exploits observed in Facebook applications (3rd party applications and widgets) were 228 times greater in number than in other social networking applications.

Exploits continue to target enterprises' most valued assets.

- Out of 1,395 applications found, 10 were responsible for 97% of all exploit logs observed.
- Of the 10 applications, 9 are internal applications and they represented 82% of the exploit logs.

Malware relies heavily on custom applications.

- Custom or unknown traffic was the #1 type of traffic associated with malware communications, as leading malware families continue to customize their command-and-control traffic.
- Control of unknown and custom traffic provided an intriguing option for controlling botnet communications.

The use of SSL – both a security mechanism and a masking agent.

- 356 applications used SSL in some way, shape or form - 85 of them did not use standard SSL ports.
- SSL by itself represented 5% of all bandwidth and the sixth highest volume of malware logs within known applications.
- HTTP proxy, used both as a security component and to evade controls, exhibited the seventh highest volume of malware logs.

The analysis and related findings in this report are generated via live network traffic observed in several thousand organizations worldwide. In that respect the report is unique in that it is not based on a survey – it is real data collected from live traffic.

Data Sources and Key Facts

A summary of the data sources, statistics observed, and the key facts are listed below. Additional commentary and analysis is included throughout the report.

- A total of 1,395 applications consumed more than 12.6 petabytes (12,640,385,037,520,200 bytes) of bandwidth across 3,056 participating organizations.
- Bandwidth consumption was roughly equivalent to 4.2 million 2-hour HD movie downloads (average download size of 3GB).
- Over 5,300 unique critical, high, and medium severity threats representing more than 268 million logs were observed.

	Threat Logs Viewed – by Severity			
Threat Type	**Critical**	**High**	**Medium**	**Total**
Malware: botnet	98,546,921	30,206,844	3,334	128,757,099
Malware: spyware	3,053,523	194,983	51,545,824	54,794,330
Malware: adware	13,475,720	628,367	40,576	14,144,663
Malware: backdoor	26,744	4,780,764	218,603	5,026,111
Malware: net-worm	1,766,940			1,766,940
Malware: keylogger		1,936		1,936
Malware: total logs	**116,869,848**	**35,812,894**	**51,808,337**	**204,491,079**
Exploit: code-execution	9,403,354	18,851,800	8,681,830	36,936,984
Exploit: overflow	1,560,304	7,880,736	15,874,245	25,315,285
Exploit: sql-injection		5,589	1,408,599	1,414,188
Exploit: total logs	**10,963,658**	**26,738,125**	**25,964,674**	**63,666,457**
Grand total	**127,833,506**	**62,551,019**	**77,773,011**	**268,157,536**

- Collectively, social networking, filesharing, and photo-video applications represented 25% of the applications (339) and 20% of total bandwidth (890,000+ 2hr high-definition movie downloads), but only 0.4% of all threat logs observed.
- The number of application variants found in each category were: social networking (75), filesharing (152), and photo-video (112).
- Each network analyzed had an average of 17 social networking, 19 filesharing, and 30 photo-video application variants.
- Of the 75 social networking applications found, the four Facebook functions (-base, -apps, -social-plugins, and -posting) represented 75% of all social networking bandwidth.
- Facebook applications (3rd party applications and widgets) represented 97% of **all** social networking threat logs and 99% of **all** social networking exploit logs yet only 0.2% of the respective bandwidth.
- Exploit logs observed in Facebook-apps were 228 times greater than the application with the second highest volume of exploits (Facebook-base).
- Myspace-posting was found in only 3% of the 3,056 organizations, yet it has the highest byte-per-session consumption within social networking (1.8MB per session).
- Google-plus-posting was nearly non-existent in enterprise environments – found in only 5 of the 3,056 organizations. Comparatively, posting activity for LinkedIn and Facebook were found in 1,471 and 2,550 organizations respectively.

- The 152 filesharing applications found consumed a 6.2% of the total bandwidth observed with BitTorrent representing roughly half that amount (3%).
- The top 10 filesharing applications represented 92% of the respective bandwidth; 98% of the respective threat logs observed; yet they are distributed across the all three technologies (3 P2P, 3 client/server, and 4 browser-based).
- FTP and WebDAV displayed the highest number of filesharing threat logs (primarily exploits) and were the fourth and sixth most heavily used filesharing applications.
- 97% of all exploit logs were found in ten applications; nine of those applications were internal/infrastructure applications (databases, active directory, RPC, etc.).
- 99.99% of all malware logs were found in only seven (out of 1,395) applications with *custom/unknown-UDP* representing the highest volume at 55%.
- In contrast, exploits were a small percentage of custom or unknown traffic: *custom/unknown-TCP* displayed a mere 0.3% of exploit logs while *custom/unknown-UDP* displayed even fewer logs.
- Exploits in custom traffic were high risk: 83% were classified as "critical" and the remaining 17% as "high" or "medium" severity.
- 26% of the applications (356) used SSL in some way shape or form; these applications represented roughly 7% of total bandwidth.
- SSL by itself represented 5% of total bandwidth, 7% of all sessions and the sixth highest number of malware logs – primarily command and control traffic.
- 85 of the 356 applications that used SSL never used port 443, nor did they use SSL defined ports (37 hop ports, 28 use TCP/80, 20 use ports other than TCP/443).

When Does 25% + 20% = 0.4%?

The concept of "share everything" that has permeated the user population has instilled an unreasonable level of trust, which in turn encourages the assumption that social networking, filesharing and video applications are the primary source of threats. The analysis found that these 339 applications do in fact act as threat vectors, however when compared to their popularity and more importantly to the other applications found, the volume of threats observed was significantly less than expected.

Collectively, the 339 social networking, filesharing and video applications found represent 25% of the applications, 20% of the total bandwidth observed but only 0.4% of **all** threat logs observed. Put another way, merely blocking all of these applications will indeed improve the security posture of any organization, but not in the massive leaps and bounds that one would hope. The other aspect of "*block them all*" to consider is the user backlash and potentially isolating the customers by offering them fewer communications vehicles.

Facebook Domination Continues

The Facebook domination within social networking continues unabated, despite the emergence of new offerings and the re-emergence of older offerings.

- 75 different social network applications were found that consumed 1% of the total bandwidth; on average 17 variants were found on 92% of the 3,056 networks observed.
- Out of the 5,303 individual threats observed, 117 were found within social networking applications.
- The total threat logs within social networking was surprisingly low - a mere 0.18% of all threat logs viewed. Within the social networking threats logs observed, the ratio of exploit to malware logs was 49:1.

Of the 75 social networking applications, four of the Facebook functions are identified separately, (Facebook-base, -apps, -posting, and social-plugins) allowing organizations to enable some, while disabling others. These four Facebook functions consumed 75% of **all** social networking bandwidth.

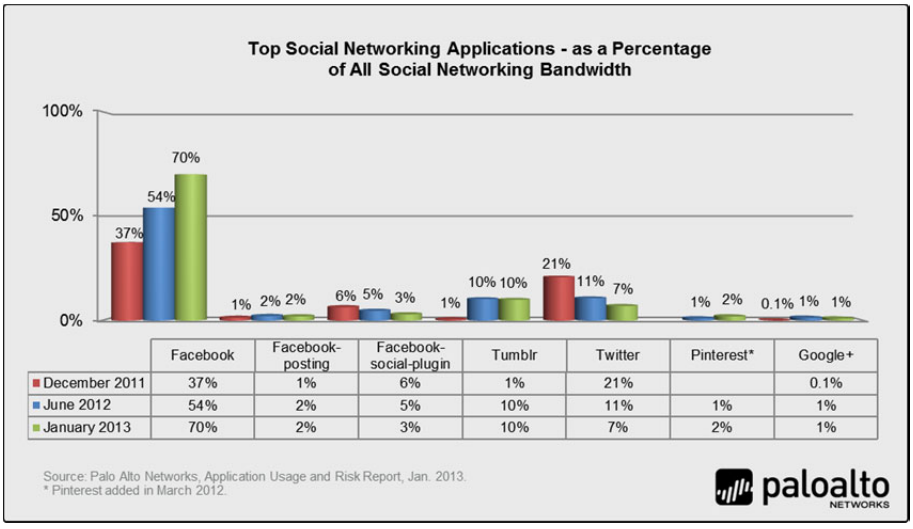

Figure 1: Top social networking application bandwidth consumption comparison.

Put another way, the other 71 social networking applications are left to divide the remaining 25% bandwidth. While Facebook is the most frequently and heavily used social media application, new participants are slowly establishing their own niche; others are showing signs of re-emergence.

- **Tumblr:** Tumblr maintained its position behind Facebook – it was detected in 85% of the organizations and it consumed 10% of the social networking bandwidth. From a business use case perspective, Tumblr currently still has fairly limited appeal. The data in this report was collected before the Tumblr worm outbreak in December 2012, and thus that event is not reflected in this report., It does however highlight the balance organizations will need to take when deciding whether or not to allow Tumblr usage.

- **Pinterest:** Pinterest showed steady growth as more organizations add it to their social media applications (and as more users sign up). 85% of the participating organizations had Pinterest on their networks, up 5-fold from 15% in Spring 2012. Social networking bandwidth consumed by Pinterest doubled to 2%.

- **MySpace and Google+:** Lost in the Facebook tidal wave is the re-emergence of MySpace and Google+. MySpace-browsing continues to be found in roughly 65% of the participating organizations when viewed in year-over-year comparisons. Google+ exhibited very little growth in terms of usage by enterprises, or by users while at work. The two data points that support this are that the frequency of use (85% of the organizations had at least one user) and the bandwidth consumption remains unchanged at 1% if all social networking bandwidth. Most of the usage would appear to be users checking updates since Google+ posting was found in only five organizations (out of 3,056), compared to 2,550 instances of Facebook-posting and 1,471 instances of LinkedIn-posting.

Facebook-apps Accounts for 97% of all Social Networking Threat Logs

The Facebook-apps App-ID collectively identifies the many thousands of custom Facebook applications and widgets, many of which are designed for personal use. A quick scan of the list shows an emphasis on computer games, music, video, and travel applications – most of which will have limited business value to enterprises. Of the 75 social networking applications, Facebook-apps was detected in 74% of the participating organizations and displayed 97% of **all** social networking threat logs, the majority of which being a high severity *HTTP cross-site scripting* exploit. Separating exploits and malware within the social networking group, Facebook-apps represents 99% of the social networking exploit logs viewed.

The heavy presence of cross-site scripting attacks within Facebook-apps was of particular note. These attacks allow an attacker to deliver malicious input to an insecure web application, which in turn can then reflect malicious content to an unsuspecting user, who also uses the web application. These attacks are well known and web application developers spend significant effort to ensure their applications are not vulnerable. However, Facebook has a vast number of applications that are often developed by enthusiasts who may not appreciate the security consequences of their application. Even though such problems were not found to be common in Facebook apps, a small number of applications were responsible for a very large number of cross-site scripting attacks.

Filesharing Usage Diversifies

During the analysis, there were 152 different file sharing applications detected, with an average of 19 variants on 98% of the 3,056 networks analyzed. Total bandwidth consumed by all file sharing applications was 6%, with BitTorrent exhibiting the heaviest usage (3% of total). In terms of two hour-high-definition movies, the 152 file sharing applications consumed roughly the equivalent of 264,000 movie downloads.

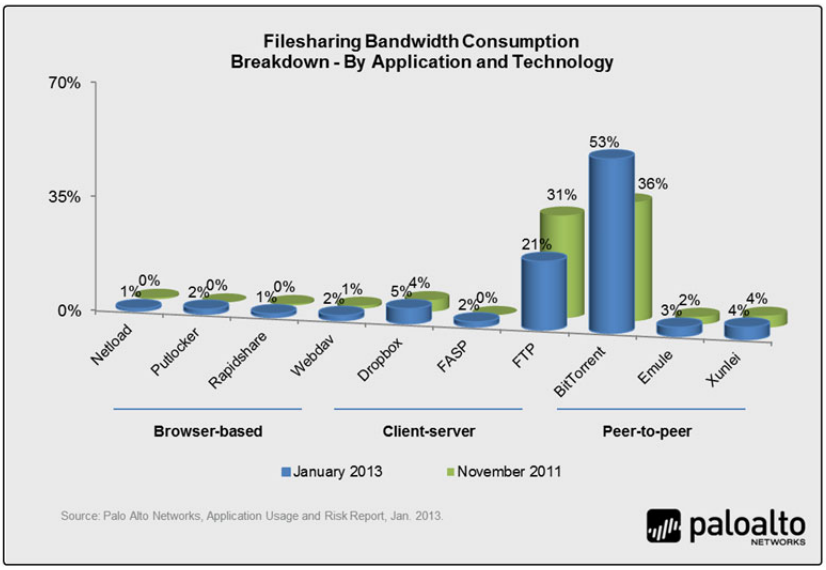

Figure 2: Top 10 filesharing applications based on total bandwidth consumption.

In the past, the filesharing usage patterns were clearly delineated between those that use high amounts of bandwidth (peer-to-peer, client-server) and those that were used with high frequency (browser-based). Looking at the bandwidth consumed by the top-10 file sharing applications, the usage patterns seen in this dataset are spread more evenly across the three technologies. The largest contributor to the balancing of the three technologies is the migration of Dropbox from a browser-based application to a client-server application. The very early versions of Dropbox did not require a user to install a client upon registration – it does so now. Once the account is established, a user can then access their folders via a browser. This, combined with the steady growth of Dropbox, balances the technology breakdown of the top-10 file sharing applications.

Browser-based Filesharing: A Very Crowded Market

Since 2008, this segment of applications has expanded incessantly to where there are 97 different variants identified by Palo Alto Networks and other new offerings being announced regularly. The density of offerings means that some manner of segmentation will continue. Currently, there are two coarsely defined segments – those browser-based file sharing applications that are used for productivity, and those that are used for entertainment.

- **Facebook-file-sharing:** This application was found in only 190 organizations (6%) and compared to the other Facebook properties observed, the usage is low. The low usage may be due to the fact that it is not a standalone feature – it is a feature within Facebook-groups, making it less appealing for sharing files with those users outside of their respective Facebook-group.

- **Skydrive:** This application continues to grow in terms of usage as Microsoft enhances it and the surrounding cloud-based productivity tools. It was found in 85% of the participating organizations, up from 65% in previous reports. Bandwidth consumption within the browser-based segment showed Skydrive consumed 6% - up from less than 1% in previous reports. The integration with Microsoft tools would indicate that Skydrive is more of a productivity offering than one used for entertainment.

- **Putlocker, Megaupload and Mega:** Putlocker continues to grow in terms of frequency and bandwidth consumption as it acts as an entertainment-oriented replacement for the now defunct Megaupload. Mega, the new offering recently announced by the founders of Megaupload will encrypt every file that is stored in its service, making it difficult for outside sources to view and access them.

- **EMC Syncplicity:** EMC recently announced Syncplicity, which is a private, cloud-based synch-n-share application integrates with Documentum and other EMC storage components. Syncplicity will give organizations access to a private cloud with corporate visibility and security enabling more granular control over their internal documents; the files don't have to go off-site to a third-party synch-n-share.

Threats Target Internal Filesharing Applications

File sharing applications displayed the 3rd highest volume of unique threats with 325 discovered across 38 of the 152 application variants. The majority of the unique threats found were exploits.

- **Browser-based filesharing:** There were 153 threats detected within the twenty browser-based applications (out of 73 variants total) with WebDAV being the most heavily targeted with 98 of the (153) threats and 28% of all threat logs viewed. The most commonly detected threat was a high severity *PHP remote file inclusion* exploit.

- **Client-Server filesharing:** Nine of the 37 client-server applications displayed 150 threats with FTP being the most heavily targeted based on unique threats (131) and logs viewed (64% of file sharing threat logs viewed). The most common threat found was a critical severity *FTP evasion exploit*.

- **Peer-to-peer filesharing:** These applications used the most bandwidth but exhibited the fewest number of unique threats (22), possibly because some of them use proprietary encryption. Emule and Xunlei exhibited the highest volume of threat logs. Within emule, *Win32.Conficker.C p2p*, a critical severity malware threat, was the most commonly detected, while *ZeroAccess.Gen [botnet] Command and Control Traffic*, also a critical severity malware threat, was found in both emule and Xunlei. Overall, emule, Xunlei and BitTorrent were strongly associated with malware and accounted for the vast majority of malware logs in the file sharing category of applications.

Within file sharing applications, critical, high, or medium severity threat logs were observed in roughly 25% of the 152 variants observed. Interestingly, those that displayed the highest volume – FTP and WebDAV – are commonly used internally to enable the business. This indicates again that while personal use applications do indeed present organizations with business and security risks, protecting the internal variants is of equal or greater importance.

What Is The Business Value of 30 Photo-Video Applications per Network?

The use of video for business purposes is known and proven; it's used for marketing promotions, lead generation, product announcements, training, and education to name just a few examples. For business purposes, the most common applications are YouTube and HTTP video, and perhaps Vimeo.

- There were 112 (out of 1,395) photo-video applications found and they consumed 13% of all bandwidth – roughly the equivalent of 557,000 two-hour high definition movie downloads.
- 96% of the 3,056 networks analyzed had an ***average*** of 30 photo-video application variants on their network.. The question organizations should ask is this: what business value do Netflix-streaming and Hulu Networks have on their network?
- Within the photo-video category, 92 unique threats were detected, ranking it the 10th highest in terms of total threats.

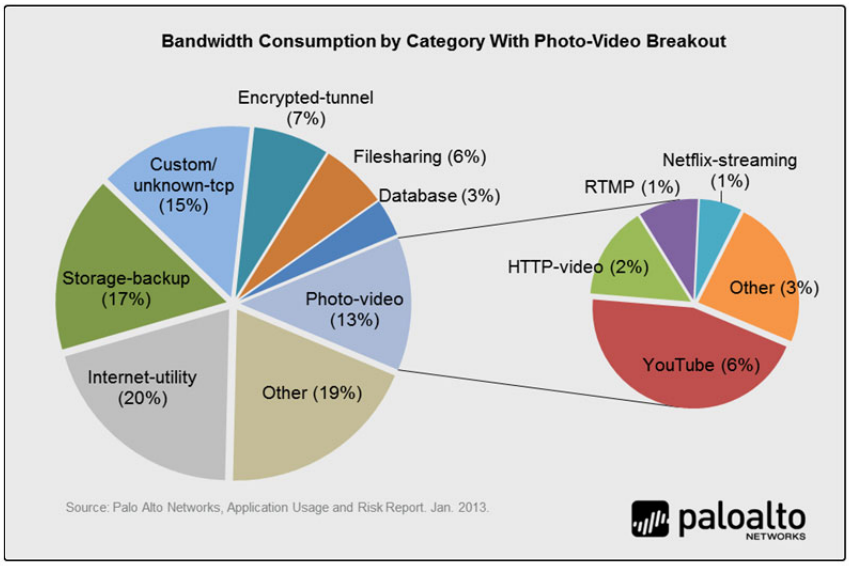

Figure 3: Top photo-video applications, based on bandwidth consumption.

The voluminous bandwidth consumption differences between the top two video applications – YouTube and HTTP video – and the other 100+ variants is consistent within the threat logs observed. There were 25 unique threats found within YouTube and 12 unique threats within HTTP video. Collectively they represent 97% (52% and 45%) of all of all photo-video threat logs observed. This is interesting given that these applications accounted for similar numbers of threat logs, even though YouTube generated about 3 times as much traffic. Within the photo-video category, exploit logs outnumbered malware logs 12:1. Many of the exploits in this category targeted well-known buffer overflow attacks against vulnerable applications such as *VUPlayer*.

Crunchy on the Outside, Tender on the Inside

In the previous section, the applications discussed were primarily Internet based and as such, the traffic will traverse the perimeter firewall and IPS as the first line of defense. The key finding was that the volume of threat logs viewed was relatively low when compared to the bandwidth, the frequency of use and the overall assumption that personal Internet applications represent the highest volume of risk. This finding implies that perimeter security efforts are "crunchy on the outside" effectively stopping some of the threats. Shifting the focus to internal applications, the findings show that security is "tender on the inside", with 97% of the vulnerability exploit logs found in only 10 applications (out of 1,395 found). Nine of these applications are considered high-value assets; they are internal or infrastructure related applications that are integral to many business functions. This data indicates that the strategy of attacking critical resources from inside the network continues to become the rule and not the exception, and will force enterprises to monitor their internal traffic for threats in addition to the perimeter.

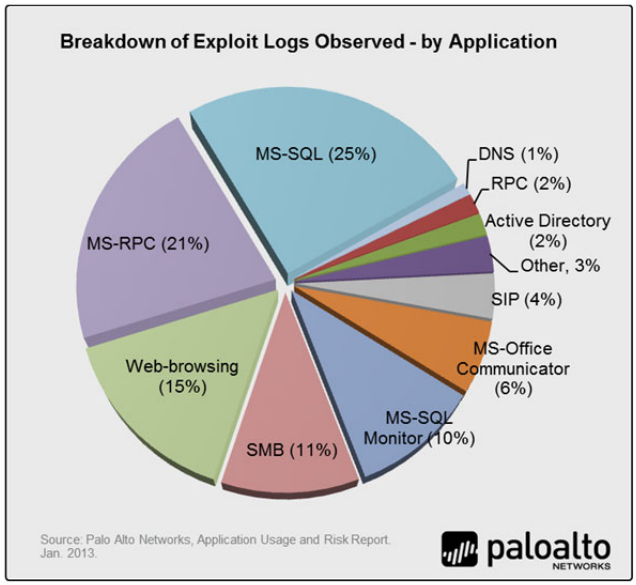

Figure 4: Top 10 applications based on critical, high and medium severity exploit logs observed.

Analyzing the exploits logs a bit further, the data shows that there were 2,016 unique critical, high, and medium severity exploits distributed across roughly 60 million logs. Not surprisingly, the application with the highest concentration of unique exploits was web-browsing, while the highest concentration of logs viewed was within MSSQL and MSRPC, indicating a greater volume of activity within those specific exploits.

	Critical Severity		High Severity		Med. Severity		Totals	
	Exploits	Logs	Exploits	Logs	Exploits	Logs	Exploits	Logs
Active Directory	4	1,146,863	3	436			7	1,147,299
DNS	23	330,351	9	140,670	12	66,456	44	537,477
SMB	81	454,654	109	496,679	32	5,827,984	222	6,779,317
MS-Office Communicator			3	3,775,694			3	3,775,694
MS-RPC	41	263,398	48	12,291,810	13	824,763	102	13,379,971
MS-SQL	1	41			36	15,665,597	37	15,665,638
MS-SQL Monitor	5	6,380,406					5	6,380,406
RPC	2	42	21	1,042,296			23	1,042,338
SIP			16	2,340,384	1	1	17	2,340,385
Web-browsing	469	1,425,067	760	5,325,661	327	2,658,419	1,556	9,409,147
Totals	**626**	**10,000,822**	**969**	**25,413,630**	**421**	**25,043,220**	**2,016**	**60,457,672**

Table 1: Applications with heaviest concentration of unique exploits and related logs.

Interestingly, the concentration of exploits and related logs shown in the previous table does not necessarily correlate to the volume of use shown below, with the exception of web-browsing. The applications with the highest concentration of log activity (MS-SQL, MS-RPC) were neither the most frequently used nor the most heavily used applications.

Application	Frequency of use (n=3,056)	Bytes Consumed (GB)	% of total bytes	Sessions	% of total sessions
Active Directory	60%	8,323	0.1%	178,195,472	0.1%
DNS	99%	51,990	0.4%	46,950,014,080	25.4%
SMB	84%	1,071,798.85	8%	1,370,938,544	1%
MS-Office Communicator	3%	73	0.0%	68,230	0.0%
MS-RPC	76%	77,062	0.6%	1,303,650,727	0.7%
MS-SQL	58%	256,742	2.0%	568,857,780	0.3%
MS-SQL Monitor	56%	17	0.0%	33,633,926	0.0%
RPC	28%	27,615	0.2%	72,479,876	0.0%
SIP	65%	1,468	0.0%	130,363,436	0.1%
Web-browsing	99%	2,121,568	16.8%	44,697,951,094	24.2%

Table 2: Bandwidth consumption and frequency of use for applications with the most exploits and logs.

The data confirms that while there is a significant amount of justifiable concern around the risks of commonly used social media, filesharing and video applications, the real targets are the internal applications that house your most valued assets. SQL databases, SMB file services, Active Directory, and RPC all represent the soft underbelly of the corporate business infrastructure where the intellectual property, corporate information, credit card data, or perhaps social security numbers are stored.

Examples of the exploits viewed within this segment of applications include *Microsoft Windows Server Service Remote Buffer Overflow Vulnerability,* a code execution attack targeting *MS-DS-SMB; Microsoft SMTP Service and Exchange Routing Engine Buffer Overflow Vulnerability* an overflow attack targeting DNS; and *HTTP SQL Injection Attempt* targeting database applications.

As attackers become more advanced, it is apparent that many exploits against critical systems may come from devices inside the network that have been infected with malware. Such an approach allows an attacker to exploit a system without ever crossing a perimeter IPS, underscoring the importance of organizations bringing IPS and threat prevention measures deeper into the network and not exclusively monitoring at the perimeter.

Unknown/Custom Applications Epitomize the 80%-20% Rule

Correlating threats with specific applications (or application features) allows security teams to directly see and control the risk in their networks. However, this full classification of traffic also has the often-unanticipated benefit of revealing custom or unknown application traffic on the network. Custom or unknown traffic is classified as either *custom/unknown-TCP* or *custom/unknown-UDP*, which allows organizations to quickly analyze and systematically manage the traffic and associated risk. Unknown or custom traffic exists on every network observed, typically in the range of 8-10% of all traffic; it can be an internal (custom) application, it can be a commercial application not yet identified, or it can be a threat (custom application). While small in volume, unknown/custom traffic is very high in risk – it epitomizes the 80%-20% rule – a high volume of risk from a low volume of traffic. Determining and then managing this small amount of high risk traffic is particularly significant for controlling threats as attackers and their malware will often customize existing applications and protocols to fit the attacker's needs. The next section of this report provides the first analysis of how unknown or custom application traffic contributes to the risk facing an enterprise.

Custom Traffic Used By Threats

The Palo Alto Networks next-generation firewall begins every analysis by positively classifying the traffic at the application layer. This process can involve multiple layers of decoding and heuristic analysis. Traffic that doesn't match any known application is classified as either *custom/unknown-TCP* or *custom/unknown-UDP* and then analyzed for threats. This process allows organizations to not only identify anomalous traffic but also to directly measure the risk of that anomalous traffic to a network.

Custom Traffic and Malware

Custom/unknown-UDP is the leading application classification associated with malware, accounting for 55% of all malware logs observed. Observing the command and control traffic of malware already present in a particular network generated many of the malware logs.

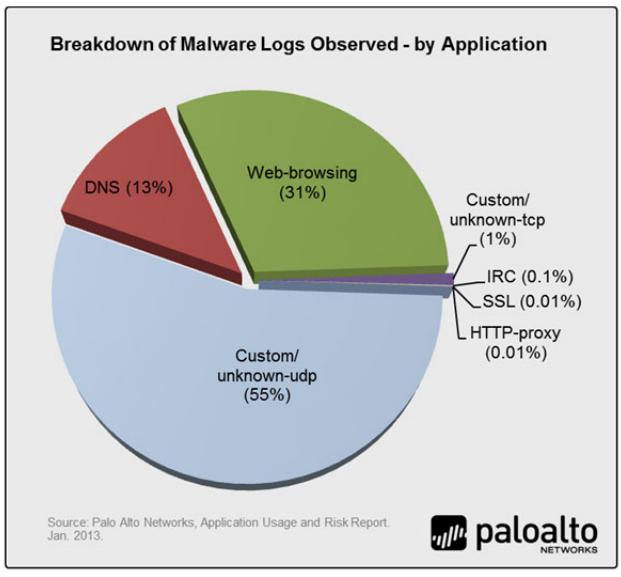

Figure 5: Applications with the highest concentration of malware logs.

As a result of the presence of malware prior to the analysis, a large number of command and control applications were observed within unknown/custom-UDP and TCP, but also masking themselves as more traditionally known malware communication paths such as DNS, IRC, SSL and web-proxies.

The use of custom traffic was observed in use by variety of very popular malware families including the *ZeroAccess Botnet, Conficker, the Poison Ivy RAT, and the IMDDOS denial of service botnet*. The table below displays the frequency and volume of use for the applications with the highest concentration of malware logs. When viewed in tandem with the malware data, the top-3 applications in both cases are applications that are found commonly and can utilize any port.

Application	Frequency of use (n=3,056)	Bytes consumed (GB)	% of total bytes	Sessions	% of total sessions
Web-browsing	99%	2,121,568 GB	16.8%	44,697,951,094	24.2%
DNS	99%	51,990 GB	0.4%	46,950,014,080	25.4%
Custom/unknown-UDP	95%	214,356 GB	1.7%	1,453,837,656	0.8%
SSL	99%	647,372 GB	5.1%	12,761,554,347	6.9%
Custom/unknown-TCP	97%	1,837,853 GB	14.5%	1,126,090,760	0.6%
HTTP-proxy	90%	300,515 GB	2.4%	8,763,997,973	4.7%
IRC	28%	116 GB	0.0%	19,472,012	0.0%

Table 3: Bandwidth consumption and frequency of use for applications with the most malware logs.

Custom Traffic and Exploits

Taken as whole, the number of malware logs was considerably larger than the total number of observed exploits. This trend was consistent within the subset of custom traffic.

- Exploits were a small percentage of custom or unknown traffic: *custom/unknown-TCP* displayed a mere 0.3% of exploit logs (9th overall).
- The overwhelming majority of exploits observed in unknown or custom traffic were critical, with 83% of the exploits classified as "critical" and the remaining 17% classified as either "high" or "medium" threats.

This makes logical sense given that a successful exploit would only be seen once, while malware could potentially be observed many times. However the smaller number of exploits shouldn't undermine the significance. An exploit is critical not only for taking advantage of a target, but also for infecting the target with malware as part of an ongoing persistent attack. Critical exploits are those exploits that can provide an attacker with near-total control over a compromised machine. These critical exploits were specifically the type of exploits that were overwhelmingly observed in customized traffic, underscoring the importance of being able to proactively find and control such traffic. The majority of the exploits used custom traffic targeting IIS web-servers, SQL databases, or were used as part of cross-site scripting attacks.

The Connection Between Attackers and Custom or "Unknown" Traffic

It is no secret that modern attackers have become highly adaptable and customized in order to avoid traditional security, constantly modifying their malware executables in order to bypass existing anti-malware signatures. What is much less understood is that attackers will also heavily modify and customize their communications not only to confuse traditional security, but also for more functional purposes.

Malware will regularly modify peer-to-peer protocols in order to create their own resilient command-and-control communications. For example, the *Zero Access botnet* (and its rootkit) is one of the most popular pieces of malware in the wild, and likewise was the leading malware observed in our data.

This particular malware uses customized peer-to-peer traffic as well as other customized UDP and TCP traffic for communicating with its command and control infrastructure. This traffic is critically important to the reliability and survivability of the botnet in the wild. The malware, having delivered its payload, is sacrificed and the botnet survives to execute the next phase of the attack.

However while this traffic works perfectly well from the attacker's point of view, it does not match any known applications, and was thus classified as being custom or unknown traffic. This technique is incredibly common in malware traffic and is one of the key reasons why custom traffic (especially custom UDP traffic) was the #1 type of traffic associated with malware in the report.

In other cases, attackers may be forced into modifying a protocol as part of the attack itself. For instance, an exploit against a web-server may include malicious content within the HTTP header or frame. This can be as simple as a header overflow against a web-server, or more complex modifications to enable cross-site scripting attacks. However, these very same modifications, which can provide benefits to the attacker, can also help the firewall identify the traffic as customized, which can provide an early indicator that something is wrong with the traffic. The prevalence of these techniques put customized traffic in the top ten sources of observed exploits, even beating out FTP and SMTP in the process.

Potential for Proactive Controls

The analysis clearly shows that customized or modified traffic is highly correlated with threats. This indicates that proactively controlling or blocking "unknown" traffic could easily provide a powerful and untapped strategy for controlling modern threats. This does not imply a replacement of threat signatures, but an augmentation of it. Attackers are in a constant struggle to find new ways of breaking into networks, and security companies are likewise in a constant exercise of delivering new protections for new threats. However, the same creativity that attackers use to find new attack vectors can also be used against them. By blocking or tightly controlling unknown traffic, security teams can greatly reduce their attack surface and proactively manage new, evolving threats in real time. This real-time management of anomalous traffic has become increasingly important for controlling advanced and highly adaptable network threats: an enterprise that blocks *unknown/custom-UDP* would block a huge chunk of malware while stopping the use of Facebook-apps would quickly make the use of social media applications more secure.

Hiding in Plain Sight: SSL and Proxies

The Application Usage and Risk Report has regularly tracked the prevalence of applications that have the potential to circumvent security. This broad category includes applications such as SSL and other encrypted tunnel applications, external proxies, and remote desktop access tools that are indispensable to an enterprise, but also create the potential for threats to enter a network without detection. From an attacker's perspective, these applications are common enough to blend in to a normal network, while allowing the attacker to hide their attacks. As a result, the control of circumventing applications and their contents has become a very important hotspot in the fight between attackers and security, and this tendency was readily observable in the data.

HTTP-proxy was found on 97% of the organizations observed and it represented the third highest volume of malware logs. HTTP-proxy was used by a wide range of threats, including *TDL-4, Rustock, Gozi,* and *Citadel*. Many of these malware families include their own proxy servers for obscuring their true sources as well as tunneling traffic through a network of compromised users. Needless to say, encryption and proxies are critical applications for enterprises, but the data shows that they are also critical applications for attackers.

SSL was the second largest source of malware logs, a significant finding considering that many of the organizations that participated in the analysis did not enable the SSL decryption functionality, implying that the findings almost certainly under-report the true number of malware carried within SSL. In particular, SSL was used by the *Ramnit botnet* as well as variants of the *Poison Ivy RAT*.

When looking at SSL, a common assumption is that it equates primarily to HTTPs (and TCP/443), when in reality, any application can use SSL across any port as a means of security. The analysis

showed that there are 356 applications that can use SSL (on a range of ports) and collectively, they represent 7% of the total bandwidth observed.

The figure below highlights the wide range of applications that can use SSL along with a lack of consistency.

- Of the 26 application subcategories, 21 of them had at least one application that can use SSL.
- Collaborative applications (email, instant messaging, file sharing, and conferencing) showed inconsistency in the use of SSL - more than 50% of the email applications observed do not use SSL while instant messaging had the highest number of applications using SSL across non-standard ports (19).
- 85 of the 356 applications that use SSL never use port 443, nor do they use SSL defined ports (37 hop ports, 28 use port 80, 20 use other ports).

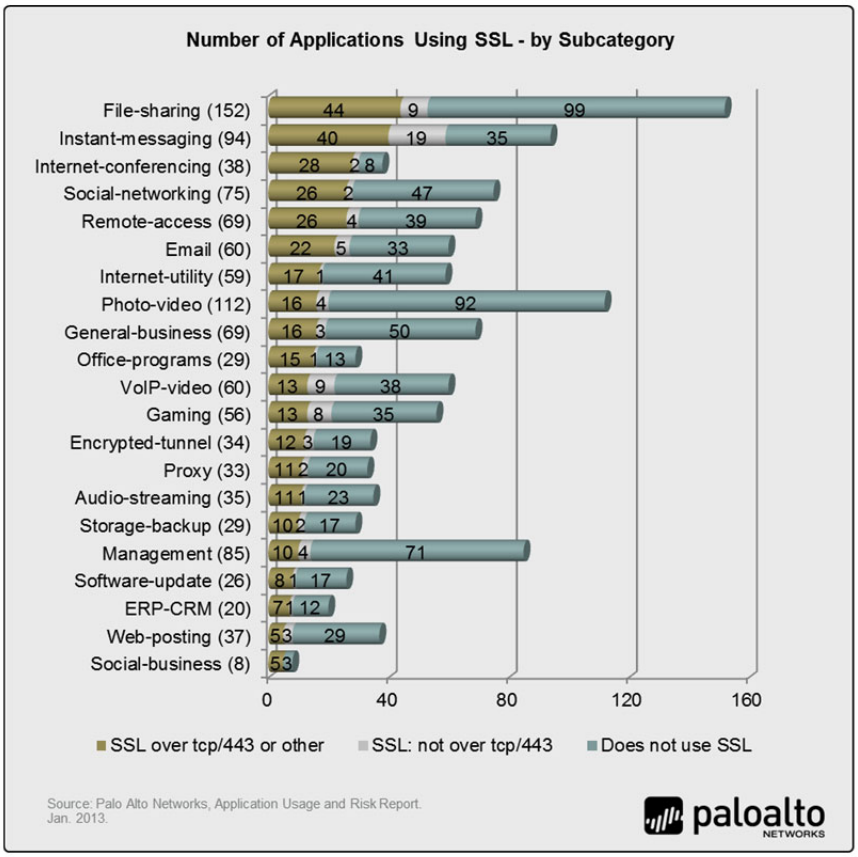

Figure 6: Port and subcategory breakdown of the applications that can use SSL.

Encrypted Tunnels – Security or Evasion?

Expanding the view beyond just SSL to the broader category of encrypted tunnel applications, there are generally two use cases: those that are endorsed by the organization and used for secure communications (e.g., SSH, SSL, IPSec, IKE, ESP, and Secure Access) and those that may not be endorsed and may be used to mask activity or evade controls (e.g., Hamachi, Tor, UltraSurf, and Freegate). The latter group of applications has little business value on most enterprise networks.

The figure below displays how frequently some of the more common encrypted tunnel applications were detected over the last three Application Usage and Risk Reports. The only application that increased in frequency of use was Freegate, an application that is described as *an anti-censorship software for secure and fast Internet access.*

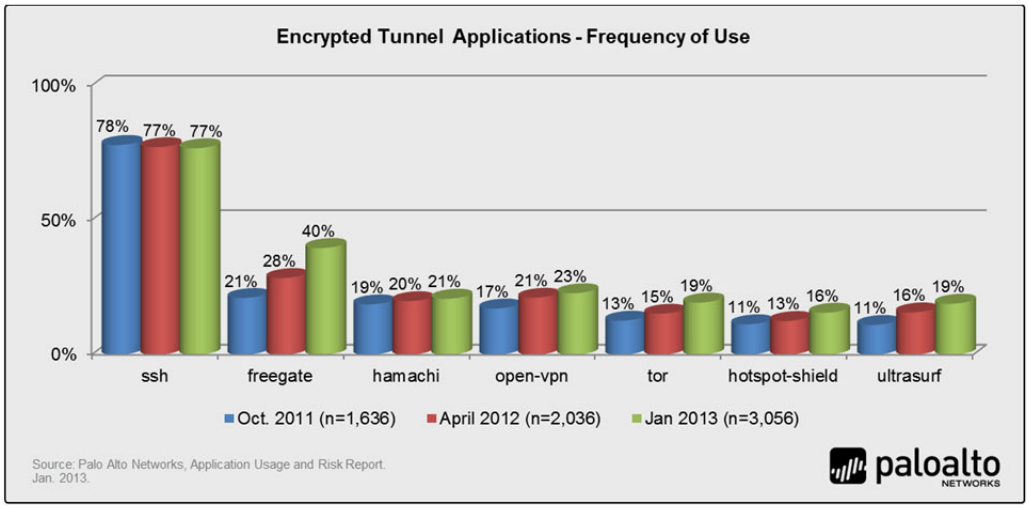

Figure 7: Encrypted tunnel frequency of use – over time.

- OpenVPN, HotSpot Shield, and Hamachi describe themselves more aggressively as VPN security tools, with privacy from censorship (or other controls) as a secondary message. They each have free-to-registered user versions with pay options also available. On a university network, a student using these services may be taking the correct approach to their online activity, which is to stay protected. They may also be trying to mask their activity.

- Tor, UltraSurf, and Freegate all place a much greater emphasis on protecting the user from censorship, with privacy and security as its byproduct. The users of these applications, in most cases, are making a more concerted attempt to stay anonymous by masking their activity.

An additional data point to consider when evaluating the use case for these applications is the default port they use as shown in Table 4. The generally accepted port breakdown is as follows; well known ports (0-1023), registered ports (1024-49151), dynamic/private (49152-65535). Three of the VPN offerings use the default port for SSL (TCP/443) and only one of them, Tor, uses it exclusively.

The applications that use a non-standard or private port make it harder for the port-based security infrastructure used by most universities to know which applications are traversing the network and that lack of knowledge introduces risks.

Target Use is Security/VPN Focused		Target Use is to Evade Controls	
Application	**Default Port**	**Application**	**Default Port**
OpenVPN	TCP/UDP/1194, TCP/443	Tor	TCP/443
HotSpot Shield	TCP/443, 80, dynamic	Freegate	TCP/dynamic
Hamachi	TCP/12975, 10080, UDP 17771	UltraSurf	TCP/dynamic

Table 4: Default port breakdown for encrypted tunnel applications found on participating networks.

To be clear on what this data indicates: these applications are in use within the participating organizations university networks, in some cases, with increasing frequency. It is impossible to determine, from the data collected, which use case is most common – security or masking of activities. The purpose for the discussion is to pose the question: what is the business value?

Summary

The findings within this report are somewhat surprising as they shed some light on the actual levels of risk based on the types of applications and threats therein. Few would have predicted that the 339 social networking, file sharing, and video applications would represent so few threat logs, when compared to the volume of use, the number of variants found in each organization and the overarching assumption that it is a target rich environment. Shifting focus to the internal applications, the most significant surprise was not that these applications are targets; after all, they are integral pieces of most organizations infrastructure. The surprise here was the sheer volume of exploit logs viewed in only 9 applications. The last and final surprise was how malware creators surreptitiously hide their tracks within UDP traffic, which is often viewed as benign when in fact it is the exact opposite. From a security practitioner's perspective, the takeaway would be to first exert some added control over the personal-use applications that represent the highest volume of risk (e.g., Facebook-apps); second, apply significantly more effort to protecting the internal, infrastructure applications from vulnerability exploits; and third, find and block anything that resembles *unknown/custom-UDP*.

About Palo Alto Networks

Palo Alto Networks™ is the network security company. Its innovative platform enables enterprises, service providers, and government entities to secure their networks and safely enable the increasingly complex and rapidly growing number of applications running on their networks. The core of Palo Alto Networks platform is its next-generation firewall, which delivers application, user, and content visibility and control integrated within the firewall through its proprietary hardware and software architecture. Palo Alto Networks products and services can address a broad range of network security requirements, from the data center to the network perimeter, as well as the distributed enterprise, which includes branch offices and a growing number of mobile devices. Palo Alto Networks products are used by more than 10,000 customers in over 100 countries. For more information, visit www.paloaltonetworks.com.

Demographics and Methodology

The latest edition of the Application Usage and Risk Report summarizes 3,056 traffic assessments performed worldwide. The distribution of the participating organizations is distributed fairly equally across three geographic regions: Americas, Mexico, Canada, Asia Pacific/Japan, and Europe. The findings within this report will focus solely on the global view of application traffic with any regional specific variations in usage patterns discussed separately.

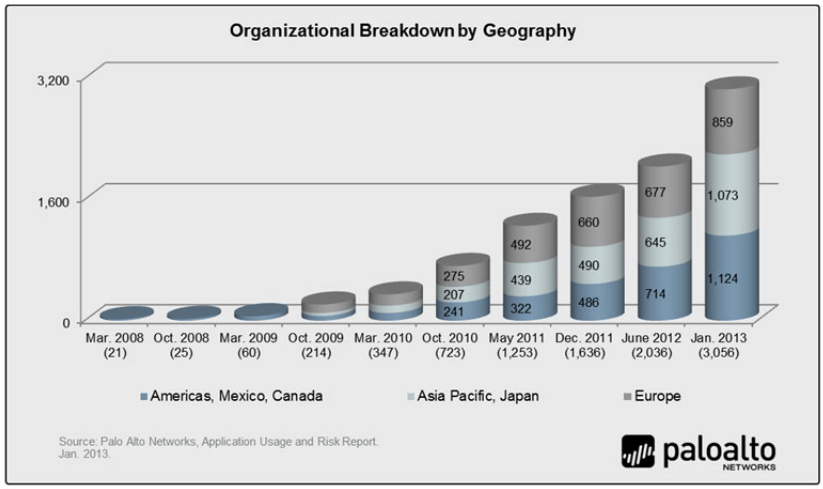

Figure 1: Geographic distribution of participating organizations.

The data in this report is generated via the Palo Alto Networks Application Visibility and Risk assessment process where a Palo Alto Networks next-generation firewall is deployed within the network, where it monitors traffic traversing the network. At the end of the data collection period, usually up to seven days, an Application Visibility and Risk Report is generated that presents the findings along with the associated business risks, and a more accurate picture of how the network is being used. The data from each of the AVR Reports is then aggregated and analyzed, resulting in The Application Usage and Risk Report.