SYMANTEC ENTERPRISE SECURITY

Symantec Internet Security Threat Report

Trends for July–December 06

Volume XI, Published March 2007

Dean Turner
Executive Editor
Symantec Security Response

Stephen Entwisle
Senior Editor
Symantec Security Response

Marci Denesiuk
Editor
Symantec Security Response

Marc Fossi
Analyst
Symantec Security Response

Joseph Blackbird
Analyst
Symantec Security Response

David McKinney
Analyst
Symantec Security Response

Ronald Bowes
Analyst
Symantec Security Response

Nicholas Sullivan
Analyst
Symantec Security Response

Peter Coogan
Analyst
Symantec Security Response

Candid Wueest
Analyst
Symantec Security Response

Ollie Whitehouse
Security Architect—Advanced Threat Research
Symantec Security Response

Zulfikar Ramzan
Analyst—Advanced Threat Research
Symantec Security Response

Contributors

David Cole
Director Product Management
Symantec Security Response

Peter Szor
Security Architect
Symantec Security Response

David Cowings
Sr. Business Intelligence Manager
Symantec Business Intelligence

Shravan Shashikant
Pr. Business Intelligence Manager
Symantec Business Intelligence

Igor Moochnick
Sr. Software Engineer
Symantec Instant Messaging Security

Volume XI, Published March 2007

Symantec Internet Security Threat Report

Contents

Internet Security Threat Report Volume XI Executive Summary

Over the past two reporting periods, Symantec has observed a fundamental shift in Internet security activity. The current threat environment is characterized by an increase in data theft and data leakage, and the creation of malicious code that targets specific organizations for information that can be used for financial gain.

Instead of exploiting high-severity vulnerabilities in direct attacks, attackers are now discovering and exploiting medium-severity vulnerabilities in third-party applications, such as Web applications and Web browsers. Those vulnerabilities are often used in "gateway" attacks, in which an initial exploitation takes place not to breach data immediately, but to establish a foothold from which subsequent, more malicious attacks can be launched.

Symantec has observed high levels of malicious activity across the Internet, with increases in phishing, spam, bot networks, Trojans, and zero-day threats. However, whereas in the past these threats were often used separately, attackers are now refining their methods and consolidating their assets to create global networks that support coordinated criminal activity.

This has resulted in an increasing interoperability between diverse threats and methods. For example, targeted malicious code may take advantage of Web-enabled technologies and third-party applications to install a back door, which then downloads and installs bot software. These bots can, in turn, be used to distribute spam, host phishing sites, or launch attacks in such a way as to create a single coordinated network of malicious activity. Once entrenched, these networks can be used in concert as global networks of malicious activity that support their own continued growth.

This volume of the *Internet Security Threat Report* will offer an overview of threat activity that took place between July 1 and December 31, 2006. This brief summary and the discussion that follows will offer a synopsis of the data and trends that are presented in the main report. Symantec will continue to monitor and assess threat activity in order to best prepare consumers and enterprises for the complex Internet security issues to come.

Internet Security Threat Report Overview

The Symantec *Internet Security Threat Report* provides a six-month update of Internet threat activity. It includes analysis of network-based attacks, a review of known vulnerabilities, and highlights of malicious code. It also assesses numerous issues related to online fraud, including phishing, spam, and security risks such as adware, spyware, and misleading applications. This summary of the *Internet Security Threat Report* will alert readers to current trends and impending threats. In addition, it will offer recommendations for protection against and mitigation of these concerns. This volume covers the six-month period from July 1 to December 31, 2006.

Symantec has established some of the most comprehensive sources of Internet threat data in the world. The Symantec™ Global Intelligence Network, which includes Symantec DeepSight™ Threat Management System and Symantec™ Managed Security Services, tracks attack activity across the entire Internet. It consists of over 40,000 sensors monitoring network activity in over 180 countries. As well, Symantec gathers malicious code data along with spyware and adware reports from over 120 million client, server, and gateway systems that have deployed Symantec's antivirus products.

Symantec operates one of the most popular forums for the disclosure and discussion of vulnerabilities on the Internet, the BugTraq™ mailing list, which has approximately 50,000 direct subscribers who contribute, receive, and discuss vulnerability research on a daily basis.[1] Symantec also maintains one of the world's most comprehensive vulnerability databases, currently consisting of over 20,000 vulnerabilities (spanning more than a decade) affecting more than 45,000 technologies from over 7,000 vendors. Symantec also tracks and assesses certain criminal activities using online fraud monitoring tools.

Finally, the Symantec Probe Network, a system of over two million decoy accounts, attracts email messages from 20 different countries around the world, allowing Symantec to gauge global spam and phishing activity. These resources give Symantec analysts unparalleled sources of data with which to identify emerging trends in attacks and malicious code activity. Symantec also gathers phishing information through the Symantec Phish Report Network, an extensive antifraud community of enterprises and consumers. Members of the network contribute and receive fraudulent Web site addresses for alerting and filtering across a broad range of solutions.

The Symantec *Internet Security Threat Report* is grounded principally on the expert analysis of data provided by all of these sources. Based on Symantec's expertise and experience, this analysis yields a highly informed commentary on current Internet threat activity. By publishing the analysis of Internet security activity in the Symantec *Internet Security Threat Report*, Symantec hopes to provide enterprises and consumers with the information they need to help effectively secure their systems now and in the future.

Executive Summary Highlights

The following section will offer a brief summary of the security trends that Symantec observed during this period based on data provided by the sources listed above. This summary includes all of the metrics that are included in the main report. Following this overview, the Executive Summary will discuss selected metrics in greater depth.

Attack Trends Highlights

- The government sector accounted for 25 percent of all identity theft-related data breaches, more than any other sector.
- The theft or loss of a computer or other data-storage medium made up 54 percent of all identity theft-related data breaches during this period.
- The United States was the top country of attack origin, accounting for 33 percent of worldwide attack activity.
- Symantec recorded an average of 5,213 denial of service (DoS) attacks per day, down from 6,110 in the first half of the year.
- The United States was the target of most DoS attacks, accounting for 52 percent of the worldwide total.
- The government sector was the sector most frequently targeted by DoS attacks, accounting for 30 percent of all detected attacks.
- Microsoft Internet Explorer was targeted by 77 percent of all attacks specifically targeting Web browsers.
- Home users were the most highly targeted sector, accounting for 93 percent of all targeted attacks.

[1] The BugTraq mailing list is hosted by SecurityFocus (http://www.securityfocus.com). Archives are available at http://www.securityfocus.com/archive/1

- Symantec observed an average of 63,912 active bot-infected computers per day, an 11 percent increase from the previous period.
- China had 26 percent of the world's bot-infected computers, more than any other country.
- The United States had the highest number of bot command-and-control computers, accounting for 40 percent of the worldwide total.
- Beijing was the city with the most bot-infected computers in the world, accounting for just over five percent of the worldwide total.
- The United States accounted for 31 percent of all malicious activity during this period, more than any other country.
- Israel was the highest ranked country for malicious activity per Internet user, followed by Taiwan and Poland.
- Fifty-one percent of all underground economy servers known to Symantec were located in the United States, the highest total of any country.
- Eighty-six percent of the credit and debit cards advertised for sale on underground economy servers known to Symantec were issued by banks in the United States.

Vulnerability Trends Highlights

- Symantec documented 2,526 vulnerabilities in the second half of 2006, 12 percent higher than the first half of 2006, and a higher volume than in any other previous six-month period.[2]
- Symantec classified four percent of all vulnerabilities disclosed during this period as high severity, 69 percent were medium severity, and 27 percent were low severity.
- Sixty-six percent of vulnerabilities disclosed during this period affected Web applications.
- Seventy-nine percent of all vulnerabilities documented in this reporting period were considered to be easily exploitable.
- Seventy-seven percent of all easily exploitable vulnerabilities affected Web applications, and seven percent affected servers.
- Ninety-four percent of all easily exploitable vulnerabilities disclosed in the second half of 2006 were remotely exploitable.
- In the second half of 2006, all the operating system vendors that were studied had longer average patch development times than in the first half of the year.
- Sun Solaris had an average patch development time of 122 days in the second half of 2006, the highest of any operating system.
- Sixty-eight percent of the vulnerabilities documented during this period were not confirmed by the affected vendor.
- The window of exposure for vulnerabilities affecting enterprise vendors was 47 days.
- Symantec documented 54 vulnerabilities in Microsoft Internet Explorer, 40 in the Mozilla browsers, and four each in Apple Safari and Opera.

[2] The Symantec *Internet Security Threat Report* has been tracking vulnerabilities in six-month periods since January 2002.

- Mozilla had a window of exposure of two days, the shortest of any Web browser during this period.
- Twenty-five percent of exploit code was released less than one day after vulnerability publication. Thirty-one percent was released in one to six days after vulnerability publication.
- Symantec documented 12 zero-day vulnerabilities during this period, a significant increase from the one documented in the first half of 2006.
- Symantec documented 168 vulnerabilities in Oracle database implementations, more than any other database.

Malicious Code Trends Highlights

- Of the top ten new malicious code families detected in the last six months of 2006, five were Trojans, four were worms, and one was a virus.
- The most widely reported new malicious code family this period was that of the Stration worm.[3]
- Symantec honeypot computers captured a total of 136 previously unseen malicious code threats between July 1 and December 31, 2006.
- During this period, 8,258 new Win32 variants were reported to Symantec, an increase of 22 percent over the first half of 2006.
- Worms made up 52 percent of the volume of malicious code threats, down from 75 percent in the previous period.
- The volume of Trojans in the top 50 malicious code samples reported to Symantec increased from 23 percent to 45 percent.
- Trojans accounted for 60 percent of the top 50 malicious code samples when measured by potential infections.
- Polymorphic threats accounted for three percent of the volume of top 50 malicious code reports this period, up from one percent in the two previous periods.
- Bots made up only 14 percent of the volume of the top 50 malicious code reports.
- Threats to confidential information made up 66 percent of the top 50 malicious code reported to Symantec.
- Keystroke logging threats made up 79 percent of confidential information threats by volume of reports, up from 57 percent in the first half of the year and 66 percent in the second half of 2005.
- Seventy-eight percent of malicious code that propagated did so over SMTP, making it the most commonly used propagation mechanism.
- Malicious code using peer-to-peer to propagate rose from 23 percent of all propagating malicious code in the first six months of 2006 to 29 percent in the last half of the year.
- The majority of malicious code reports during this period originated in the United States.

[3] http://www.symantec.com/security_response/writeup.jsp?docid=2006-092111-0525-99

- During the second half of 2006, 23 percent of the 1,318 documented malicious code instances exploited vulnerabilities.
- MSN Messenger was affected by 35 percent of new instant messaging threats in the second half of the year.

Phishing, Spam, and Security Risks Highlights

- The Symantec Probe Network detected a total of 166,248 unique phishing messages, a six percent increase over the first six months of 2006. This equates to an average of 904 unique phishing messages per day for the second half of 2006.
- Symantec blocked over 1.5 billion phishing messages, an increase of 19 percent over the first half of 2006.
- Throughout 2006, Symantec detected an average of 27 percent fewer unique phishing messages on weekends than the weekday average of 961.
- On weekends, the number of blocked phishing attempts was seven percent lower than the weekday average of 7,958,323 attempts per day.
- Organizations in the financial services sector accounted for 84 percent of the unique brands that were phished during this period.
- Forty-six percent of all known phishing Web sites were located in the United States, a much higher proportion than in any other country.
- Between July 1 and December 31, 2006, spam made up 59 percent of all monitored email traffic. This is an increase over the first six months of 2006 when 54 percent of email was classified as spam.
- Sixty-five percent of all spam detected during this period was written in English.
- In the last six months of 2006, 0.68 percent of all spam email contained malicious code. This means that one out of every 147 spam messages blocked by Symantec Brightmail AntiSpam contained malicious code.
- Spam related to financial services made up 30 percent of all spam during this period, the most of any category.
- During the last six months of 2006, 44 percent of all spam detected worldwide originated in the United States.
- The United States hosted the largest proportion of spam zombies, with 10 percent of the worldwide total.
- The most commonly reported security risk was an adware program named ZangoSearch.
- All of the top ten security risks reported in the last six months of 2006 employ at least one anti-removal technique compared to only five of the top ten security risks in the last reporting period.
- All of the top ten security risks reported during this period employ self-updating.
- Potentially unwanted applications accounted for 41 percent of reports in the top ten new security risks in the second half of 2006.
- Misleading application detections increased by 40 percent in the second half of 2006.

Executive Summary Discussion

This section will discuss selected metrics from the *Internet Security Threat Report* in greater depth, providing analysis and discussion of the trends indicated by the data. The following metrics will be discussed:

- Malicious activity by country
- Data breaches that could lead to identity theft
- Underground economy servers
- Zero-day vulnerabilities
- Threats to confidential information
- Malicious code types
- Phishing
- Spam
- Bot-infected computers

Malicious activity by country

For the first time, in this volume of the *Internet Security Threat Report*, Symantec is evaluating the countries in which malicious activity takes place or originates. To determine this, Symantec has compiled geographical data on numerous malicious activities, namely: bot-infected computers, bot command-and-control servers, phishing Web sites, malicious code reports, spam relay hosts, and Internet attacks.

Between July 1 and December 31, 2006, the United States was the top country for malicious activity, accounting for 31 percent of the worldwide total (table 1). For each of the malicious activities taken into account for this measurement, the United States ranked number one by a large margin with the exception of bot-infected computers. It ranked second for that criterion, 12 percentage points lower than China.

Overall Rank	Country	Overall Proportion	Malicious Code Rank	Spam Host Rank	Command and Control Server Rank	Phishing Host Rank	Bot Rank	Attack Rank
1	United States	31%	1	1	1	1	2	1
2	China	10%	3	2	4	8	1	2
3	Germany	7%	7	3	3	2	4	3
4	France	4%	9	4	14	4	3	4
5	United Kingdom	4%	4	13	9	3	6	6
6	South Korea	4%	12	9	2	9	11	9
7	Canada	3%	5	23	5	7	10	5
8	Spain	3%	13	5	15	16	5	7
9	Taiwan	3%	8	11	6	6	7	11
10	Italy	3%	2	8	10	14	12	10

Table 1. Malicious activity by country
Source: Symantec Corporation

The high degree of malicious activity originating in the United States is likely driven by the expansive Internet infrastructure there. The United States accounts for 19 percent of the world's Internet users.[4] Furthermore, the number of broadband Internet users in that country grew by 14 percent between December 2005 and July 2006.[5] Despite the relatively well developed security infrastructure in the United States, the high number of Internet-connected computers there presents more targets for attackers to compromise for malicious use. Symantec predicts that the United States will remain the highest ranked country for malicious activity until another country exceeds it in numbers of Internet users and broadband connectivity.

China was the second highest country for malicious activity during this six-month reporting period, accounting for 10 percent of all worldwide malicious activity. Germany was third, with seven percent. The prominence of both of these countries can likely be attributed to the high number of Internet users there, as well as the rapid growth in the country's Internet infrastructure.

Having determined the top countries by malicious activity, Symantec evaluated the top 25 of these countries according to the number of Internet users located there. This measure is intended to remove the bias of high numbers of Internet users from the "Malicious activity by country" measurement. The percentage assigned to each country in this discussion equates to the proportion of malicious activity that could be attributed to a single (average) Internet user in that country.

Israel was the most highly ranked country for malicious activity per Internet user. If one person from each of the top 25 countries were to represent their country's Internet-connected population, the average Internet user in Israel would carry out nine percent of the group's malicious activity. Taiwan had the second most malicious activity per Internet user, accounting for eight percent of the sample group's activity. Poland ranked third, accounting for six percent.

Data breaches that could lead to identity theft

Identity theft is an increasingly prevalent security issue. Organizations that store and manage personal identification information must take care to ensure the confidentiality and integrity of such data. Any compromise that results in the leakage of personal identity information could result in a loss of public confidence, legal liability, and/or costly litigation.

In the second half of 2006, the government sector accounted for the majority of data breaches that could lead to identity theft, making up 25 percent of the total (figure 1). Government organizations store a lot of personal information that could be used for the purposes of identity theft. Furthermore, they often consist of numerous semi-independent departments. As a consequence, sensitive personal identification information may be stored in separate locations and be available to numerous people. This increases the opportunity for attackers to gain unauthorized access to this data. Governments may also be more likely to report such breaches than private organizations, which may fear negative market reaction.

[4] http://www.internetworldstats.com
[5] http://www.oecd.org/document/9/0,2340,en_2649_34225_37529673_1_1_1_1,00.html

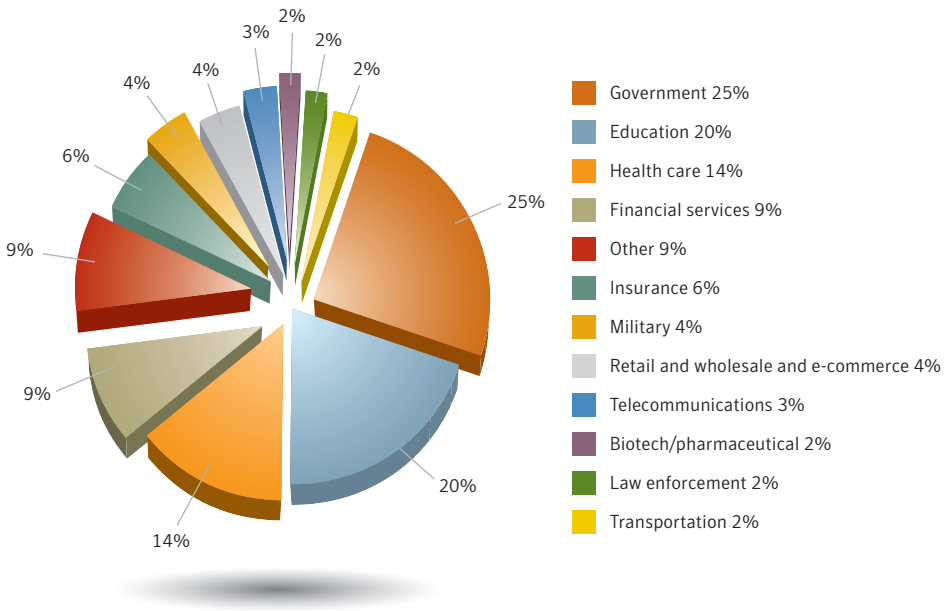

Figure 1. Data breaches that could lead to identity theft by sector
Source: Based on data provided by Privacy Rights Clearinghouse and Attrition.org

During this period, 54 percent of all data breaches that could lead to identity theft were caused by the theft or loss of a computer or data-storage medium (such as a USB memory key or back-up media). Twenty-eight percent of such breaches were caused by insecure policy, which includes a failure to develop, implement, and/or comply with adequate security policy. For example, this could mean posting personal identification information on a publicly available Web site or sending it through unencrypted email.

Most breaches of this type are avoidable. In the case of theft or loss, the compromise of data could be averted by encrypting all sensitive data. This would ensure that even if the data were lost or stolen, it would not be accessible to unauthorized third parties. This step should be part of a broader security policy that organizations should develop, implement, and enforce in order to ensure that all sensitive data is protected from unauthorized access.

Underground economy servers

Underground economy servers are used by criminals and criminal organizations to sell stolen information, typically for subsequent use in identity theft. This data can include government-issued identity numbers, credit cards, bank cards and personal identification numbers (PINs), user accounts, and email address lists.

During the second half of 2006, 51 percent of all underground economy servers known to Symantec were located in the United States, the highest total of any country (figure 2). The prominence of the United States is no surprise, as the expansive Internet infrastructure and continual broadband growth there create numerous opportunities for criminals to carry out malicious activities. Sweden ranked second, accounting for 15 percent of the worldwide total, and Canada ranked third, accounting for seven percent.

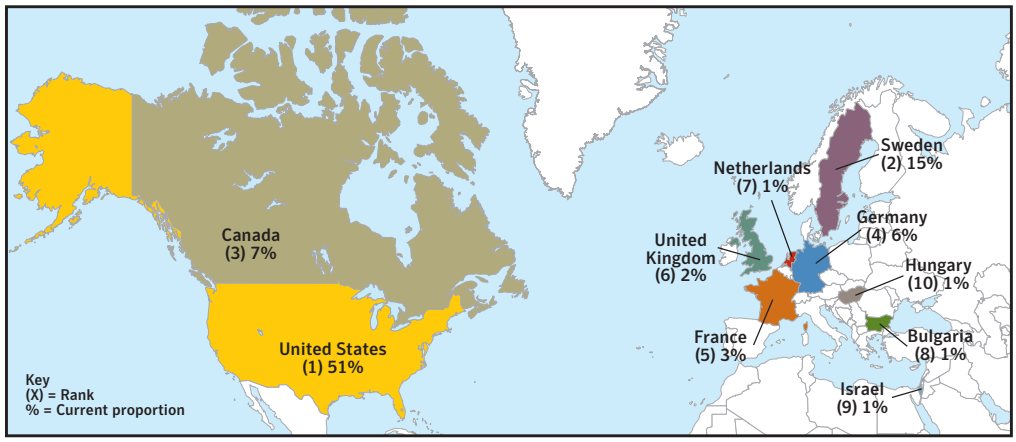

Figure 2. Location of underground economy servers
Source: Symantec Corporation

By far the most credit and debit cards advertised for sale on underground economy servers were issued by banks in the United States. The prominence of the United States is not entirely unexpected, as the vast majority of the data breaches that could lead to identity theft reported during this period took place there.

In order to reduce the likelihood of facilitating identity theft, it is important that organizations take the necessary steps to protect data stored on their computers or transmitted over networks. This should include the development and implementation of a policy requiring that all sensitive data is encrypted. This would ensure that, even if the data were lost or stolen, it would not be accessible. This step should be part of a broader security policy that organizations should develop and implement in order to ensure that any sensitive data is protected from unauthorized access.

Zero-day vulnerabilities

A zero-day vulnerability is one for which there is sufficient public evidence to indicate that the vulnerability has been exploited in the wild prior to being publicly known. It may not have been known to the vendor prior to exploitation, and the vendor had not released a patch at the time of the exploit activity.

Zero-day vulnerabilities represent a serious threat in many cases because there is no patch available for them and because they will likely be able to evade purely signature-based detection. They may be used in targeted attacks and in the propagation of malicious code. As Symantec predicted in Volume IX of the *Internet Security Threat Report*, a black market for zero-day vulnerabilities has emerged that has the potential to put them into the hands of criminals and other interested parties.[6]

In the second half of 2006, Symantec documented 12 zero-day vulnerabilities. This is a significant increase over the first half of 2006 and the second half of 2005 when only one zero-day vulnerability was documented for each reporting period.

The second half of 2006 saw a large number of high-profile zero-day vulnerabilities. This activity peaked in September of 2006, when four zero-day vulnerabilities were made known. The majority of these were client-side vulnerabilities that affected Office applications, Internet Explorer, and ActiveX controls. Many of these may have been discovered through the use of fuzzing technologies.

[6] Symantec *Internet Security Threat Report*, Volume IX (March 2006): http://eval.veritas.com/mktginfo/enterprise/white_papers/ent-whitepaper_symantec_internet_security_threat_report_ix.pdf : p. 21

Zero-day threats appear to be occurring more frequently than in the past. While it is believed that zero-day vulnerabilities have previously posed a threat, the recent rise in incidents may be partially accounted for by increasing capabilities to detect these attacks in the wild. Such capabilities include improved vulnerability-handling procedures within organizations, improved cooperation between enterprises and vendors, and better technologies for the detection and analysis of exploits and malicious code.

In order to protect against zero-day vulnerabilities, Symantec recommends that administrators deploy intrusion detection/intrusion prevention systems (IDS/IPS) and regularly updated antivirus software. Security vendors may be able to provide rapid response to recently discovered zero-day vulnerabilities in the wild by developing and implementing new or updated IDS/IPS and antivirus signatures before the affected vendor has released a patch. Generic signatures may also block zero-day threats, as may behavior-blocking solutions and heuristic technologies.

Threats to confidential information

Some malicious code programs are designed specifically to expose confidential information that is stored on an infected computer. Threats to confidential information are a particular concern because of their potential use in criminal activities. Compromises of this nature can result in significant financial loss, particularly if credit card information or banking details are exposed.

Exposure of confidential information within the enterprise can lead to significant data leakage. If it involves customer-related data—such as credit card information—it can severely undermine customer confidence as well as violate local laws. Sensitive corporate information, including financial details, business plans, and proprietary technologies, could also be leaked from compromised computers.

In the last six months of 2006, threats to confidential information made up 66 percent of the volume of the top 50 malicious code reported to Symantec (figure 3). This is an increase over the 48 percent reported in the first half of the year and the 55 percent reported during the second half of 2005.

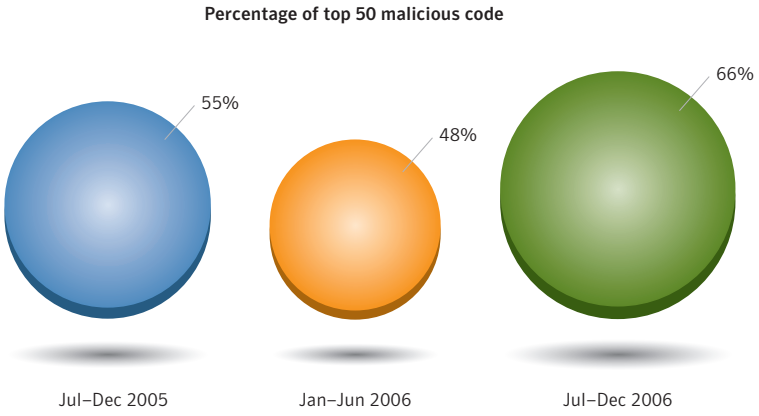

Figure 3. Threats to confidential information
Source: Symantec Corporation

In the second half of the 2006, threats that allow remote access, such as back doors, made up 84 percent of the volume of confidential information threats. Keystroke logging threats made up 79 percent of confidential information threats by volume of reports, and threats that could be used to export user data accounted for 62 percent of confidential information threats during this reporting period.

Malicious code types

During the current reporting period, worms made up 52 percent of the volume of malicious code threats, down from 75 percent in the previous period.[7] However, the number of unique samples of worms in the top 50 malicious code reports remained fairly constant over the last six months of 2006. During this period, 36 worms were reported to Symantec, compared to 38 in the previous period.

The volume of Trojans in the top 50 malicious code samples reported to Symantec increased significantly in the last six months of 2006. During this period, they constituted 45 percent of the volume of the top 50 malicious code samples, a significant increase over the 23 percent last period and the 38 percent reported in the second half of 2005.

As is discussed in the "Future Watch" section of this report, attackers are moving towards staged downloaders, also referred to as modular malicious code. These are small, specialized Trojans that download and install other malicious programs such as a back door or worm. During the current period, 75 percent of the volume of the top 50 malicious code reports contained a modular component such as this.

For the first time, in this edition of the *Internet Security Threat Report*, Symantec is assessing malicious code according to the number of unique samples reported to Symantec and the number of potential infections. This is an important distinction. In some cases, a threat that may create a large number of reports may not cause a large number of potential infections and *vice versa*.

For instance, worms made up 52 percent of malicious code reports in the second half of 2006, but caused only 37 percent of potential infections (figure 4). The main reason for this is that mass-mailing worms generate a significant number of email messages to which they attach their malicious code. Each message that is detected will generate a malicious code report. Because of the high volume of email that one worm can generate, a single infection can result in many reports. However, once a malicious code sample is detected, antivirus signatures are quickly developed that can protect against subsequent infections by that sample. Furthermore, gateway policies and technologies can block the executable attachments that also come with a mass mailer. So, only a small percentage of the high volume of email messages will result in additional infections.

[7] It is important to note that a malicious code sample can be classified in more than one threat type category. For example, bots such as variants of the Mytob family are classified as both a worm and a back door. As a result, cumulative percentages of threat types in the top 50 malicious code reports may exceed 100.

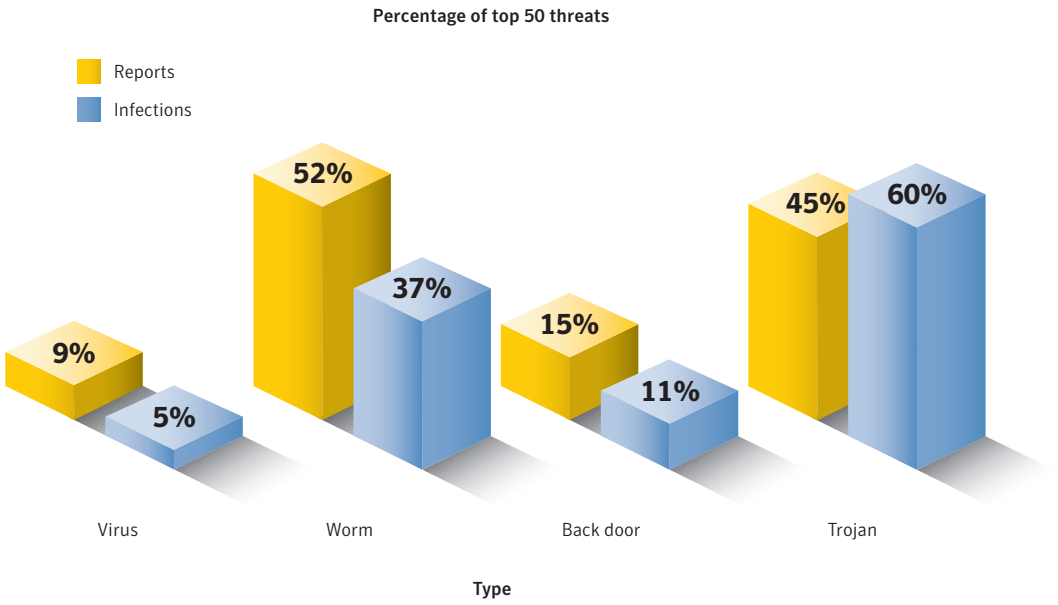

Figure 4. Malicious code types, by reports and by potential infections, July–December 2006
Source: Symantec Corporation

Trojans, on the other hand, only constituted 45 percent of the volume of the top 50 malicious code reports during the last six months of 2006. However, they accounted for 60 percent of potential infections by the top 50 malicious code samples during the same period. Since Trojans do not contain any propagation mechanisms, they do not proliferate as widely as mass-mailing worms, resulting in fewer reports. Because they are frequently installed by exploiting Web browser and zero-day vulnerabilities, a Trojan report is more likely to be the result of an infection. Consequently, the ratio of potential infections to reports is likely to be higher for Trojans than for worms.

Phishing

Over the last six months of 2006, the Symantec Probe Network detected a total of 166,248 unique phishing messages, an average of 904 per day. This total is a six percent increase over the first six months of 2006 when 157,477 unique phishing messages were detected.

In the second half of 2006, Symantec blocked over 1.5 billion phishing messages, an increase of 19 percent over the first half of 2006, and a six percent increase over the second half of 2005. This means that Symantec blocked an average of 8.48 million phishing emails per day over the last six months of 2006.

In the second half of 2006, 46 percent of all known phishing Web sites were located in the United States, a much higher proportion than in any other country. This is likely because a large number of Web-hosting providers—particularly free Web hosts—are located in the United States. Furthermore, the United States has the highest number of Internet users in the world, and it is home to a large number of Internet-connected organizations, both large and small.

Most of the unique brands phished in the last six months of 2006 were in the financial services sector. Organizations in that sector accounted for 84 percent of the brands that were used in phishing attacks this period. This is not surprising, as most phishing attacks are motivated by profit. A successful phishing attack on a financial entity is likely to yield information that an attacker could subsequently use for financial gain.

Spam

Between July 1 and December 31, 2006, spam made up 59 percent of all email traffic monitored by Symantec. This is an increase over the first six months of 2006 when Symantec classified 54 percent of email as spam.

The most common type of spam detected in the latter half of 2006 was related to financial services, which made up 30 percent of all spam on the Internet during this period. Spam related to health services and products made up 23 percent of all spam, while spam related to commercial products was the third most common type of spam, accounting for 21 percent of the total.

The rise in financially-related spam was due mainly to a noticeable increase in stock market "pump and dump" spam. Pump and dump is the name given to schemes in which criminals profit by creating an artificial interest in a stock they own. They buy a penny stock when the price is low. They then artificially pump up demand for the stock by sending out spam that appears to be from a respected stock advisor, but that actually contains false predictions of high performance for the stock. Recipients of the message, trusting the spam content, buy the stock, creating demand for it and thereby raising the price. When the prices are high, the perpetrators of the scheme sell their stock for a profit.[8]

This type of spam has been proven to allow the spammers to generate revenue directly and almost immediately.[9] This alone is likely to make it more appealing than other types of spam.

A spam zombie is a computer infected with a bot or some other malicious code that allows email messages to be relayed through it. Between July 1 and December 31, 2006, ten percent of all spam zombies were located in the United States, making it the highest country in this category. During this period, the United States was one of the top reporting countries for bots such as Spybot and Mytob, which are commonly used to send spam.

China and Germany were the second and third highest countries for spam zombies, hosting nine and eight percent of the worldwide total, respectively. The small variance between the top countries hosting spam zombies is quite different from the distribution of bots during this period. This indicates that not all spam zombies are necessarily bots and that not all bots are used to send spam.

Bot-infected computers

Bots are programs that are covertly installed on a user's machine in order to allow an unauthorized user to control the computer remotely through a communication channel such as IRC. These channels allow the remote attacker to control a large number of compromised computers over a single, reliable channel in a bot network, which can then be used to launch coordinated attacks.

[8] http://www.sec.gov/answers/pumpdump.htm
[9] http://papers.ssrn.com/sol3/papers.cfm?abstract_id=920553

Bots allow for a wide range of functionality and most can be updated to assume new functionality by downloading new code and features. Bots can be used by external attackers to perform DoS attacks against an organization's Web site. Furthermore, bots within an organization's network can be used to attack other organizations' Web sites, which can have serious business and legal consequences. Bots can be used by attackers to harvest confidential information from compromised computers, which can lead to identity theft. Bots can also be used to distribute spam and phishing attacks, as well as spyware, adware, and misleading applications.

Between July 1 and December 31, 2006, Symantec observed an average of 63,912 active bot-infected computers per day. This is an 11 percent increase over the previous six-month period. Furthermore, Symantec observed 6,049,594 distinct bot-infected computers during the current reporting period, a 29 percent increase from the previous period. This increase is largely driven by a peak in bot activity in September when a number of vulnerabilities were disclosed that were actively exploited by bots.

Command-and-control servers are computers that bot network owners use to relay commands to bot-infected computers on their networks. In the last six months of 2006, Symantec identified 4,746 bot command-and-control servers, a 25 percent decrease from the first six months of 2006.

A drop in the number of command-and-control servers combined with a rise in the number of bot-infected computers indicates that, on average, bot networks are increasing in size. Bot networks are thus becoming more consolidated. Consolidated bot networks will likely mean that organizations will have to deal with a well entrenched, experienced, and dedicated group of bot network owners instead of a population of hobby hackers.

It could also signal a fundamental change in the way bots communicate with one another. Symantec has seen bots that are structured on a peer-to-peer model, in which the machines connect together rather than connecting to a central command-and-control server. Symantec has also observed that command-and-control servers are beginning to adopt encryption, so that they are less visible.

China had the highest number of bot-infected computers during the second half of 2006, accounting for 26 percent of the worldwide total (figure 5). This is an increase of six percentage points over the previous six months. This increase was driven by a rise in the number of bots in the country rather than a decrease in other countries. This coincides with and illustrates a trend that Symantec first discussed in 2005, in which bot activity in China appeared to be increasing.[10] During the second half of 2006, the United States had the second highest number of bot-infected computers, accounting for 14 percent of the worldwide total.

[10] Symantec *Internet Security Threat Report*, Volume VII (March 2005): http://eval.veritas.com/mktginfo/enterprise/white_papers/ent-whitepaper_symantec_internet_security_threat_report_vii.pdf : p. 26

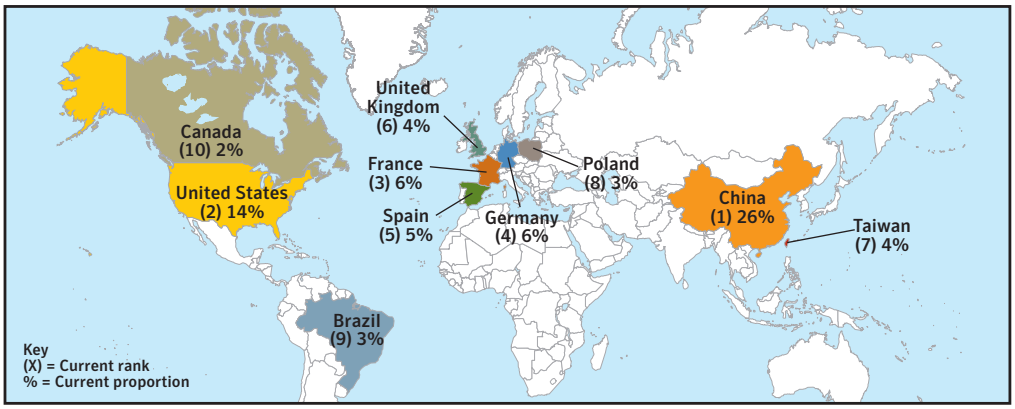

Figure 5. Bot-infected computers by country
Source: Symantec Corporation

The United States was the site of 40 percent of all known command-and-control servers, making it the highest ranked country in this category. The high proportion of command-and-control servers likely indicates that servers in the United States control not only bot networks within the country but offshore as well.

Organizations should monitor all network-connected computers for signs of bot infection, ensuring that any infections are detected and isolated as soon as possible. They should also ensure that all antivirus definitions are updated regularly. As compromised computers can be a threat to other systems, Symantec also recommends that the enterprises notify their ISPs of any potentially malicious activity. Creating and enforcing policies that identify and limit applications that can access the network may also be helpful in limiting the spread of bot infections.

To prevent bot infections, Symantec recommends that ISPs perform both ingress and egress filtering to block known bot traffic.[11] ISPs should also filter out potentially malicious email attachments to reduce exposure to enterprises and end users.

End users should employ defense-in-depth strategies, including the deployment of antivirus software and a firewall.[12] They should update antivirus definitions regularly and ensure that all desktop, laptop, and server computers are updated with all necessary security patches from their operating system vendor. Symantec also advises that users never view, open, or execute any email attachments unless the attachment is expected and comes from a known and trusted source, and unless the purpose of the attachment is known.

[11] Ingress traffic refers to traffic that is coming into a network from the Internet or another network. Egress traffic refers to traffic that is leaving a network, bound for the Internet or another network.
[12] Defense in-depth emphasizes multiple, overlapping, and mutually supportive defensive systems to guard against single-point failures in any specific technology or protection methodology. Defense in-depth should include the deployment of antivirus, firewalls, and intrusion detection systems, among other security measures.

Future Watch

This section of the *Internet Security Threat Report* will discuss emerging trends and issues that Symantec believes will become prominent over the next six to twenty-four months. These forecasts are based on emerging research that Symantec has collected during the current reporting period and are speculative in nature. In discussing potential future trends, Symantec hopes to provide organizations and end users with an opportunity to prepare themselves for rapidly evolving and complex security issues. This section will discuss potential security issues associated with the following:

- Windows Vista™
- Windows Vista and third-party software
- New phishing targets and methods
- Spam and phishing targeting mobile devices
- Virtualization

Threats posed to Windows Vista becoming evident

Microsoft's latest operating system, Windows Vista, was released publicly in January 2007. The release of an operating system that is expected to be widely adopted will likely have a significant effect on the security landscape. The previous *Internet Security Threat Report* discussed some of the general security concerns that may be associated with Windows Vista.[13] Over the past six months, Symantec has continued to research potential issues associated with the new Microsoft operating system, which this section will discuss. These issues fall into three categories: vulnerabilities, malicious code, and attacks against the Teredo protocol.

In December 2006, Symantec reported a vulnerability in previous versions of Windows that also affects the version of Windows Vista that was released to consumers in January.[14] This indicates that Microsoft's Security Development Lifecycle,[15] while thorough, does not necessarily identify all potential vulnerabilities. This may be because some vulnerabilities can be extremely subtle.

That said, it appears that Microsoft's implementation of mitigating technologies such as address space layout randomization (ASLR), GS,[16] and data execution prevention (DEP) could reduce the successful exploitation of any vulnerabilities that are discovered. Nevertheless, Symantec expects that new threats for Windows Vista will utilize older exploitation techniques that have been previously successful—such as those developed to successfully exploit Windows XP SP2—in order to bypass improvements in Windows Vista. For example, attackers may revert to attacks that utilize email, P2P, and other social engineering techniques.

Existing malicious code may also pose a problem for Windows Vista. According to research conducted by Symantec, some malicious code that did not originally target Windows Vista may affect the new operating system. This could be problematic because some enterprises may act on the belief that their installations of Windows Vista are immune from older malicious code samples. As a result, they may not deploy appropriate security solutions on new Windows Vista hosts, potentially leaving them vulnerable to infection by older malicious code samples. For instance, Symantec has already noted that some malicious code samples can infect Windows Vista.[17]

[13] Symantec *Internet Security Threat Report*, Volume X (September 2006): http://www.symantec.com/specprog/threatreport/ent-whitepaper_symantec_internet_security_threat_report_x_09_2006.en-us.pdf : p. 28
[14] http://www.symantec.com/enterprise/security_response/weblog/2006/12/vista_vulnerable.html
[15] The Secure Development Lifecycle is a development paradigm that incorporates security at every stage from the initial architecture to programming and in the quality assurance/testing phases. Threat modeling is a security auditing methodology that involves formally identifying and mapping out all possible attack vectors for an application. See the following for more information: http://www.microsoft.com/presspass/features/2005/nov05/11-21SecurityDevelopmentLifecycle.mspx
[16] GS is a compiler technology. The name is derived from the compiler parameter that is used to enable this functionality. The use of GS will enable stack cookies to be placed around vulnerable functions in order to mitigate stack-based overflows.
[17] For example, please see: http://www.symantec.com/enterprise/security_response/weblog/2006/12/hit_or_miss_vista_and_current.html

The third potential Windows Vista security issue identified by Symantec for this discussion is Teredo. Teredo is a protocol developed by Microsoft to enable the transition between versions of Internet protocol (IP), one of the protocols underlying all Internet-based communications. Teredo is enabled by default in Windows Vista. Computers using Windows Vista can easily be identified through Teredo.

Attacks sent over Teredo will often bypass organizations' network security controls since the protocol is tunneled through network address translation (NAT) over an IPv4 UDP connection. Many security products don't support Teredo and thus would not inspect it. This could make Windows Vista susceptible to attacks through Teredo.[18]

Symantec recommends that enterprises planning a migration to Windows Vista do so first in small, non-critical environments, and that thorough security audits be conducted to reduce possible exposure to attack. In addition, enterprises should ensure that any third-party security solutions they currently use will run on Windows Vista and are deployed in accordance with any existing security policies. Organizations contemplating using IPv6 within Windows Vista rather than Teredo should plan the IPv6 transition carefully, including native access and updated security controls.

Windows Vista release makes third-party software security paramount

With the advent of Windows Vista and the continued use of the Security Development Lifecycle, it is likely that Microsoft-authored code will become more difficult to exploit. As a result, attackers may turn their focus to common third-party applications that are authored by companies that have not employed the Security Development Lifecycle. These third-party applications may not use accepted best software-development practices, such as secure design, secure coding practices, code reviews, or secure developer tools such as Microsoft's Visual Studio.[19] As a result, they may be less secure than Microsoft applications or the Windows Vista platform on which they are deployed.

These third-party applications could include third-party security software (such as antivirus), Web browsers, instant message clients, email clients, and office suites. They may include applications that have a significant user base, either globally or regionally. Symantec has already observed the emergence of a number of zero-day vulnerabilities being exploited in targeted attacks against office suites that are deployed in particular regions.[20]

Due to the security improvements in Windows Vista, third-party drivers may be targeted as a means of gaining kernel-level access on compromised hosts. This is because these applications may not have been developed using the Security Development Lifecycle or other secure development practices. As a result, they may be susceptible to compromise. This could allow attackers to bypass the security improvements in Windows Vista, which are designed to prevent complete compromises, by running applications with non-administrative user privileges.

Only by implementing secure development practices can developers ensure the optimal security of their applications. Failure to employ all available secure coding measures will likely increase the probability of the discovery and successful exploitation of vulnerabilities.

[18] For a more in-depth discussion on the security consequences of Teredo, please visit: http://www.symantec.com/avcenter/reference/Teredo_Security.pdf
[19] Microsoft Visual Studio is important as it introduces a number of security features that can be enabled for unmanaged code. These features include enabling key security features for the application when executed under Windows Vista.
[20] A zero-day attack is one that attacks a vulnerability for which there is no available patch. It also generally means an attack against a vulnerability that is not yet publicly known or known of by the vendor of the affected technology. For example, Justsystem's Ichitaro zero-day was used to transmit a Trojan: http://www.symantec.com/enterprise/security_response/weblog/2006/08/justsystems_ichitaro_0day_used.html

New phishing economies

As phishing becomes entrenched as a mainstream attack activity, antiphishing techniques are improving and phishers are being forced to focus on new targets and adopt new methods. Symantec believes that, in the near future, phishers will expand the scope of their targets to include new industry sectors. For example, they will likely start to target a number of the secondary economies introduced through so-called massively multiplayer online games (MMOGs).[21] MMOGs have become big business and are already attracting large groups of organized criminals who are using digital attacks for financial gain. In December 2006, forty-four suspects were arrested for stealing $90,000 USD worth of digital assets from a single game.[22]

Symantec also expects that phishers will develop new techniques to evade antiphishing solutions. Symantec has already started to see techniques to counter the effectiveness of block lists. For instance, phishers can use multiple unique URLs to direct users to a single Web site. Each URL is discarded after one use, so that even if they are placed on a block list, the lists still will not be able to block other URLs that direct potential victims to the same Web site. In some cases, Symantec has observed thousands of distinct URLs directing users to a single Web site.[23] Finally, attackers may already be using ready-made phishing kits. A phishing kit is a set of tools that an attacker can use to easily construct phishing email messages and Web sites based on a template.

Symantec has also observed that phishers are starting to adopt a technique known as intelligence lead phishing. This is a practice in which the phisher compromises a database or social networking site to obtain user information. This information is then used in a targeted phishing attempt against the user in question. The high degree of personalization made possible by the illicitly gained information can increase the effectiveness of the phishing attempt significantly. As widespread phishing attempts are increasingly choked off by antiphishing technologies, Symantec expects to see more phishing attacks that use these intelligence techniques.

In addition to the evolved phishing techniques outlined above, Symantec expects to see more generic phishing attacks; that is, attacks that are not restricted to spoofing a particular brand. For instance, instead of being required to know which bank the targeted user currently uses, a generic phishing attack could instead prompt the victim to “switch to Bank XYZ.” These more generic phishing attempts can be restricted to a particular country if the phished institution is nationally based, thereby increasing the phisher’s chance of success.

These recently evolved techniques illustrate the need for enterprises and end users to deploy effective antiphishing and antifraud solutions. Enterprises should be aware of and implement effective antiphishing technologies and practices. Enterprises that engage their clients over the Internet should continue to stay abreast of new phishing methods and techniques.[24] They should also monitor abuses of their brand in order to react appropriately and minimize potential damage to the company’s reputation.

End users should follow best security practices, including the use of regularly updated antivirus software, antispam software, firewalls, toolbar blockers, and other software detection methods. Symantec also advises end users to never disclose any confidential personal or financial information unless and until they can confirm that any request for such information is legitimate.

[21] A massively multiplayer online game is an Internet-based computer game on which hundreds to thousands of players are capable of participating simultaneously.
[22] Please see “Virtual Item Theft Ring Busted” http://playnoevil.com/serendipity/index.php?/archives/1051-Virtual-Item-Theft-Ring-Busted.html
[23] http://www.symantec.com/enterprise/security_response/weblog/2006/12/phishing_2006_the_year_in_revi.html
[24] See the Symantec Phish Report Network, an extensive antifraud community where members contribute and receive fraudulent Web site addresses for alerting and blocking attacks across a broad range of solutions. It is available at: http://www.phishreport.net

Enterprises that use the Internet for any transaction-based activity should ensure that they have implemented phishing detection and response processes and procedures. In addition to providing a structured, standardized response to a phishing incident, this will also ensure that information is passed on to the appropriate resources, thereby protecting against subsequent use of the same attack.

Enterprises should ensure that their users are educated about phishing techniques and are informed of the latest phishing scams. For further information, the Internet Fraud Complaint Center (IFCC) has released a set of guidelines on how to avoid Internet-related scams.[25]

SMiShing—Spam and phishing go mobile

In July 2006, Symantec reported that SMS and MMS had emerged as new vectors for spam and phishing activity.[26] Subsequently, the term SMiShing was coined by the industry to describe this class of threat.

There is a logical evolution from email to SMS and MMS as transport mechanisms for spam and phishing attacks. This is due in part to the fact that the technological and procedural defenses for devices deploying these services may not be as well developed or as widely deployed as those for other platforms. Additionally, users of mobile devices typically perceive messages received by SMS and MMS as being more personal than those received by email on a desktop computer. Furthermore, threats against these surfaces have been rare thus far. As a result, users are more likely to trust those messages and to act on them.

Targeting SMS and MMS may also offer attackers a significant benefit over targeting a specific mobile operating system. SMS and MMS are sufficiently well established and are deployed widely enough that they are available on almost all handsets on all networks. Most legacy and proprietary operating system handsets will support both of these technologies. As a result, they have a much larger target user base than smartphones.

There has been a rise in the amount of SMS-based premium-rate spam over the past few years since the introduction of subscriber-billed SMS.[27] This is a payment model in which the subscriber is billed a considerably higher cost for receiving a message than for sending one. This mechanism is typically used lawfully by the suppliers of ring tones, wallpapers, and other mobile content such as games. It is a convenient way of making micro-payments without having to introduce another payment tool such as a debit or credit card. However, some criminals have utilized the technology to obtain money, which has resulted in a number of national telecommunications regulators stamping out the practice.[28]

Symantec speculates that SMS- and MMS-based phishing and spam will continue to increase. Cellular operators will likely be forced to invest in filtering technologies to combat this growing problem. This issue will be compounded by the fact that there are a number of different Internet-based SMS gateways that could allow users to supply their own originating number or name, which could be spoofed and used to send spam. As the costs of SMS services goes down, the likelihood that these gateways will be used for spam activities will increase.

[25] http://www.fbi.gov/majcases/fraud/internetschemes.htm
[26] SMS (short messaging service) is a service that is used for sending short text messages to mobile phones and other mobile text devices such as pagers. MMS (multimedia messaging service) is a service that allows mobile devices to send phone messages as well as multimedia files, such as images, audio, and video.
[27] http://www.grumbletext.co.uk
[28] http://news.bbc.co.uk/1/hi/technology/4708167.stm

Software virtualization brings new security threats

Software virtualization is a technology that allows one computer (the host) to run one or more distinct virtual computers (the guests). These virtual computers each run independently of the others and have their own virtual hardware, allowing the user to run multiple different operating systems on the same physical hardware.

Software virtualization has become a very powerful tool, bringing with it numerous benefits. However, many users assume that virtual machines provide a foolproof security barrier, leading to a false sense of security. While it is true that virtual machines can insulate against some current attacks, there are others against which they offer no protection. Further, they could potentially make new classes of attack possible. Symantec believes that the potential security implications of software virtualization have not yet been fully investigated and understood.

Guest virtual machines may not run the same security software as the host. For instance, they may not include antivirus software, personal firewalls, or host-based intrusion prevention products. As a result of these omissions, the virtual machines may be more exposed to threats than if they were run on independent hardware. Furthermore, virtual machines will do little to protect the data on the host. Consequently, virtualization technology may not diminish or protect against the threat of application-oriented threats such as phishing and data theft.

Symantec also believes that threats that are specific to virtualization technologies could emerge. With many different virtual machines being used, Symantec believes that these virtualization-specific threats could fall into two distinct classes of threat.

The first type of threat targets the use of real hardware in virtualized machines. Hardware drivers that provide software emulation of hardware acceleration outside of the virtual machine in the host operating system could be targeted from inside the guest operating system. An example of a vulnerability that illustrated this principle was the NVIDIA Binary Graphics Driver for Linux Buffer Overflow Vulnerability.[29] Symantec speculates that this type of vulnerability could be exploited from within the guest operating system to break into the host system. For enterprises that rely on separation through the use of software virtualization technology, the impact of this type of threat could be considerable.

The second type of threat that Symantec believes could emerge is related to the impact that software-virtualized computers may have on random number generators that are used inside guest operating systems on virtual machines. This speculation is based on some initial work done by Symantec Advanced Threat Research in a paper on GS and ASLR in Windows Vista. This research showed that the method used to generate the random locations employed in some security technologies would, under certain circumstances, differ wildly in a software-virtualized instance of the operating system. If this proves to be true, it could have considerable implications for a number of different technologies that rely on good randomness, such as unique identifiers, as well as the seeds used in encryption.

In the short to medium term, enterprises need to fully understand any potential impact that the use of software virtualization may have on the security of their environment and plan accordingly. They should control and monitor host operating systems very strictly, as the expected activity would likely be limited to the starting and stopping of virtual machines. Symantec feels that these threats constitute an important area of research and will continue to investigate and monitor these issues.

[29] http://www.securityfocus.com/bid/20559

Attack Trends

This section of the *Internet Security Threat Report* will provide an analysis of attack activity, identity theft-related data breaches, and the trade of illegal information that Symantec observed between July 1 and December 31, 2006. Attacks are defined as any malicious activity carried out over a network that has been detected by an intrusion detection system (IDS) or firewall.

Attack activity is monitored by the Symantec™ Global Intelligence Network across the entire Internet. Over 40,000 sensors deployed in more than 180 countries by Symantec DeepSight™ Threat Management System and Symantec™ Managed Security Services gather this data.

Furthermore, Symantec uses proprietary technologies to monitor bot command-and-control servers and underground economy servers across the Internet. Finally, Symantec uses publicly available information to assess identity theft-related data breaches.[30] These resources combine to give Symantec an unparalleled ability to identify, investigate, and respond to emerging threats. This discussion will be based on data provided by all of these sources.

Attack Trends Highlights

The following section will offer a brief summary of some of the attack trends that Symantec observed during this period based on data provided by the sources listed above. Following this overview, the *Internet Security Threat Report* will discuss selected metrics in greater depth, providing analysis and discussion of the trends indicated by the data.

- The government sector accounted for 25 percent of all identity theft-related data breaches, more than any other sector.
- The theft or loss of a computer or other data-storage medium made up 54 percent of all identity theft-related data breaches during this period.
- The United States was the top country of attack origin, accounting for 33 percent of worldwide attack activity.
- Symantec recorded an average of 5,213 denial of service (DoS) attacks per day, down from 6,110 in the first half of the year.
- The United States was the target of most DoS attacks, accounting for 52 percent of the worldwide total.
- The government sector was the sector most frequently targeted by DoS attacks, accounting for 30 percent of all detected attacks.
- Microsoft Internet Explorer was targeted by 77 percent of all attacks specifically targeting Web browsers.
- Home users were the most highly targeted sector, accounting for 93 percent of all targeted attacks.
- Symantec observed an average of 63,912 active bot-infected computers per day, an 11 percent increase from the previous period.

[30] Data is made available by Privacy Rights Clearinghouse, a non-profit consumer information and advocacy organization: http://www.privacyrights.org

- China had 26 percent of the world's bot-infected computers, more than any other country.
- The United States had the highest number of bot command-and-control computers, accounting for 40 percent of the worldwide total.
- Beijing was the city with the most bot-infected computers in the world, accounting for just over five percent of the worldwide total.
- The United States accounted for 31 percent of all malicious activity during this period, more than any other country.
- Israel was the most highly ranked country for malicious activity per Internet user followed by Taiwan and Poland.
- Fifty-one percent of all underground economy servers known to Symantec were located in the United States, the highest total of any country.
- Eighty-six percent of the credit and debit cards advertised for sale on underground economy servers known to Symantec were issued by banks in the United States.

Attack Trends Discussion

This section will discuss selected "Attack Trends" metrics in greater depth, providing analysis and discussion of the trends indicated by the data. The following metrics will be discussed:

- Malicious activity by country
- Malicious activity by country per Internet user
- Data breaches that could lead to identity theft
- Underground economy servers
- Bot-infected computers
- Bot-infected computers by country

Malicious activity by country

In this volume of the *Internet Security Threat Report*, Symantec is evaluating the countries in which the highest amount of malicious activity takes place or originates. To determine this, Symantec has compiled geographical data on numerous malicious activities, namely: bot-infected computers, bot command-and-control servers, phishing Web sites, malicious code reports, spam relay hosts, and Internet attacks.

To determine the proportion of Internet-wide malicious activity that originated in each country, the mean of the proportion of all of the considered malicious activities that originated in each country was calculated. This average determined the proportion of overall malicious activity that originated from the country in question and was used to rank each country. This section will discuss those findings.

Between July 1 and December 31, 2006 the United States was the top country for malicious activity, making up 31 percent of worldwide malicious activity (table 2). For each of the malicious activities taken into account for this measurement, the United States ranked number one by a large margin with the exception of bot-infected computers. It ranked second for that criteria, 12 percentage points lower than China.

Overall Rank	Country	Overall Proportion	Malicious Code Rank	Spam Host Rank	Command and Control Server Rank	Phishing Host Rank	Bot Rank	Attack Rank
1	United States	31%	1	1	1	1	2	1
2	China	10%	3	2	4	8	1	2
3	Germany	7%	7	3	3	2	4	3
4	France	4%	9	4	14	4	3	4
5	United Kingdom	4%	4	13	9	3	6	6
6	South Korea	4%	12	9	2	9	11	9
7	Canada	3%	5	23	5	7	10	5
8	Spain	3%	13	5	15	16	5	7
9	Taiwan	3%	8	11	6	6	7	11
10	Italy	3%	2	8	10	14	12	10

Table 2. Malicious activity by country
Source: Symantec Corporation

The high degree of malicious activity originating in the United States is likely driven by the expansive Internet infrastructure there. The United States accounts for 19 percent of the world's Internet users.[31] Furthermore, the number of broadband Internet users in that country grew by 14 percent between December 2005 and July 2006.[32] Despite the relatively well developed security infrastructure in the United States, the high number of Internet-connected computers there presents more opportunities for malicious activities to take place. Symantec predicts that the United States will remain the highest ranked country for malicious activity until another country exceeds it in numbers of Internet users and broadband connectivity.

China was the second highest country for malicious activity during this six-month reporting period, accounting for 10 percent. China's prominence, like that of the United States, is likely driven by the high number of Internet users there, as well as the rapid growth in the country's Internet infrastructure. China has the second highest Internet usership in the world, accounting for 11 percent of the world's users and, as has been stated previously in this report, is expected to surpass the United States in usership in the next year.[33]

In the last six months of 2006, Germany was the third ranked country for malicious activity. Seven percent of all Internet-wide malicious activity originated there during this period. Germany ranked highly in all the malicious activities considered for this metric. The prominence of Germany, like that of both China and the United States, is likely influenced by its Internet infrastructure and growth. Germany has the fifth highest Internet usership in the world, boasting five percent of the world's users.[34] Furthermore, the number of broadband users increased by 16 percent between December 2005 and July 2006.[35]

[31] http://www.internetworldstats.com
[32] http://www.oecd.org/document/9/0,2340,en_2649_34225_37529673_1_1_1_1,00.html
[33] http://www.internetworldstats.com
[34] http://www.internetworldstats.com
[35] http://www.oecd.org/document/9/0,2340,en_2649_34225_37529673_1_1_1_1,00.html

Malicious activity by country per Internet user

Having assessed the top countries by malicious activity, Symantec also evaluated the top 25 of these countries according to the number of Internet users located there. This measure is intended to remove the bias of high Internet users from the consideration of the "Malicious activity by country" metric.

In order to determine this, Symantec divided the amount of malicious activity originating in each of the top 25 countries by the number of worldwide Internet users who are located in that country. The proportion assigned to each country in this discussion thus equates to the proportion of malicious activity that could be attributed to a single (average) Internet user in that country.

Israel was the most highly ranked country for malicious activity per Internet user. If one person from each of the top 25 countries were assessed as a representation of their country's Internet users, the average user in Israel would carry out nine percent of the group's malicious activity (figure 6).

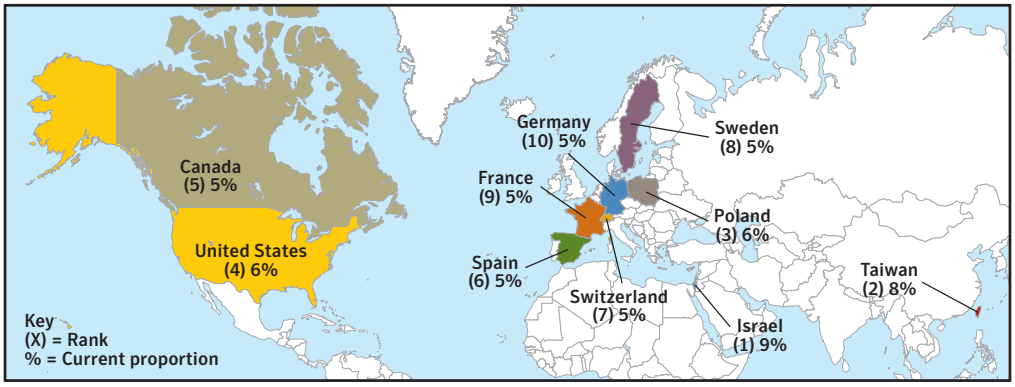

Figure 6. Malicious activity by country per Internet user
Source: Symantec Corporation

Taiwan ranked second, accounting for eight percent of malicious activity per Internet user. Poland ranked third, accounting for six percent. Although these countries both have a high proportion of malicious activity per Internet user, they account for a relatively low proportion of malicious activity worldwide.

Data breaches that could lead to identity theft

For the first time, in this volume of the *Internet Security Threat Report*, Symantec is assessing data breaches that have exposed information that could lead to identity theft. Using publicly available data,[36] Symantec has determined the sectors that were most often affected by these data breaches, as well as the most common causes of data loss.

Identity theft is an increasingly prevalent security issue. Many organizations manage information that could facilitate identity theft. Compromises that result in the loss of personal data could be quite costly, not only to the people whose identity may be at risk and to their respective financial institutions, but also to the organization. Data leaks that lead to identity theft could damage the organization's reputation, thereby potentially undermining customer confidence. They could also result in criminal charges and/or litigation.

[36] http://www.privacyrights.org/ar/ChronDataBreaches.htm

In the second half of 2006, the government sector accounted for the majority of data breaches that could lead to identity theft, making up 25 percent of the total (figure 7). Government organizations store a lot of personal information that could be used for the purposes of identity theft. These organizations often consist of numerous semi-independent departments. As a consequence, sensitive personally identifiable information may be stored in separate locations and be accessible by numerous people. This increases the opportunities for attackers to gain unauthorized access to this data.

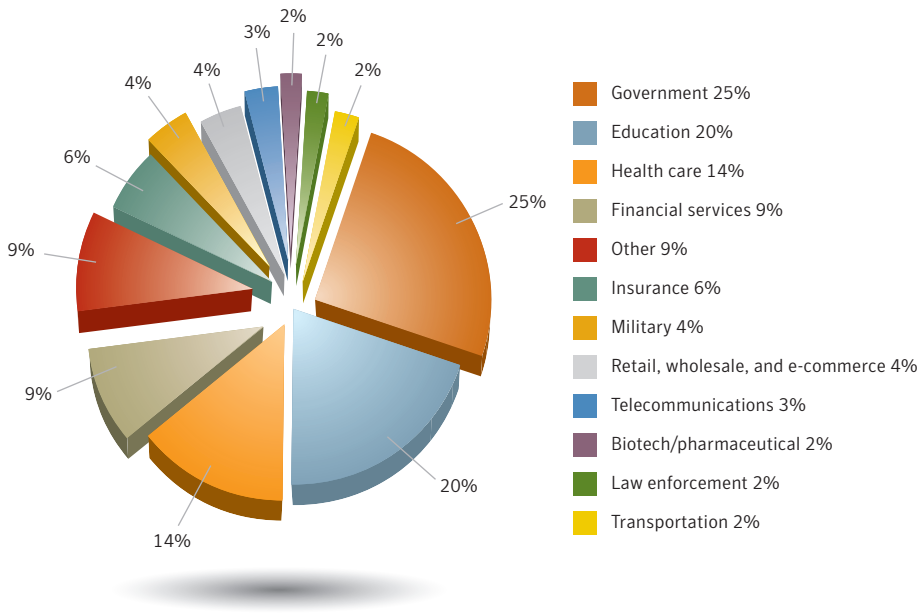

Figure 7. Data breaches that could lead to identity theft by sector
Source: Based on data provided by Privacy Rights Clearinghouse and Attrition.org

The second factor relates to the reporting of such breaches. Government organizations are more likely to report data breaches, either due to legislative obligation,[37] or due to publicly accessible audits and performance reports. As well, companies that rely on consumer confidence may be less inclined to report such breaches for fear of negative market reaction.

During this reporting period, the education sector accounted for 20 percent of data breaches that could lead to identity theft. Health care accounted for 14 percent of the total, the third highest number. Organizations in both of these sectors store and manage a significant amount of sensitive personal information that can be used for the purposes of identity theft. Furthermore, health organizations store information related to personal health that could result in damaging breaches of privacy if viewed by unauthorized personnel. Educational organizations such as research hospitals may also store such information.

[37] An example is the Fair and Accurate Credit Transactions Act of 2003 (FACTA) of California. For more on this act, please see: http://www.privacyrights.org/fs/fs6a-facta.htm

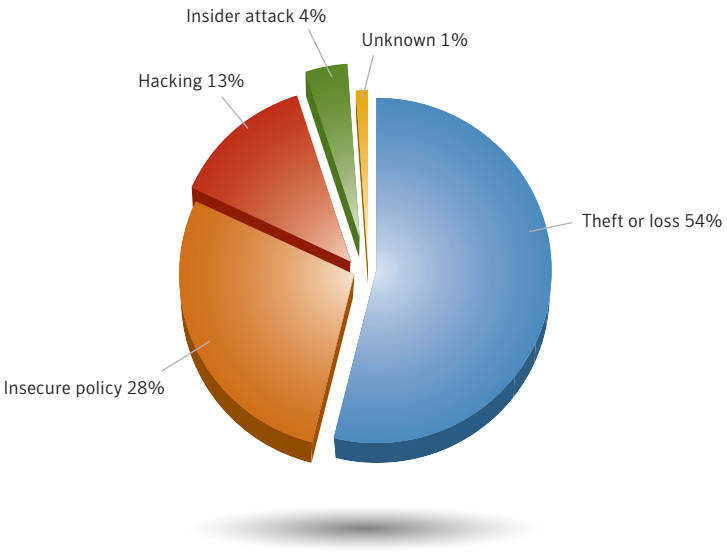

Figure 8. Data breaches that could lead to identity theft by cause
Source: Based on data provided by Privacy Rights Clearinghouse and Attrition.org

In the second half of 2006, the primary cause of data breaches that could facilitate identity theft was theft or loss of a computer or other medium on which the data is stored or transmitted, such as a USB key or back-up medium (figure 8). These made up 54 percent of all identity theft-related data breaches during this period. In many cases, computers that were lost or stolen were laptop computers.

The second most common cause of data breaches that could lead to identity theft during this period was insecure policy, which made up 28 percent of all incidents. A data breach is considered to be caused by insecure policy if it can be attributed to a failure to develop, implement, and/or comply with adequate security policy. For example, this could mean posting personally identifiable information on a publicly available Web site, sending it through unencrypted email, or storing it in unencrypted form.

Together, theft and loss along with insecure policy made up 82 percent of all data breaches in the second half of 2006. Most breaches of this type are avoidable. In the case of theft or loss, the compromise of data could be averted by encrypting all sensitive data. This would ensure that even if the data were stolen, it would not be accessible to unauthorized third parties. This step should be part of a broader security policy that organizations should develop, implement, and enforce in order to ensure that all sensitive data is protected from unauthorized access.

Underground economy servers

Underground economy servers are used by criminals and criminal organizations to sell stolen information, typically for subsequent use in identity theft. This data can include government-issued identification numbers, credit cards, bank cards and personal identification numbers (PINs), user accounts, and email address lists. Symantec tracks and assesses underground economy servers across the Internet using proprietary online fraud monitoring tools.

For the first time, in this volume of the *Internet Security Threat Report*, Symantec will assess underground economy servers. It will do so in two ways: according to their geographic location and according to the location of banks that issued credit and debit cards that were being sold on underground economy servers. This discussion will also look at the types of information that are being exchanged through underground economy servers.

During the last six months of 2006, 51 percent of all known underground economy servers in the world were located in the United States, the highest total of any country (figure 9). The prominence of the United States is no surprise, as the expansive Internet infrastructure and continual broadband growth there create numerous opportunities for criminals to carry out malicious activities. Sweden ranked second, accounting for 15 percent, and Canada ranked third, accounting for seven percent of all underground economy servers.

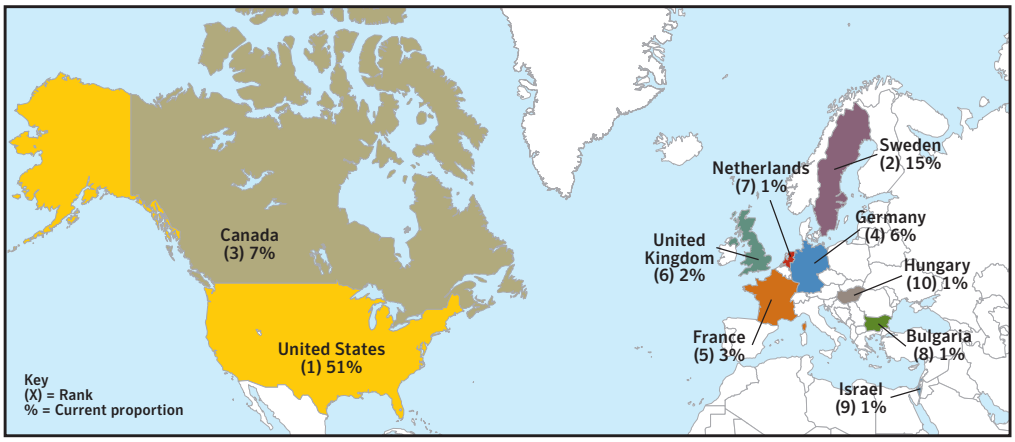

Figure 9. Location of underground economy servers
Source: Symantec Corporation

During the last six months of 2006, Symantec observed 4,943 credit cards being traded on underground economy servers.[38] Symantec also determined that, by far, most of the credit and debit cards advertised for sale on underground economy servers were issued by banks in the United States (figure 10). The prominence of the United States is not entirely unexpected. As was discussed earlier in this report, the vast majority of the identity theft-related data breaches reported during the last six months of 2006 took place in the United States.

[38] It should be noted that this discussion is not necessarily representative of Internet-wide activity; rather, it is intended as a snapshot of the limited activity that Symantec monitored during this period.

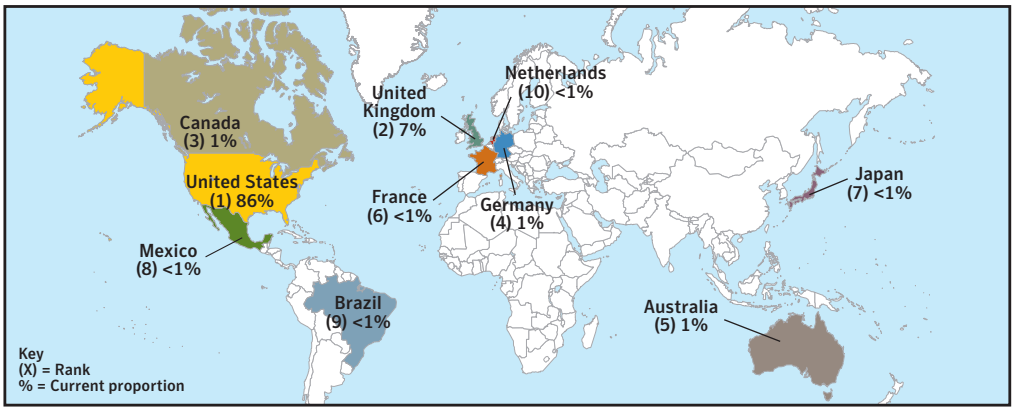

Figure 10. Location of banks whose cards were sold on underground economy servers
Source: Symantec Corporation

Cards from the United States are generally advertised for about half as much as those from the United Kingdom. For instance, credit cards from United States-based banks have been advertised for $3.00 USD, while credit cards from United Kingdom-based banks are advertised for $6.00 USD. This could be because there are a much higher number of cards from the United States available for sale. It could also be because the UK pound is currently stronger than the US dollar. Finally, it could indicate that buyers in the United Kingdom, and other countries, are unlikely to want to purchase cards from the United States.

In addition to bank and credit cards, Symantec has discovered other items that are being sold on underground economy servers (table 3). These include full identities, which typically involve government-issued identification numbers (such as social security numbers), bank account information (including passwords), personal information (such as date of birth), as well as identity verification information (such as a person's mother's maiden name).

Item	Advertised Price (in US Dollars)
United States-based credit card with card verification value	$1–$6
United Kingdom-based credit card with card verification value	$2–$12
An identity (including US bank account, credit card, date of birth, and government issued identification number)	$14–$18
List of 29,000 emails	$5
Online banking account with a $9,900 balance	$300
Yahoo Mail cookie exploit—advertised to facilitate full access when successful	$3
Valid Yahoo and Hotmail email cookies	$3
Compromised computer	$6–$20
Phishing Web site hosting—per site	$3–5
Verified PayPal account with balance (balance varies)	$50–$500
Unverified PayPal account with balance (balance varies)	$10–$50
Skype account	$12
World of Warcraft account—one month duration	$10

Table 3. Advertised prices of items traded on underground economy servers
Source: Symantec Corporation

Advertised prices for identities range from $14.00 USD to $18.00 USD. Other items that can be purchased on underground economy servers include lists of email addresses, stolen gift certificates, compromised computers, one-month "World of Warcraft" accounts, as well as a number of email and Web-based accounts, which include usernames and passwords.

In order to reduce the likelihood of identity theft, organizations that store personal information should take the necessary steps to protect data transmitted over the Internet or stored on their computers. This should include the development, implementation, and enforcement of a secure policy requiring that all sensitive data is encrypted. This would ensure that even if the computer or medium on which the data was stored were lost or stolen, the data would not be accessible. This step should be part of a broader security policy that organizations should develop and implement in order to ensure that any sensitive data is protected from unauthorized access.

Bot-infected computers

Bots are programs that are covertly installed on a user's machine in order to allow an unauthorized user to control the computer remotely. They allow an attacker to remotely control the targeted system through a communication channel such as IRC. These channels allow the remote attacker to control a large number of compromised computers over a single, reliable channel in a bot network, which can then be used to launch coordinated attacks.

Bots allow for a wide range of functionality and most can be updated to assume new functionality by downloading new code and features. Bots can be used by external attackers to perform DoS attacks against an organization's Web site. Furthermore, bots within an organization's network can be used to attack other organizations' Web sites, which can have serious business and legal consequences. Bots can be used by attackers to harvest confidential information from compromised computers, which can lead to identity theft. Bots can also be used to distribute spam and phishing attacks, as well as spyware, adware, and misleading applications.

Between July 1 and December 31, 2006, Symantec observed an average of 63,912 active bot-infected computers per day (figure 11). This is an 11 percent increase from the previous period when Symantec observed an average of 57,717 active bots per day. Furthermore, Symantec observed 6,049,594 distinct bot-infected computers during the current reporting period, a 29 percent increase from the previous period when 4,696,903 distinct bot-instinct computers were identified.

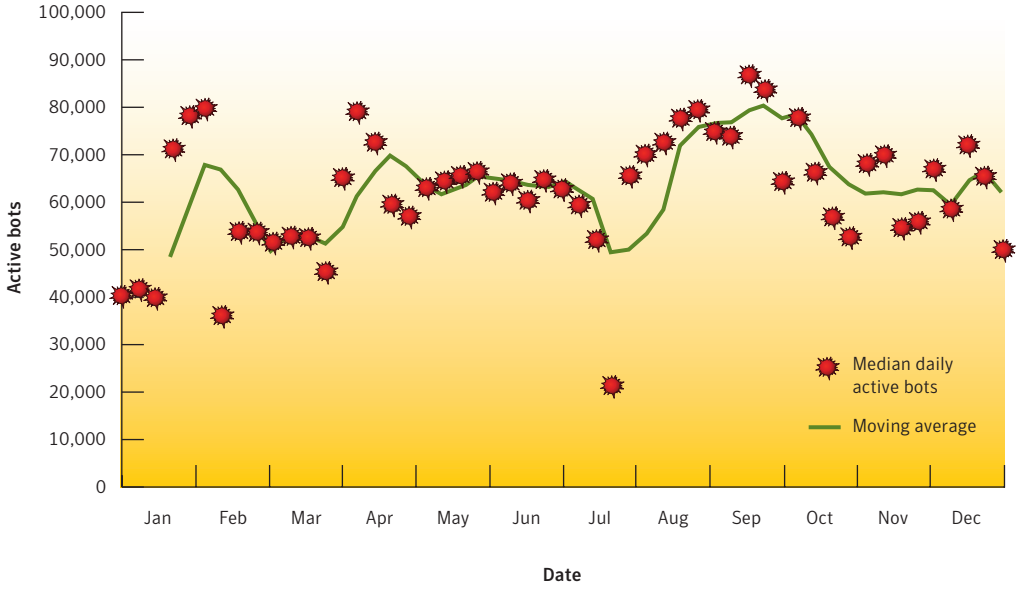

Figure 11. Active bot-infected computers per day
Source: Symantec Corporation

This increase is largely driven by a peak in bot activity in September. During this month, a number of vulnerabilities that had previously been disclosed were actively exploited by bots, including the Microsoft Windows Server Service Remote Buffer Overflow Vulnerability as well as the Microsoft Internet Explorer Vector Markup Language Buffer Overflow Vulnerability.[39]

[39] Please see http://www.securityfocus.com/bid/19409 and http://www.securityfocus.com/bid/20096, respectively.

In Volume IX (March 2006) of the *Internet Security Threat Report*, Symantec speculated that bot networks had likely reached their saturation point.[40] The slight increase in bot computers in the second half of 2006, along with the lack of an increase in attacks, likely indicates that bots are slowly making up a greater proportion of attacking computers. That is, it is likely that attackers who use means other than bot-infected computers in coordinated bot networks are becoming less common.

It may also be possible that the increase in bot-infected computers is a sign of an impending boom cycle in bots. In the same discussion cited in the previous paragraph, Symantec speculated that if bots begin to exploit an attack vector that bypasses firewalls and perimeter defenses, the population of bot-infected computers could rapidly increase.[41] It is possible that the increased focus on Web browsers may be facilitating this. As a result, Symantec believes that it is reasonable to speculate that a boom period for bots is possible in the near future.

Symantec also tracks the number of bot command-and-control servers worldwide. Command-and-control servers are computers that bot network owners use to relay commands to bot-infected computers on their networks. In the last six months of 2006, Symantec identified 4,746 bot command-and-control servers. This is a 25 percent decrease from the 6,337 detected during the first six months of 2006.

A drop in the number of command-and-control servers combined with a rise in the number of bot-infected computers indicates that, on average, bot networks are increasing in size. Bot networks are thus becoming more consolidated. Consolidated bot networks will likely mean that organizations will have to deal with a well entrenched, experienced, and dedicated group of bot network owners instead of a population of hobby hackers.

It could also signal a fundamental change in the way bots communicate with one another. Symantec has seen bots that are structured on a peer-to-peer model, in which the machines connect together rather than connecting to a central command-and-control server. Symantec has also observed that command-and-control servers are beginning to adopt encryption so that they are less visible.

To prevent bot infections, Symantec recommends that ISPs perform both ingress and egress filtering to block known bot traffic.[42] ISPs should also filter out potentially malicious email attachments to reduce exposure to enterprises and end users.

Organizations should monitor all network-connected computers for signs of bot infection, ensuring that any infections are detected and isolated as soon as possible. They should also ensure that all antivirus definitions are updated regularly. As compromised computers can be a threat to other systems, Symantec also recommends that enterprises notify their ISPs of any potentially malicious activity.

Organizations should also perform egress filtering on outgoing network traffic, ensuring that malicious activity and unauthorized communications are not taking place. They should also create and enforce policies that identify and restrict applications that can access the network.

[40] Symantec *Internet Security Threat Report*, Volume IX (March 2006): http://eval.veritas.com/mktginfo/enterprise/white_papers/ent-whitepaper_symantec_internet_security_threat_report_ix.pdf : p.36
[41] Symantec *Internet Security Threat Report*, Volume IX (March 2006): http://eval.veritas.com/mktginfo/enterprise/white_papers/ent-whitepaper_symantec_internet_security_threat_report_ix.pdf : p.36
[42] Ingress traffic refers to traffic that is coming into a network from the Internet or another network. Egress traffic refers to traffic that is leaving a network, bound for the Internet or another network.

To reduce exposure to bot-related attacks, end users should employ defense-in-depth strategies, including the deployment of antivirus software and a firewall. Users should update antivirus definitions regularly and ensure that all desktop, laptop, and server computers are updated with all necessary security patches from their operating system vendor. Symantec also advises that users never view, open, or execute any email attachment unless the attachment is expected and comes from a known and trusted source, and unless the purpose of the attachment is known.

Bot-infected computers by country

Recognizing the ongoing threat posed by bot networks, Symantec tracks the distribution of bot-infected computers worldwide. This can help analysts understand how bot-infected computers are distributed globally. This is important, as a high percentage of bot-infected computers likely indicates a greater potential for bot-related attacks. It could also give insight into the level of patching and security awareness amongst computer administrators and users in a given region.

China had the highest number of bot-infected computers during the second half of 2006, accounting for 26 percent of the worldwide total (figure 12). This is an increase of six percentage points over the previous six months. This increase was driven by a rise in the number of bots in the country rather than a decrease in other countries. This coincides with and illustrates a trend that Symantec first discussed in 2005, which saw an increase in bot activity in China during that period.[43]

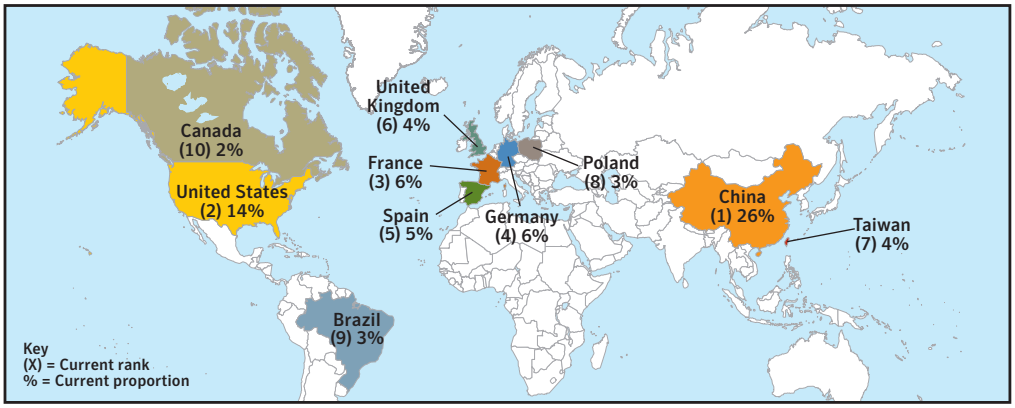

Figure 12. Bot-infected computers by country
Source: Symantec Corporation

[43] Symantec *Internet Security Threat Report*, Volume VII (March 2005): http://eval.veritas.com/mktginfo/enterprise/white_papers/ent-whitepaper_symantec_internet_security_threat_report_vii.pdf : p. 26

Symantec has observed that bots usually infect computers that are connected to high-speed broadband Internet through large ISPs and that the expansion of broadband connectivity often facilitates the spread of bots. China's Internet infrastructure is currently expanding rapidly. The number of broadband subscribers located there is expected to surpass that of the United States in the next year.[44]

Frequently, rapidly expanding ISPs will focus their resources on meeting growing broadband demand at the expense of implementing adequate security measures, such as port blocking and ingress and egress filtering. As a result, they may have security infrastructures and practices that are insufficient for their needs. Furthermore, it is also likely that home users and system administrators in China are also struggling to adapt their security practices and policies to deal with broadband Internet. Symantec believes that bot activity in China will continue to rise as long as broadband Internet in China continues to be adopted at a rapid rate.

Symantec also tracks the global distribution of bot command-and-control servers. These are computers that bot network owners use to relay commands and instructions to the bot-infected computers that make up their networks. This analysis will allow administrators to identify and understand the locations from which bot networks are being controlled as well as the geographic distribution of bot networks.

Although China had the most bot-infected computers worldwide, it had only the fourth highest number of known command-and-control servers worldwide (table 4). This discrepancy likely indicates that the majority of bot-infected computers in China are being controlled from servers in other countries. While it is simple to trace a command-and-control server to its location, the server may not reside in the same location as the person who controls it. For example, an attacker in the United States could control a command-and-control server in the United Kingdom to administer bot-infected computers all over the world.

Current Rank	Previous Rank	Country	Current Proportion	Previous Proportion
1	1	United States	40%	42%
2	2	South Korea	10%	8%
3	5	Germany	6%	5%
4	4	China	5%	6%
5	3	Canada	4%	7%
6	6	Taiwan	3%	3%
7	7	Sweden	3%	3%
8	8	Japan	2%	3%
9	12	United Kingdom	2%	2%
10	10	Italy	2%	2%

Table 4. Command-and-control servers by country
Source: Symantec Corporation

[44] http://www.vnunet.com/vnunet/news/2163552/china-lead-broadband-world

In the last six months of 2006, the United States had the second highest number of bot-infected computers, accounting for 14 percent of the worldwide total. The United States was also the site of 40 percent of all known command-and-control servers, making it the highest ranked country in that category.

The high proportion of command-and-control servers in the United States likely indicates that servers there control not only bot networks within the country but offshore as well. The high proportion of bot-infected computers and command-and-control servers in the United States is driven by that country's extensive Internet and technology infrastructure. As of June 2006, more than 57 million broadband Internet users were located there, the highest number in the world.[45]

France had the third highest proportion of bot-infected computers, accounting for six percent of the worldwide total. The rise of France to the third position is driven primarily by a drop in the percentage of bot-infected computers located in the United Kingdom. This drop likely indicates that the security awareness and infrastructure in the United Kingdom are catching up to the growth of Internet connectivity there.

[45] http://www.oecd.org/document/9/0,2340,en_2825_495656_37529673_1_1_1_1,00.html

Vulnerability Trends

Vulnerabilities are design or implementation errors in information systems that can result in a compromise of the confidentiality, integrity, or availability of information stored upon or transmitted over the affected system. They are most often found in software, although they exist in all layers of information systems, from design or protocol specifications to physical hardware implementations. Vulnerabilities may be triggered actively, either by malicious users or automated malicious code, or passively during system operation. The discovery and disclosure of a single vulnerability in a critical asset can seriously undermine the security posture of an organization.

New vulnerabilities are discovered and disclosed regularly by a sizeable community of end users, security researchers, hackers, security vendors, and occasionally by the software vendors themselves. Symantec carefully monitors vulnerability research, tracking vulnerabilities throughout their lifecycle, from initial disclosure and discussion to the development and release of a patch or other remediation measure.

Symantec operates one of the most popular forums for the disclosure and discussion of vulnerabilities on the Internet, the BugTraq™ mailing list, which has approximately 50,000 direct subscribers who contribute, receive, and discuss vulnerability research on a daily basis.[46] Symantec also maintains one of the world's most comprehensive vulnerability databases, currently consisting of over 20,000 vulnerabilities (spanning more than a decade) affecting more than 45,000 technologies from over 7,000 vendors. The following discussion of vulnerability trends is based on a thorough analysis of that data.

This section of the Symantec *Internet Security Threat Report* will discuss vulnerabilities that have been disclosed between July 1 and December 31, 2006. It will compare them with those disclosed in the previous six-month period and discuss how current vulnerability trends may affect potential future Internet security activity.

Vulnerability Trends Highlights

The following section will offer a brief summary of some of the vulnerability trends that Symantec observed during this period based on data provided by the sources listed above. Following this overview, the *Internet Security Threat Report* will discuss selected metrics in greater depth, providing analysis and discussion of the trends indicated by the data.

- Symantec documented 2,526 vulnerabilities in the second half of 2006, 12 percent higher than the first half of 2006, and a higher volume than in any other previous six-month period.[47]
- Symantec classified four percent of all vulnerabilities disclosed during this period as high severity, 69 percent were medium severity, and 27 percent were low severity.
- Sixty-six percent of vulnerabilities disclosed during this period affected Web applications.
- Seventy-nine percent of all vulnerabilities documented in this reporting period were considered to be easily exploitable.
- Seventy-seven percent of all easily exploitable vulnerabilities affected Web applications, and seven percent affected servers.

[46] The BugTraq mailing list is hosted by SecurityFocus (http://www.securityfocus.com). Archives are available at http://www.securityfocus.com/archive/1
[47] The Symantec *Internet Security Threat Report* has been tracking vulnerabilities in six-month periods since January 2002.

- Ninety-four percent of all easily exploitable vulnerabilities disclosed in the second half of 2006 were remotely exploitable.
- In the second half of 2006, all the operating system vendors that were studied had longer average patch development times than in the first half of the year.
- Sun Solaris had an average patch development time of 122 days in the second half of 2006, the highest of any operating system.
- Sixty-eight percent of the vulnerabilities documented during this period were not confirmed by the affected vendor.
- The window of exposure for vulnerabilities affecting enterprise vendors was 47 days.
- Symantec documented 54 vulnerabilities in Microsoft Internet Explorer, 40 in the Mozilla browsers, and four each in Apple Safari and Opera.
- Mozilla had a window of exposure of two days, the shortest of any Web browser during this period.
- Twenty-five percent of exploit code was released less than one day after vulnerability publication. Thirty-one percent was released in one to six days after vulnerability publication.
- Symantec documented 12 zero-day vulnerabilities during this period, a significant increase from the one documented in the first half of 2006.
- Symantec documented 168 vulnerabilities in Oracle database implementations, more than any other database.

Vulnerability Trends Discussion

This section will discuss selected "Vulnerability Trends" metrics in greater depth, providing analysis and discussion of the trends indicated by the data. The following metrics will be discussed:

- Patch development time for operating systems
- Vendor responsiveness
- Web browser vulnerabilities
- Window of exposure for Web browsers
- Zero-day vulnerabilities
- Database vulnerabilities
- Vulnerabilities—protection and mitigation

Patch development time for operating systems

The time period between the disclosure date of a vulnerability and the release date of an associated patch is known as the "patch development time." If exploit code is created and made public during this time, computers may be immediately vulnerable to widespread attack. This metric will assess and compare the average patch development times for five different operating systems: Apple Mac OS X, Hewlett-Packard HP-UX, Microsoft Windows, Red Hat Linux (including enterprise versions and Red Hat Fedora), and Sun Microsystems Solaris.

Microsoft Windows had the shortest average patch development time of the five operating systems in the last six months of 2006. During this period, Windows had an average patch development time of 21 days based on a sample set of 39 patched vulnerabilities (figure 13). This represents an increase over the first six months of 2006, when Windows had an average patch development time of 13 days based on a sample set of 22 vulnerabilities.

Of the 39 Microsoft vulnerabilities disclosed during this period, 12 were considered high severity, 20 were medium severity, and seven were low. In the first half of 2006, of the 22 Microsoft vulnerabilities, five were considered high severity, 11 were medium severity and six were low.

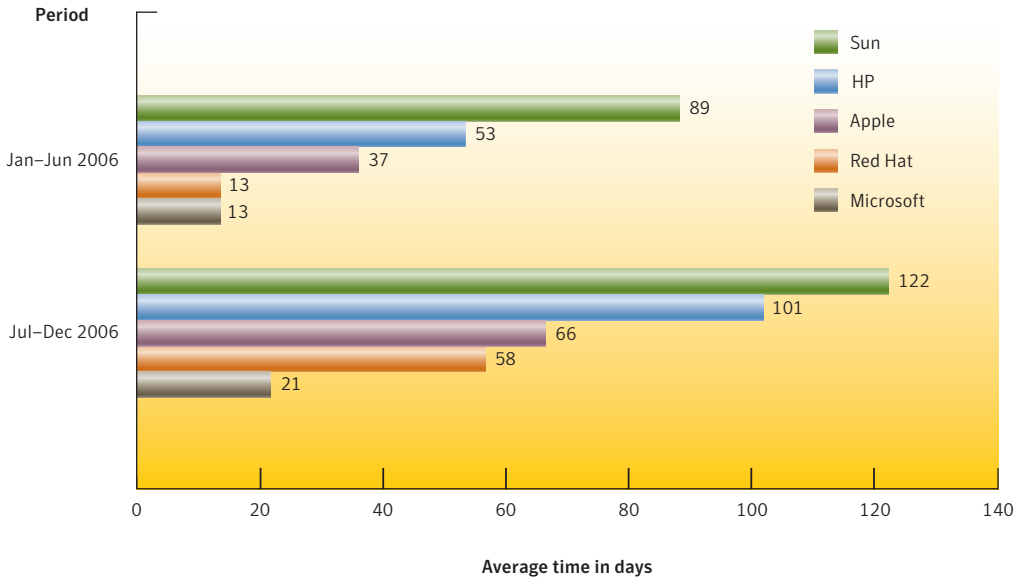

Figure 13. Patch development time for operating systems
Source: Symantec Corporation

Red Hat Linux had the second shortest average patch development time during this reporting period, with an average of 58 days for a sample set of 208 vulnerabilities. However, this is a significant increase from the 13-day average in the first half of 2006, when there were 42 patched vulnerabilities in Red Hat.

Of the 208 Red Hat vulnerabilities during the second half of 2006, two were considered high severity, 130 were medium severity, and 76 were considered low. During the first half of 2006, of the 42 vulnerabilities in Red Hat, one was considered high severity, 21 were medium severity, and 20 were low severity.

Apple Mac OS X had the third shortest average patch development time in the second half of 2006, at 66 days for a sample set of 43 vulnerabilities. This is an increase from the 37-day average in the first half of 2006 for a sample set of 21 vulnerabilities.

Out of 43 vulnerabilities in Mac OS X during the current period, one was considered high severity, 31 were medium severity, and 11 were low. In the first half of 2006, 21 vulnerabilities were documented for Apple. Of these, three were considered high severity, 12 were medium severity, and six were low.

Hewlett Packard HP-UX and Sun Solaris were ranked fourth and fifth respectively for patch development times during this period. HP had an average patch development time of 101 days for a sample set of 98 vulnerabilities. This is a significant increase from 53 days in the first half of 2006 for a sample set of seven vulnerabilities.

Of the 98 HP-UX vulnerabilities disclosed in the second half of 2006, two were considered high severity, 55 were medium severity and 41 were low. During the first half of 2006, seven vulnerabilities were disclosed for HP-UX. Of these, one was considered high severity, one was medium severity, and five were low.

Sun Solaris had an average patch development time of 122 days in the second half of 2006 for a sample set of 63 vulnerabilities. This is an increase over the 89 days documented in the first half of 2006 for a sample set of 16 vulnerabilities.

Of the 63 Sun Solaris vulnerabilities in the second half of 2006, one was considered high severity, 34 were medium severity, and 28 were low. In the first half of 2006, 16 vulnerabilities were disclosed for Sun Solaris. Of these, two vulnerabilities were considered high severity, five were medium severity, and nine were low.

In the second half of 2006, all the vendors that were studied had longer average patch development times than in the first half of the year. This corresponds to an increase in the number of patched vulnerabilities in this period. This may be because as more vulnerabilities are discovered and need to be addressed, more time is required to develop, test, and roll out patches. It may also be because vulnerabilities of higher complexity result in more complex patch development processes.

With the exception of Microsoft, all vendors were affected by longer turnarounds for patches for third-party components that are distributed with each operating system. Upon examining the sample set of vulnerabilities during this period, Symantec has observed that vulnerabilities with longer patch development times generally affected third-party components. The previous issue of the Symantec *Internet Security Threat Report* commented on the relevance of this issue for commercial UNIX vendors such as HP and Sun,[48] but it holds true for all vendors of UNIX/Linux–based operating systems.

The data suggests that third-party components are considered a lower priority than those components that are developed by the operating system vendor. However, the third-party components in question are often open source, and security patches are often provided from an upstream vendor, such as the main developer of the component. Depending on the specific operating system, many third-party components provide core functionality and are enabled by default. These components can, therefore, provide a means by which attackers can compromise computers on which they are deployed. However, administrators have some recourse if a third-party component vendor has released a patch before the operating system vendor.

The risk of exploitation in the wild is a major driving force in the development of patches. As with previous periods, Microsoft Windows was the operating system that had the most vulnerabilities with associated exploit code and exploit activity in the wild. This may have pressured Microsoft to develop and issue patches more quickly than other vendors. Another pressure that may have influenced Microsoft's relatively short patch development time is the development of unofficial patches by third-parties in response to high-profile vulnerabilities.[49]

[48] Symantec *Internet Security Threat Report*, Volume X (September 2006): http://www.symantec.com/specprog/threatreport/ent-whitepaper_symantec_internet_security_threat_report_x_09_2006.en-us.pdf : p.58
[49] http://www.securityfocus.com/brief/318

Vendor responsiveness

Vendor responsiveness is measured by the proportion of vulnerabilities that remains unconfirmed by the vendor and, therefore, unpatched over time.[50] This metric takes into account all vendors who were affected by vulnerabilities during the last three six-month reporting periods, including large-scale enterprise vendors as well as hobbyist and small commercial vendors.[51]

Vendor responsiveness is an important security consideration because, in many cases, unsanctioned, unsupported, and unmaintained software may be deployed within the organization. Software that is affected by vulnerabilities that are unconfirmed and unpatched for long periods of time may present a dormant threat to organizations.

In the second half of 2006, 68 percent of documented vulnerabilities were not confirmed by the affected vendor (figure 14). This is an increase from the first half of the year, when 61 percent of vulnerabilities were not confirmed by the vendor. In the second half of 2005, 55 percent of documented vulnerabilities were not vendor confirmed.

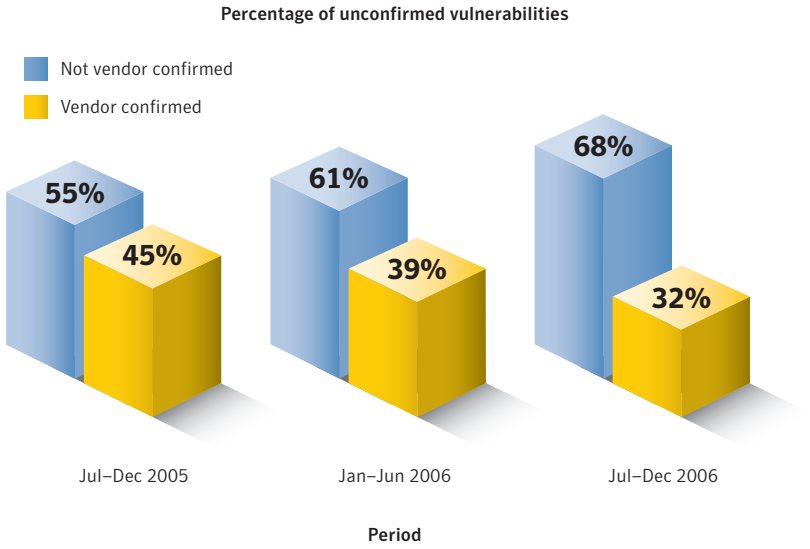

Figure 14. Vendor responsiveness
Source: Symantec Corporation

[50] Vulnerabilities that are not vendor confirmed are generally assumed to be unpatched because the vendor has not otherwise made a public statement that a patch is available. Symantec has no insight into the number of vulnerabilities that have not been publicly confirmed by the vendor but may otherwise have been patched.
[51] Hobbyist applications include non-commercial applications that often have limited development resources. With rare exception, these applications have smaller deployment than enterprise applications and present a smaller overall risk to the Internet at large; however, vulnerabilities in these applications could pose a threat to networks on which they are deployed.

In each of the last three six-month reporting periods, the majority of the documented vulnerabilities were not confirmed by vendors.[52] The proportion of non-confirmed vulnerabilities is highest in the current reporting period, likely because less time has elapsed since the vulnerabilities were initially published. However, the vulnerabilities that remain unconfirmed from previous periods still have not been confirmed in the time since initial disclosure.

It is worrisome that vendors can let vulnerabilities go unconfirmed for prolonged periods. The bulk of vulnerabilities documented in any given reporting period are associated with smaller commercial or hobbyist vendors. So, the lack of vendor confirmation can often be attributed to smaller-scale vendors who may not have dedicated security resources. In many cases, these vendors likely do not monitor security mailing lists or other resources for reports of vulnerabilities in their products. Web application vendors are also common among the list of vendors with unconfirmed vulnerabilities. In many cases, vendors have discontinued vulnerable applications or have ceased to operate, leaving administrators with no recourse other than best practices.

Symantec recommends that administrators employ vulnerability assessment services, a vulnerability management solution, and vulnerability assessment tools to evaluate the security posture of the enterprise. Where possible, problematic applications should be removed or isolated, especially if there is no vendor-provided remediation available. IPS systems can aid in detecting known attacks against such applications.

When deploying applications, administrators should ensure that secure, up-to-date versions are used, and that applications are properly configured to avoid the exploitation of latent vulnerabilities. As much as possible, enterprises are advised to avoid deploying products that are not regularly maintained or that are not supported by the vendor.

Web browser vulnerabilities

The Web browser is a critical and ubiquitous application that has become an increasingly popular subject for vulnerability researchers over the past few years. Traditionally, the focus of security researchers has been on the perimeter: servers, firewalls, and other assets with external exposure. However, security researchers and attackers now consider client-side vulnerabilities to be a more fruitful area of research and attacks. As part of this shift toward client-side issues, vulnerabilities in Web browsers have become increasingly prominent, which in turn poses a threat to end-user desktop computers.

Browsers are complex and feature rich, traits that can expose them to vulnerabilities in newly implemented features. Due to the integration of various content-handling applications—such as productivity suites and media players—browsers are a viable attack vector for many client-side vulnerabilities. This is particularly true of operating systems in which the browser is not disassociated from many other operating system processes and features.

Web browser vulnerabilities are a serious security concern, particularly due to their role in online fraud and the propagation of spyware and adware. Web browsers are particularly prone to security concerns because they come in contact with more potentially untrusted or hostile content than other applications.

[52] This discussion is based on data that was gathered at the time of writing, so data from previous periods is representative of vulnerabilities that are still unconfirmed by the vendor.

In the second half of 2006, Symantec documented 54 vulnerabilities in Microsoft Internet Explorer (figure 15). Of these, one was considered to be high severity, 13 were medium severity, and 40 were classified as low severity. This total is an increase from the 38 vulnerabilities documented in the first half of 2006. Of these, one was considered high severity, 21 were medium severity, and 16 were low severity. In the second half of 2005, 25 Internet Explorer vulnerabilities were documented.

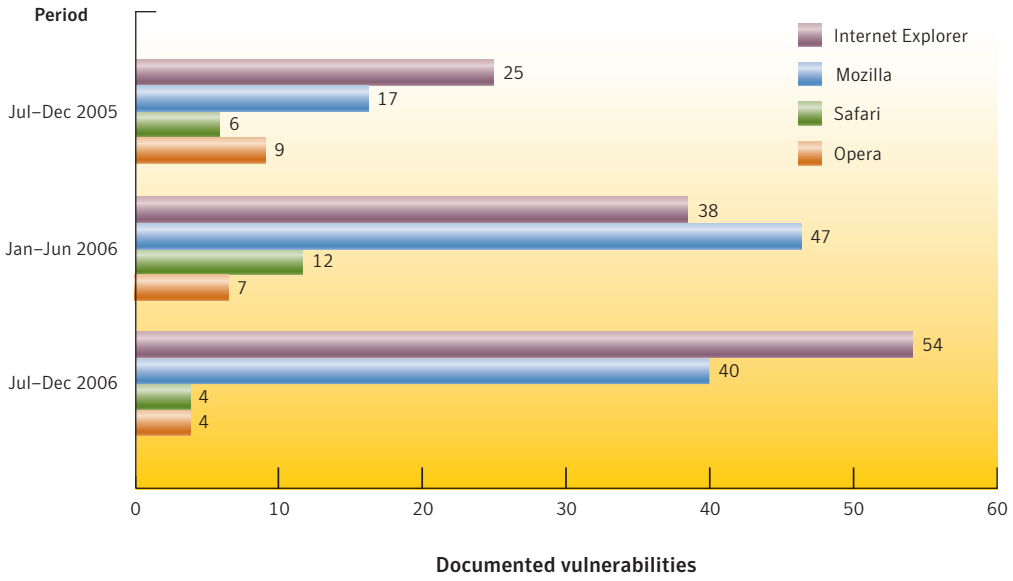

Figure 15. Web browser vulnerabilities
Source: Symantec Corporation

During the second half of 2006, 40 vulnerabilities affected the Mozilla browsers. Of these, 35 were considered to be medium severity and five were considered low. This total is a decrease from the 47 vulnerabilities that affected Mozilla browsers in the first half of 2006. Of those, 40 were considered medium severity and seven were low severity. In the second half of 2005, 17 vulnerabilities affected Mozilla browsers.

During the second half of 2006, four vulnerabilities were disclosed that affected Opera. Of these, two were low severity and the other two were medium severity. The total of four is a decrease from the seven vulnerabilities that affected Opera in the first half of 2006. Of those seven, four were considered medium severity and three were low. Symantec documented nine vulnerabilities in Opera in the second half of 2005.

Safari was also affected by four vulnerabilities in the second half of 2006. Two of these were medium-severity vulnerabilities and the other two were low severity. This is a decrease from the 12 vulnerabilities that were documented to affect Safari in the first half of 2006. Of these 12 vulnerabilities, nine were medium severity and the remaining three were low. Six vulnerabilities affected Safari in the second half of 2005.

During this reporting period, Internet Explorer was particularly affected by concerted efforts to "fuzz" the browser for new vulnerabilities. Fuzzing is a security research and quality assurance method that generally entails providing randomly generated inputs in an attempt to discover vulnerabilities and bugs. In the "Future Watch" section of the previous *Internet Security Threat Report*, Symantec predicted that the use of fuzzing technologies and techniques would result in the discovery and disclosure of new vulnerabilities.[53]

It appears that prediction is being borne out. In July 2006, security researchers embarked on a "Month of Browser Bugs," which employed various browser fuzzing tools to generate a new vulnerability for each day of the month.[54] The majority of vulnerabilities reported as a result of this project affected Internet Explorer or Windows components that were accessible through the Web browser.

In the second half of 2006, there were numerous advisories and corresponding security upgrades to Mozilla Firefox. Most of the vulnerabilities during this period were initially reported by the vendor, but were based on audits from independent researchers working in concert with the vendor. This is in contrast with Internet Explorer, for which a large number of vulnerabilities were disclosed prior to vendor notification on security mailing lists or through the "Month of Browser Bugs" initiative.

In order to protect against successful exploitation of Web browser vulnerabilities, Symantec advises users and administrators to upgrade all browsers to the latest, patched versions. Symantec recommends that organizations educate users to be extremely cautious about visiting unknown or untrusted Web sites and viewing or following links in unsolicited emails. Administrators should also deploy Web proxies in order to block potentially malicious script code.

Window of exposure for Web browsers

The window of exposure is the difference in days between the time at which exploit code affecting a vulnerability is made public and the time at which the affected vendor makes a patch available to the public for that vulnerability. During this time, the computer or system on which the affected application is deployed may be susceptible to attack, as administrators will have no official recourse against exploitation of the vulnerability. Instead they will have to resort to best practices and workarounds to reduce the risk of successful compromise.

This metric will assess the window of exposure for vulnerabilities in selected Web browsers. For this version of the *Internet Security Threat Report*, Symantec will be supplementing the Web browser window of exposure discussion with the maximum amount of time that elapsed between the disclosure of a single vulnerability and the release of an associated patch. Maximum patch times indicate the longest period of time required for a patch to be released to the public.

In the second half of 2006, Mozilla had a window of exposure of two days based on a sample set of 36 patched vulnerabilities. This is a small increase over the window of exposure of one day in the first half of 2006, which was based on three patched vulnerabilities. In the second half of 2006, Mozilla had a maximum patch development time of 33 days. In the first half of the year, the maximum patch development time was eight days.

[53] Symantec *Internet Security Threat Report*, Volume X (September 2006): http://www.symantec.com/specprog/threatreport/ent-whitepaper_symantec_internet_security_threat_report_x_09_2006.en-us.pdf : p. 29
[54] http://www.securityfocus.com/columnists/411

Mozilla has consistently had a lower average patch development time than other vendors. This may be due to security initiatives that have been undertaken by the vendor. Open-source collaboration may have also contributed to this trend since it is possible for vulnerability researchers and other volunteers to submit security patches. The majority of Mozilla vulnerabilities disclosed in the second half of 2006 were reported by the vendor when new versions of Firefox were released, as opposed to being disclosed on security mailing lists prior to vendor notification.

Interestingly, the Mozilla Foundation is one of the few vendors to offer a "bug bounty" program, which provides monetary rewards to security researchers for discovering and reporting vulnerabilities to the vendor.[55] This can affect the patch development times because it may attract security researchers who might otherwise disclose vulnerabilities prior to notifying the vendor. It may also discourage them from pursuing monetary rewards from other legitimate or black market parties.

While this does encourage some financially motivated researchers to report vulnerabilities to the vendor, it could ultimately place the vendor in the situation of bidding against others for vulnerability information, including black market or criminal buyers. Such a trend could also negatively affect smaller vendors with limited resources who are not able to pay to obtain reports of vulnerabilities in their own products.

In the second half of 2006, Microsoft Internet Explorer had a window of exposure of 10 days based on a sample set of 15 patched vulnerabilities. This is an increase from the nine-day time period in the first half of 2006, which was based on a sample set of 20 patched vulnerabilities. The maximum patch development time during the current reporting period was 78 days. In the first half of 2006, the maximum patch development time was 71 days.

The market share of Internet Explorer and Mozilla have traditionally made them more attractive targets than Opera and Safari. As a result, trends in these browsers are based on a larger data set and are less likely to be skewed by anomalous results in one or two vulnerabilities.

In the second half of 2006, Opera had a window of exposure of 23 days based on a sample set of three patched vulnerabilities. This is an increase over the window of exposure of two days in the first six months of 2006, which was based on a sample set of four patched vulnerabilities. In the second half of 2006, Opera had maximum patch development time of 46 days. This can be attributed to a few vulnerabilities in a small sample data set that disproportionately affected the average. In the first half of 2006, a maximum of seven days elapsed before a patch was available.

During the second half of 2006, Safari had a window of exposure of 62 days, an increase over the five-day window in the first half of 2006. However, this increase is based on a sample set of only one vulnerability, a sample size that is too small to ensure valid conclusions. This vulnerability affected a third-party HTML rendering component, so it is possible that the third-party nature may have slowed the patch release time. In the first half of 2006, the maximum patch development time was 21 days for a sample set of four vulnerabilities.

All browser vendors experienced a longer window of exposure during the current reporting period. In some cases, the increase was relatively small, as was the case with Internet Explorer and Mozilla. With Opera and Safari, however, the window of exposure experienced a large increase, although this is skewed by a smaller sample set of patched vulnerabilities and exploits.

[55] http://www.mozilla.org/security/bug-bounty.html

Opera and Safari both had instances in which a relatively longer patch development time for a single vulnerability caused the average patch development time to be longer than those for the first half of 2006. None of the vulnerabilities affecting Opera and Safari during this period had any indication of exploit activity in the wild, and so the relative level of risk may be an influence in how quickly the vulnerability has been patched.

Zero-day vulnerabilities

A zero-day vulnerability is one for which there is sufficient public evidence to indicate that the vulnerability has been exploited in the wild prior to being publicly known. It may not have been known to the vendor prior to exploitation, and the vendor had not released a patch at the time of the exploit activity.

Zero-day vulnerabilities represent a serious threat in many cases because there is no patch available for them and because they will likely be able to evade purely signature-based detection. It is the unexpected nature of zero-day threats that causes concern, especially because they may be used in targeted attacks and in the propagation of malicious code. As Symantec predicted in Volume IX of the *Internet Security Threat Report*, a black market for zero-day vulnerabilities has emerged that has the potential to put them into the hands of criminals and other interested parties.[56]

In the second half of 2006, Symantec documented 12 zero-day vulnerabilities (figure 16). This is a significant increase compared to the first half of 2006 and the second half of 2005 when only one zero-day vulnerability was documented for each reporting period.

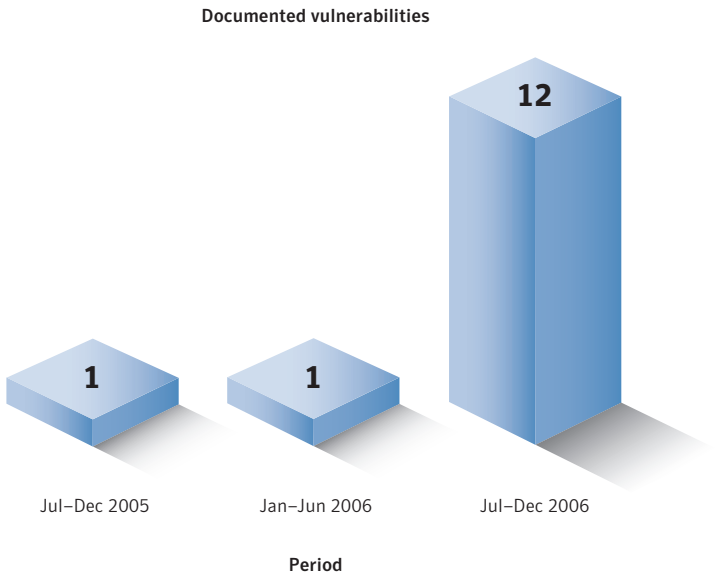

Figure 16. Zero-day vulnerabilities
Source: Symantec Corporation

[56] Symantec *Internet Security Threat Report*, Volume IX (March 2006): http://eval.veritas.com/mktginfo/enterprise/white_papers/ent-whitepaper_symantec_internet_security_threat_report_ix.pdf : p. 21

Numerous high-profile zero-day vulnerabilities were discovered in the second half of 2006. This activity peaked in September of 2006, when four zero-day vulnerabilities were documented. The majority of these were client-side vulnerabilities that affected Office applications, Internet Explorer, and ActiveX controls. Many of these may have been discovered through the use of fuzzing technologies.

In August 2006, Microsoft released a security bulletin to address a zero-day vulnerability in the Windows Server Service.[57] This vulnerability was not publicly known prior to the release of the bulletin, but Microsoft,[58] SANS,[59] and US-CERT made statements that it was being actively exploited in the wild prior to the bulletin.[60]

While it is believed that zero-day vulnerabilities have been a threat in the past, the recent increase in incidents may be partially due to improvements in capabilities to detect these attacks in the wild. Such capabilities include improved vulnerability-handling procedures within organizations, improved cooperation between enterprises and vendors, and better technologies for the detection and analysis of exploits and malicious code.

In order to protect against zero-day vulnerabilities, Symantec recommends that administrators deploy IDS/IPS systems and regularly updated antivirus software. Security vendors may provide rapid response to recently discovered zero-day vulnerabilities in the wild by developing and implementing new or updated IDS/IPS and antivirus signatures before a patch has been released by the affected vendor. Behavior-blocking solutions and heuristic signatures may also provide protection against zero-day vulnerabilities.

In addition, some IPS systems may provide further protection against memory corruption vulnerabilities in the form of ASLR and by making memory segments non-executable. These measures may complicate the exploitation of such vulnerabilities and make it more difficult for attack payloads to execute; however, they may not protect all applications by default.

Database vulnerabilities

Data is one the most important assets of any organization and is, therefore, a valuable target for attackers. Securing the confidentiality, integrity, and availability of data should be among the top priorities for enterprises. While databases are usually deployed deep within the organization's infrastructure, they are often accessed by middle-ware and third-party components that are granted a certain degree of trust.[61] This can expose database implementations to a variety of attacks that fall outside of the protection of traditional network security mechanisms, such as firewalls and IDS systems.

With this version of the *Internet Security Threat Report*, Symantec will assess database vulnerabilities for the first time. This report will discuss vulnerabilities that affected the following major database implementations during the second half of 2006: IBM DB2, Microsoft SQL Server, MySQL, Oracle, and PostgreSQL.

In the second half of 2006, 168 vulnerabilities were documented that affected Oracle databases (figure 17). This is a slight decrease from the 169 vulnerabilities disclosed in the first half of 2006 and an increase over the 131 in the second half of 2005.

[57] http://www.securityfocus.com/bid/19409/info
[58] http://www.microsoft.com/technet/security/bulletin/ms06-040.mspx
[59] http://isc.incidents.org/diary.html?storyid=1556&dshield=938c1242911bc722f0c63baf4a21df2c
[60] http://www.us-cert.gov/cas/techalerts/TA06-220A.html
[61] Database middle-ware is defined as services and applications that provide database access and interconnectivity.

During the second half of 2006, five vulnerabilities were documented in IBM DB2 databases. This is a slight increase from the four vulnerabilities documented during the first half of 2006. Seven vulnerabilities affected IBM DB2 during the second half of 2005.

Symantec documented five vulnerabilities in MySQL during the second half of 2006. This is a slight decrease from the six vulnerabilities that affected it during the first half of 2006. One vulnerability was documented in MySQL in the second half of 2005.

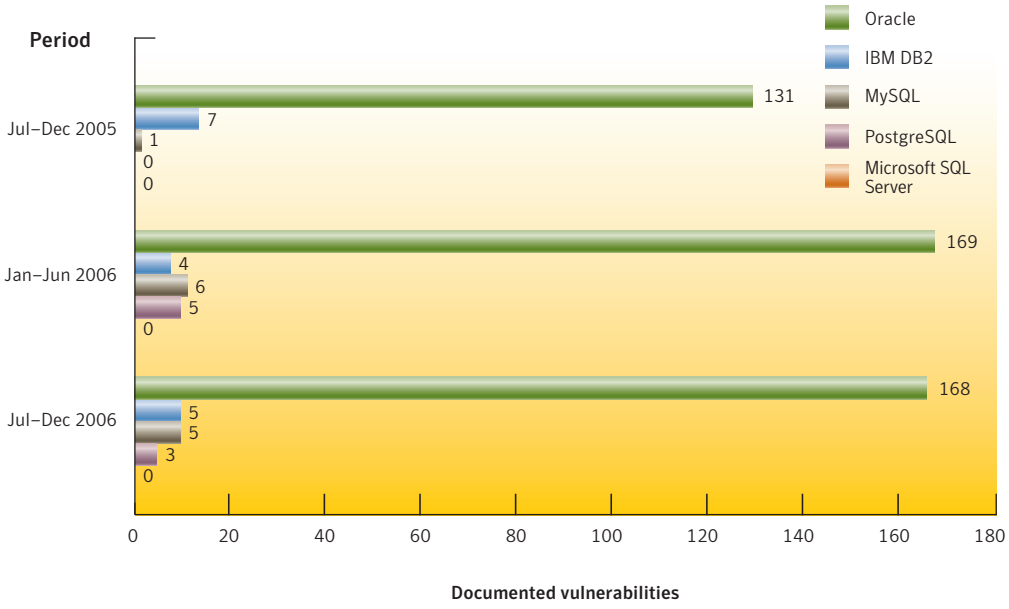

Figure 17. Database vulnerabilities
Source: Symantec Corporation

PostgreSQL was affected by three vulnerabilities in the second half of 2006. This is a decrease from the five vulnerabilities that affected PostgreSQL during the first half of 2006. No vulnerabilities in PostgreSQL were document in the second half of 2005.

Microsoft SQL Server was the final database that was assessed for this discussion. It was not affected by any vulnerabilities during any of the reporting periods.

Oracle has traditionally presented the most high-profile target to attackers due to its large market share. In addition to this, Oracle's database implementations offer a greater feature set and a broader range of database products than many of the other database vendors. The more features an application has, the more code that is available in which to find vulnerabilities, and the more code that must be audited for vulnerabilities. This can equate to a higher proportion of vulnerabilities, depending on the nature and complexity of the features.

Other database implementations such as MySQL and PostgreSQL have been more conservative when introducing new features. They have only recently adopted many features common to commercial database implementations such as Oracle, IBM DB2, and Microsoft SQL Server. This may account for the significantly lower volume of vulnerabilities in MySQL and PostgreSQL, but it also presents the possibility that, as they become more complex and more widely adopted, their share of vulnerabilities may also increase.

In contrast, Microsoft SQL Server has been free of vulnerabilities throughout the past three six-month reporting periods. Some sources have attributed this trend to Microsoft security initiatives such as the Security Development Lifecycle, which was employed during the development of SQL Server 2005.[62] That said, Microsoft SQL Server 2000 has been the only database to suffer from a widespread malicious code attack, namely SQL Slammer (also known as the W32.SQLExp.Worm), which was first detected in January 2003.[63] This may have been a major reason for developing future versions that would not be susceptible to such attacks.

Symantec recommends that administrators configure firewalls to restrict all external access to database servers and minimize the risk of direct remote attacks against databases. Database intrusion detection systems should also be deployed to detect and provide audit logs for unauthorized access attempts.

Data security and integrity is often a requirement for policy compliance, and organizations and individuals may be held responsible for data breaches that threaten personal information. Symantec recommends that enterprises engage in auditing for policy compliance and encourages the use of policy compliance tools. Databases may also be exposed to attacks in components that interface with the database, such as Web applications. Web application firewalls may help to detect and prevent Web-based attacks on the database.

Vulnerabilities—protection and mitigation

Administrators should employ a good asset management system to track what assets are deployed on the network and to determine which ones may be affected by the discovery of new vulnerabilities. Vulnerability assessment technologies should also be used to detect known vulnerabilities in deployed assets. Administrators should monitor vulnerability mailing lists and security Web sites to keep abreast of new vulnerabilities in Web applications.

Enterprises should subscribe to a vulnerability alerting service in order to be notified of new vulnerabilities. They should also manage their Web-based assets carefully. If they are developing Web applications in-house, developers should be educated about secure development practices, such as the Security Development Lifecycle and threat modeling.[64]

Symantec recommends the use of secure shared components that have been audited for common Web application vulnerabilities. If possible, all Web applications should be audited for security prior to deployment. Web application security solutions and a number of products and services are available to detect and prevent attacks against these applications.

[62] David Litchfield, "Which database is more secure? Oracle vs. Microsoft" (Nov. 21, 2006): http://www.databasesecurity.com/dbsec/comparison.pdf : p. 3
[63] http://www.symantec.com/security_response/writeup.jsp?docid=2003-012502-3306-99
[64] The Security Development Lifecycle is a development paradigm that incorporates security at every stage from the initial architecture to programming, and in the quality assurance/testing phases. Threat modeling is a security auditing methodology that involves formally identifying and mapping out all possible attack vectors for an application.

Malicious Code Trends

Symantec gathers malicious code data from over 120 million desktops that have deployed Symantec's antivirus products in consumer and corporate environments. The Symantec Digital Immune System and Scan and Deliver technologies allow customers to automate this reporting process. This discussion is based on malicious code samples reported to Symantec for analysis between July 1 and December 31, 2006.

In previous editions of the Symantec *Internet Security Threat Report*, the number and volume of threats analyzed were based upon the number of reports received by enterprise and home users. For the first time, this report will also examine malicious code types and propagation vectors based upon potential infections. This allows Symantec to determine which sample of malicious code attempted to infect a computer and the volume of potential infections worldwide.

Symantec categorizes malicious code in two ways: families and variants. A family is a new, distinct sample of malicious code. For instance, W32.Sober@mm (also known as Sober) was the founding sample, or the primary source code, of the Sober family. In some cases, a malicious code family may have variants. A variant is a new iteration of the same family, one that has minor differences but that is still based on the original. A new variant is created when the source code of a successful virus or worm is modified slightly to bypass antivirus detection definitions developed for the original sample. For instance, Sober.X is a variant of Sober.

This discussion will include any prevention and mitigation measures that might be relevant to the particular threats being discussed. However, Symantec recommends that certain best security practices always be followed to protect against malicious code infection. Administrators should keep patch levels up to date, especially on computers that host public services—such as HTTP, FTP, SMTP, and DNS servers—and are accessible through a firewall or placed in a DMZ. Email servers should be configured to only allow file attachment types that are required for business needs and to not accept email that appears to come from within the company but that actually originates from external sources. Additionally, Symantec recommends that ingress and egress filtering be put in place on perimeter devices to prevent unwanted activity.

End users should employ defense-in-depth strategies, including the deployment of antivirus software and a personal firewall. Users should update antivirus definitions regularly. They should also ensure that all desktop, laptop, and server computers are updated with all necessary security patches from their software vendors. They should never view, open, or execute any email attachment unless it is expected and comes from a trusted source, and unless the purpose of the attachment is known.

Malicious Code Trends Highlights

The following section will offer highlights of the malicious code trends that Symantec observed during this period. Following this overview, the *Internet Security Threat Report* will discuss selected metrics in greater depth, providing analysis and discussion of the trends indicated by the data.

- Of the top ten new malicious code families detected in the last six months of 2006, five were Trojans, four were worms, and one was a virus.
- The most widely reported new malicious code family this period was that of the Stration worm.[65]
- Symantec honeypot computers captured a total of 136 previously unseen malicious code threats between July 1 and December 31, 2006.
- During this period, 8,258 new Win32 variants were reported to Symantec, an increase of 22 percent over the first half of 2006.
- Worms made up 52 percent of the volume of malicious code threats, down from 75 percent in the previous period.
- The volume of Trojans in the top 50 malicious code samples reported to Symantec increased from 23 percent to 45 percent.
- Trojans accounted for 60 percent of the top 50 malicious code samples when measured by potential infections.
- Polymorphic threats accounted for three percent of the volume of top 50 malicious code reports this period, up from one percent in the two previous periods.
- Bots made up only 14 percent of the volume of the top 50 malicious code reports.
- Threats to confidential information made up 66 percent of the volume of the top 50 malicious code reported to Symantec.
- Keystroke logging threats made up 79 percent of confidential information threats by volume of reports.
- Seventy-eight percent of malicious code that propagated did so over SMTP, making it the most commonly used propagation mechanism.
- Malicious code using peer-to-peer to propagate rose from 23 percent of all propagating malicious code in the first six months of 2006 to 29 percent in the last half of the year.
- The majority of malicious code reports during this period originated in the United States.
- During the second half of 2006, 23 percent of the 1,318 documented malicious code instances exploited vulnerabilities.
- MSN Messenger was affected by 35 percent of new instant messaging threats in the second half of the year.

[65] http://www.symantec.com/security_response/writeup.jsp?docid=2006-092111-0525-99

Malicious Code Trends Discussion

This section will discuss selected malicious code metrics in greater depth, providing analysis and discussion of the trends indicated by the data. The following metrics will be discussed in depth:

- Top ten new malicious code families
- Previously unseen malicious code
- Malicious code types
- Threats to confidential information
- Propagation vectors
- Malicious code that exploits vulnerabilities

Top ten new malicious code families

Of the top ten new malicious code families detected in the last six months of 2006, five were Trojans, four were worms, and one was a virus (table 5). One of the Trojans also had back door capabilities. This indicates that attackers may be moving towards using Trojans as a means of installing malicious code on computers. As Trojans do not propagate, they allow attackers to perform targeted attacks without drawing attention to themselves. The longer a threat remains undiscovered in the wild, the more opportunity it has to compromise computers before measures can be taken to protect against it.

Rank	Sample	Type	Vectors	Impacts/Features
1	Stration	Worm	SMTP	Downloads and installs other threats
2	Gampass	Trojan	N/A	Steals online gaming passwords
3	Shufa	Worm	Yahoo! IM, SMTP	Steals passwords for Lineage online game
4	Bacalid	Virus	File sharing	Polymorphic virus that can download malicious files
5	Horst	Trojan, Back door	N/A	Relays email for spam
6	Annunci	Worm	SMTP	Dials a high-cost phone number
7	Pasobir	Worm	File sharing	Steals instant messenger passwords
8	Jakposh	Trojan	N/A	Redirects search queries to other Web sites
9	Linkmediac	Trojan	N/A	Displays pop-up ads and sends system data to an attacker
10	Zonebac	Trojan	N/A	Lowers Internet Explorer security settings

Table 5. Top ten new malicious code families
Source: Symantec Corporation

The most widely reported new malicious code family during this reporting period was that of the Stration worm.[66] More than 150 variants of this worm were discovered in the last six months of 2006. Stration sends copies of itself with various subject headers, messages, and attachment names to email addresses that are gathered from compromised computers. Once installed on a computer, the worm also attempts to download and execute remote files from predetermined Web sites.

The Gampass[67] information-stealing Trojan and the Shufa[68] worm were the second and third most common new families, respectively, in the second half of 2006. They are part of a growing trend towards threats that steal account information for online games.

[66] http://www.symantec.com/security_response/writeup.jsp?docid=2006-092111-0525-99
[67] http://www.symantec.com/security_response/writeup.jsp?docid=2006-111201-3853-99
[68] http://www.symantec.com/security_response/writeup.jsp?docid=2006-080815-5056-99

A secondary economy has emerged on various online auction sites in which users buy and sell items (such as those that players are awarded for accomplishing goals within these games). As the popularity of these games continues to grow, so does the potential for attackers to exploit them for financial gain. Once an attacker has stolen a user's account information, he or she can sell the user's items and keep the profits. Symantec expects to continue to see new threats that target online gamers.

As noted in the "Future Watch" section of the previous edition of the *Internet Security Threat Report*, there appears to be renewed interest in polymorphic viruses among malicious code authors.[69] The Bacalid family of polymorphic viruses was the fourth most prolific new threat in the second half of 2006.[70] While most viruses simply replicate by infecting executable files, Bacalid also attempts to download and execute other malicious threats from a list of Web sites contained within its code. As malicious code authors continue to counter effective antivirus defenses, Symantec anticipates that they will increasingly adopt polymorphic techniques to evade detection.

Previously unseen malicious code

Previously unseen malicious code refers to samples captured by Symantec's honeypot network that have not been previously detected and for which antivirus signatures have not yet been developed.[71] This metric is intended to give readers an understanding of the number of new threats that may exist for which there are no antivirus signatures available, which could leave users' computers susceptible to compromise.

This metric was introduced in the previous volume of the *Internet Security Threat Report*. However, since that time the methodology has been revised in order to offer a more complete picture of previously unseen malicious code threats that are captured by Symantec's honeypot computers.

Between July 1 and December 31, 2006, Symantec honeypot computers captured a total of 136 previously unseen malicious code threats. This is up from 98 new malicious code threats that were captured in the previous period. In other words, there were more than five compromises per week, on average, by previously unseen threats in the second half of 2006 compared to less than four per week in the first half of the year. Antivirus programs did not previously detect these threats, resulting in the need to create new detection signatures for them.

Previously unseen threats are particularly dangerous because traditional defenses, such as some signature-based antivirus products, are typically unable to detect them. Until antivirus signatures are developed and antivirus programs updated, computers could be susceptible to infection by these threats. Generic signatures may also block previously unseen threats. Behavior-blocking solutions and heuristic technologies may also provide protection against previously unseen malicious code.

Administrators should also maintain up-to-date antivirus definitions to ensure that their computers are protected from new threats at the earliest possible time. Enabling heuristic detection within antivirus products may also help detect previously unseen threats before traditional antivirus signatures are available. In the case of previously unseen threats that exploit vulnerabilities in order to propagate, applying appropriate patches as soon as they are available will prevent exploitation. If patches are not available, blocking access to the vulnerable service at the firewall will help protect against exploitation by previously unseen malicious code.

[69] Symantec *Internet Security Threat Report*, Volume X (September 2006): http://www.symantec.com/specprog/threatreport/ent-whitepaper_symantec_internet_security_threat_report_x_09_2006.en-us.pdf : p. 26
[70] http://www.symantec.com/security_response/writeup.jsp?docid=2006-090109-5610-99
[71] A honeypot is an Internet-connected system that acts as a decoy, allowing an attacker to enter the system so that the attacker's behavior inside the compromised system can be observed.

Malicious code types

During the current reporting period, worms made up 52 percent of the volume of the top 50 malicious code reports, down from 75 percent in the previous period (figure 18).[72] This drop can largely be attributed to the decline in reports of major worms such as Sober.X,[73] Blackmal.E,[74] and Netsky.P[75] since the first half of 2006. The longer a threat has been in the wild, the more time users will have had to update their detection signatures. The volume of these worms has likely declined because users have installed antivirus definitions that detect them. This idea is enforced by the fact the number of unique samples of worms in the top 50 malicious code reports remained fairly constant over the last six months of 2006. During this period, 36 worms were reported to Symantec, compared to 38 in the previous period.

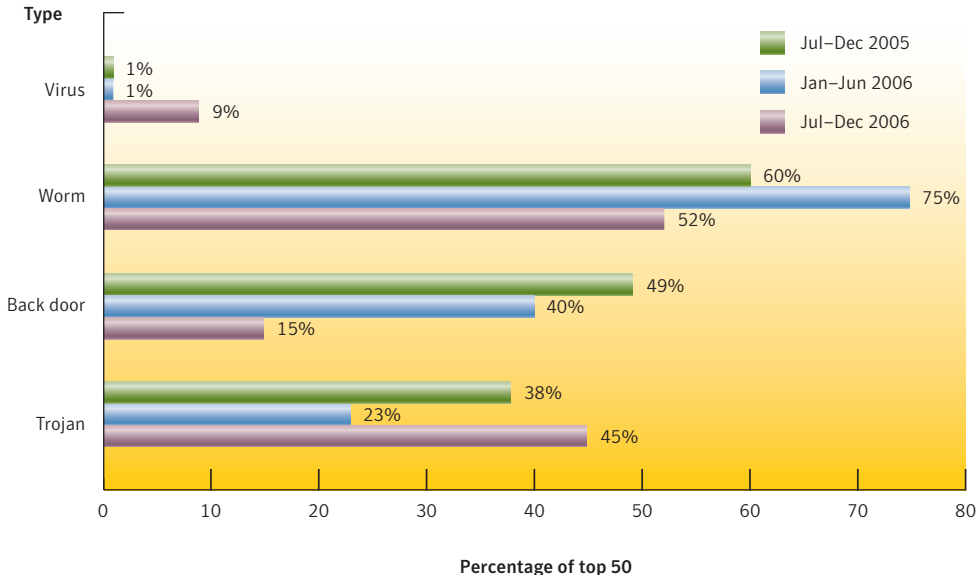

Figure 18. Malicious code types by volume
Source: Symantec Corporation

The volume of Trojans in the top 50 malicious code samples reported to Symantec increased significantly in the last six months of 2006. During this period, they constituted 45 percent of the volume of the top 50 malicious code samples, a significant increase over the 23 percent last period.

As is discussed in the "Future Watch" section of this report, attackers are moving towards staged downloaders, also referred to as modular malicious code. These are small, specialized Trojans that download and install other malicious programs, such as back doors or worms. Many of these Trojans are installed using Web browser vulnerabilities and zero-day vulnerabilities in other applications (as discussed in the "Zero-day vulnerabilities" section of this report). During the current period, 75 percent of the volume of the top 50 malicious code reports contained a modular component such as this.

[72] It is important to note that a malicious code sample can be classified in more than one threat type category. For example, bots such as variants of the Mytob family are classified as both a worm and a back door. As a result, cumulative percentages of threat types in the top 50 malicious code reports may exceed 100.
[73] http://www.symantec.com/security_response/writeup.jsp?docid=2005-111915-0848-99
[74] http://www.symantec.com/security_response/writeup.jsp?docid=2006-011712-2537-99
[75] http://www.symantec.com/security_response/writeup.jsp?docid=2004-032110-4938-99

Back doors made up only 15 percent of the volume of top 50 malicious code reports during this reporting period, down from 40 percent in the first half of 2006. While the volume of back door programs has declined, this does not necessarily mean that they are being used less. As described above, many Trojans download back doors after compromising a computer. This allows attackers to initially compromise a computer with a new, previously unseen malicious code sample in order to disable security applications and install a known back door program.

In many cases, because of security measures that are in place, only the first stage of this process will be successful. For instance, a firewall may prevent the Trojan from downloading further components or authorities may detect the additional threats and shut down the computer hosting them.

An example of this is the system consisting of Mixor.C,[76] Galapoper.A,[77] and Abwiz.F.[78] Mixor.C is a mass-mailing worm that installs a copy of the Galapoper.A Trojan on a compromised computer. This Trojan then connects to a remote Web site and downloads an instruction list of actions to perform. One of these actions is to download and execute a copy of the Abwiz.F Trojan. Abwiz.F then uploads information about the compromised computer to a Web site and allows the computer to be used by the attacker to relay spam email.

Spam can cause significant problems for a user. High volumes of email originating from a computer can cause it to be added to a block list or the user's Internet connectivity to be suspended by his or her Internet service provider. This could also affect the connectivity of other end users, as the network block to which the user's IP address belongs could also be added to the block list.

Spam can also be problematic for enterprises. If an organization's address space is blocked, it could prevent business-related email from meeting its intended destination, which could be disruptive to business communications.

In the current period, there were a large number of Mixor.C infections, but a smaller number of Galapoper.A infections and an even smaller number of Abwiz.F infections. This indicates that not all Mixor.C infections were successful in installing Galapoper.A. In turn, the Trojan was not always successful in downloading Abwiz.F. This could be due to firewalls or because the initial infection was discovered before all the components were able to perform all their functions. Additionally, authorities may have shut down the Web sites hosting the downloaded components before the earlier stages were activated.

For the first time, in this edition of the *Internet Security Threat Report*, Symantec is assessing malicious code according to the number of unique samples reported to Symantec and the number of potential infections. This is an important distinction. In some cases, a threat that may trigger a high number of reports may not cause a large number of potential infections and *vice versa*. The reasons for this will be made clear in the ensuing discussion.

The distinction between malicious code reports and infections is well illustrated by comparing worm and Trojan activity. While worms made up 52 percent of malicious code reports in the second half of 2006, they caused only 37 percent of potential infections (figure 19). The main reason for this is that mass-mailing worms generate a significant number of email messages to which they attach their malicious

[76] http://www.symantec.com/security_response/writeup.jsp?docid=2006-103115-0022-99
[77] http://www.symantec.com/security_response/writeup.jsp?docid=2006-042013-1813-99
[78] http://www.symantec.com/security_response/writeup.jsp?docid=2006-032311-1146-99

code. Each message that is detected will generate a malicious code report. Because of the high volume of email that one worm can generate, a single infection can result in many reports. However, once a malicious code sample is detected, antivirus signatures are quickly developed that can protect against subsequent potential infections by that sample. So, only a small percentage of the high volume of email messages will result in additional infections.

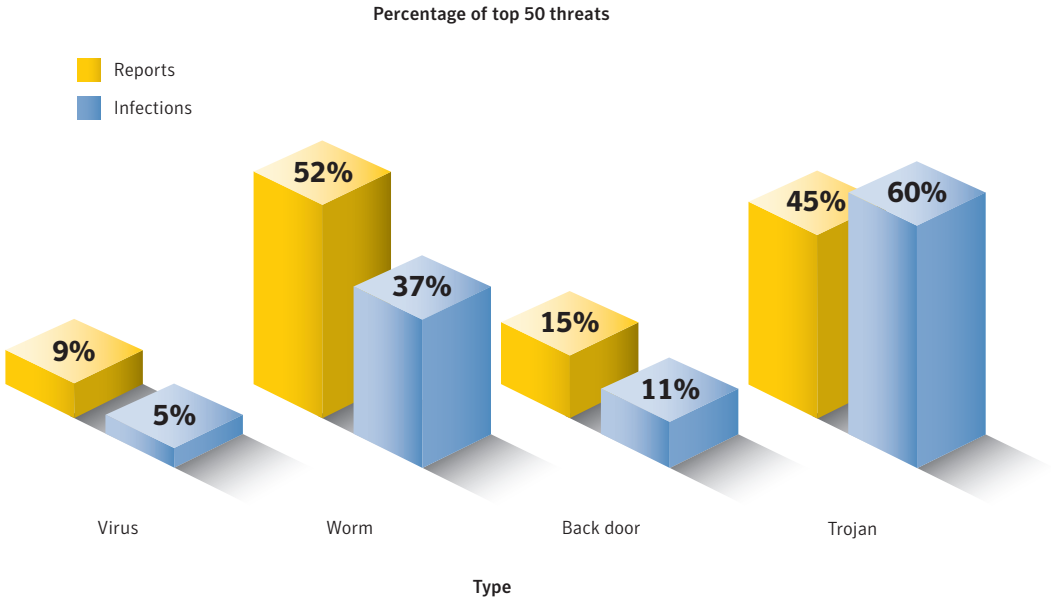

Figure 19. Malicious code types by source, July–December 2006
Source: Symantec Corporation

Trojans, on the other hand, only constituted 45 percent of the volume of the top 50 malicious code samples reported during the last six months of 2006. However, they accounted for 60 percent of potential infections by the top 50 malicious code samples during the same period. Since Trojans do not contain any propagation mechanisms, they do not proliferate as widely as mass-mailing worms, resulting in fewer reports. Because they are frequently installed by exploiting Web browser and zero-day vulnerabilities, a Trojan report is more likely to be the result of an infection. Consequently, the ratio of infections to reports is likely to be higher for Trojans than for worms.

Symantec expects the proportion of Trojans to increase as long as they remain a reliable means for attackers to compromise computers. Worms will likely continue to decline somewhat; however, a highly successful new worm could cause the proportion to increase to previous levels. For example, a new easily exploitable remote vulnerability in a network service could be exploited by a worm, resulting in rapid propagation and a high volume of infections.

Threats to confidential information

Some malicious code programs are designed specifically to expose confidential information that is stored on an infected computer. These threats may expose sensitive data such as system information, confidential files and documents, or logon credentials. Some malicious code threats, such as back doors, can give a remote attacker complete control over a compromised computer.

Threats to confidential information are a particular concern because of their potential for use in criminal activities. With the widespread use of online shopping and Internet banking, compromises of this nature can result in significant financial loss, particularly if credit card information or banking details are exposed.

Within the enterprise, exposure of confidential information can lead to significant data leakage. If it involves customer-related data—such as credit card information—this can severely undermine customer confidence as well as violate local laws.[79] Sensitive corporate information, including financial details, business plans, and proprietary technologies, could also be leaked from compromised computers.

In the last six months of 2006, threats to confidential information made up 66 percent of the volume of top 50 malicious code reported to Symantec (figure 20). This is an increase over the 48 percent reported in the first half of the year and the 55 percent reported during the second half of 2005.

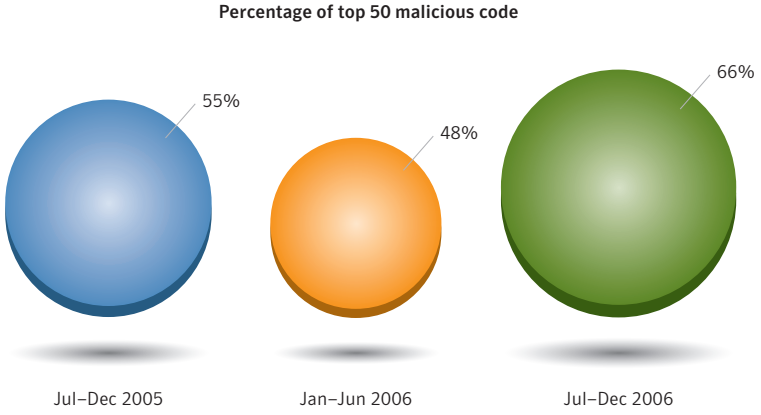

Figure 20. Threats to confidential information by volume
Source: Symantec Corporation

Malicious code can expose confidential information in a variety of ways. The most common method is by allowing remote access to the compromised computer through a back door. In this method, the attacker typically uses a specialized application to connect to the compromised computer. He or she can then perform numerous actions such as taking screenshots, changing configuration settings, and uploading, downloading, or deleting files.

[79] Many countries have implemented their own laws in this regard, such as the UK Data Protection Act, which can be found at http://www.opsi.gov.uk/ACTS/acts1998/19980029.htm

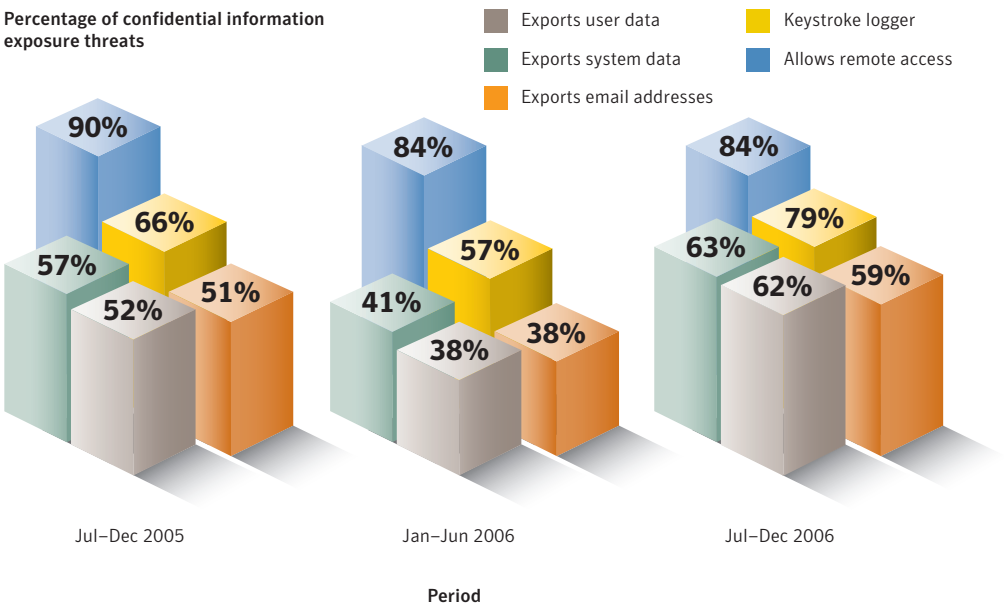

Figure 21. Threats to confidential information by type
Source: Symantec Corporation

In the second half of the 2006, threats that allow remote access, such as back doors, made up 84 percent of confidential information threats by volume of reports, the same as in the first half of the year, but a decrease from 90 percent in the second half of 2005 (figure 21). During this reporting period, threats that allow remote access made up 87 percent of threats by potential infection (figure 22). While a threat that allows remote access, such as a back door, could give an attacker full access to a computer, the attacker must typically access it manually. This likely explains why the numbers of reports (84 percent) are similar to the number of potential infections during this reporting period (87 percent).

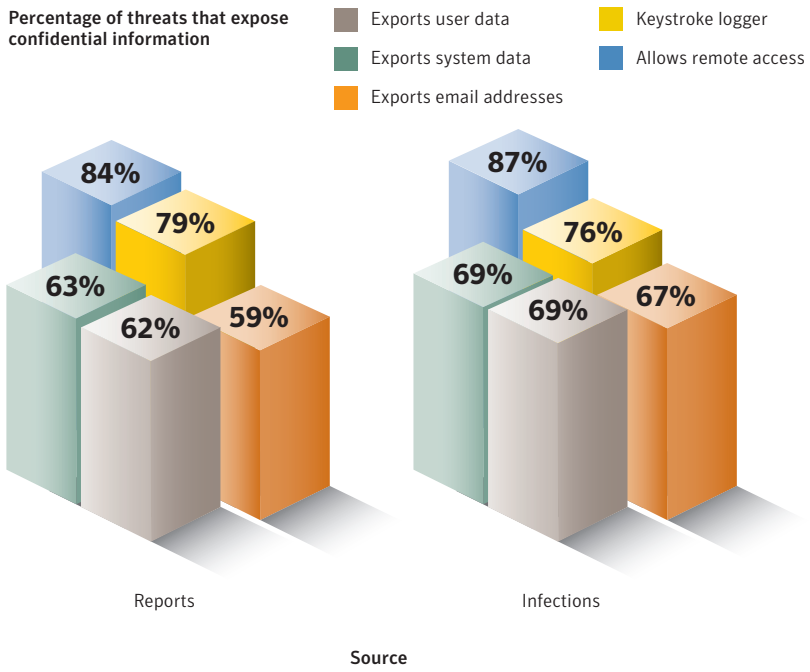

Figure 22. Threats to confidential information types by source, July–December 2006
Source: Symantec Corporation

While the volume of threats that allow remote access has decreased, the volume of threats that log keystrokes and export user and system data have all increased. Keystroke logging threats made up 79 percent of confidential information threats by volume of reports, up from 57 percent in the first half of the year and 66 percent in the second half of 2005 (figure 21). During the current reporting period, keystroke loggers made up 76 percent of confidential information threats by infection (figure 22). A keystroke logger will record keystrokes on the compromised computer. It usually either emails the log to the attacker or uploads it to a Web site that is under the attacker's control. This makes it easier for an attacker to gather confidential information from a large number of compromised computers with minimal effort.

Threats that could be employed to export user data accounted for 62 percent of confidential information threats by volume during this reporting period, up from 38 percent in the first half of the year. Furthermore, 63 percent of threats to confidential information reported during the last six months of 2006 could be used to export system data, compared to 41 percent in the first half of 2006. These forms of data leakage can aid an attacker in stealing a user's identity or launching further attacks. If the attacker has access to the user's personal and system data, they can use this to craft a more targeted social engineering attack tailored to that particular user.

hgshgs

Propagation mechanisms

Worms and viruses use various means to transfer themselves, or propagate, from one computer to another. These are collectively referred to as propagation mechanisms. This section will assess some of the propagation mechanisms used by malicious code samples reported to Symantec in the second half of 2006.[80]

In the second half of 2006, SMTP remained the most commonly used propagation mechanism. During this period, 78 percent of malicious code that propagated did so over SMTP (figure 23). This is a decrease from 98 percent in the first half of the year. It can largely be attributed to a decrease in reports of mass-mailing worms, including Blackmal.E and Sober.X, as was discussed previously in the "Malicious code types" section.

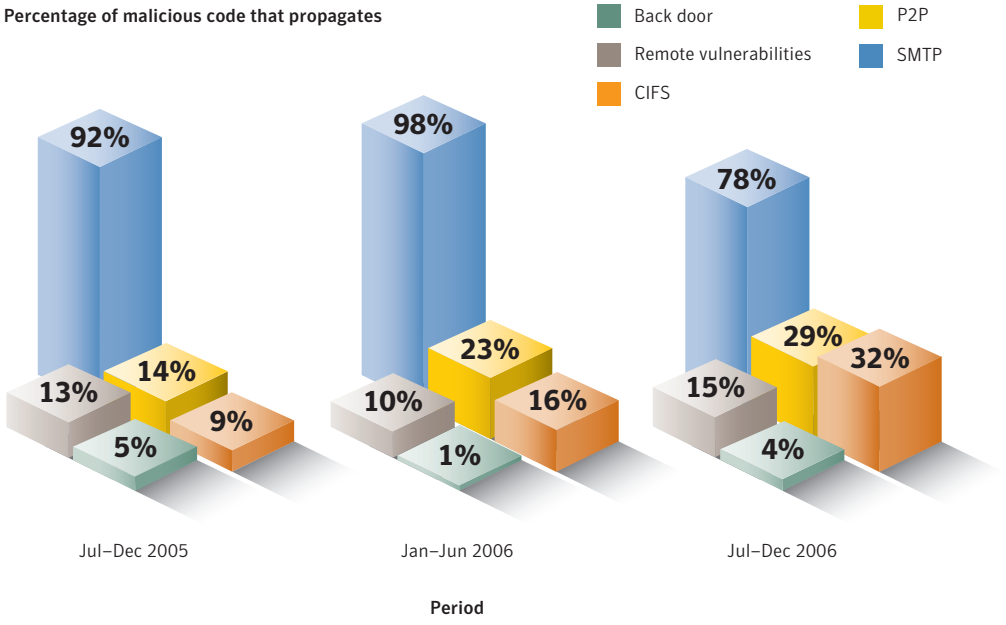

Figure 23. Propagation mechanisms
Source: Symantec Corporation

While malicious code propagating over SMTP decreased during this period, all other vectors experienced an increase. This is likely the result of an effort by attackers to diversify the way their threats proliferate. Good email scanning applications and increased user knowledge of mass-mailing threats have reduced the effectiveness of email as a propagation mechanism. As a result, some attackers are resorting to other propagation mechanisms or incorporating additional propagation mechanisms into a mass-mailing worm.

Malicious code that propagated by CIFS made up 32 percent of malicious code that propagated in the second half of 2006. This is a 100 percent increase over the first half of 2006 when 16 percent of propagating code spread by this mechanism. This is largely due to the proliferation of the Looked.P worm.[81] This worm not only copies itself to network shares with weak password protection, it also contains a viral component to infect executable files on a compromised computer. Other variants of this worm also experienced some success during this reporting period, particularly Looked.O,[82] which shared an almost identical feature set with Looked.P.

[80] It should be noted that some malicious code samples use more than one mechanism to propagate. As a result, cumulative percentages presented in this discussion may exceed 100 percent.
[81] http://www.symantec.com/security_response/writeup.jsp?docid=2006-071212-0124-99
[82] http://www.symantec.com/security_response/writeup.jsp?docid=2006-071212-0828-99

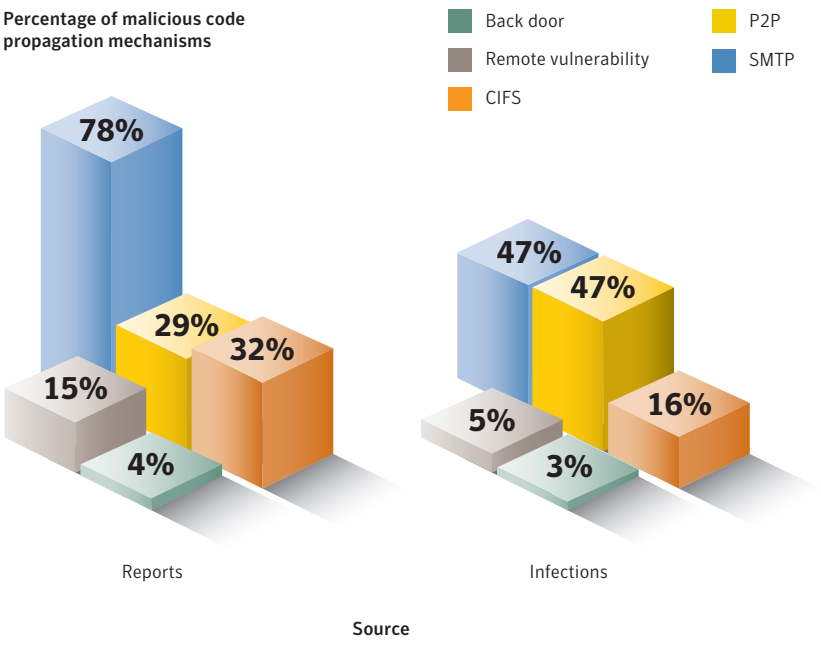

Figure 24. Propagation mechanisms by source, July–December, 2006
Source: Symantec Corporation

In the second half of 2006, there were significant differences in propagation mechanisms employed by reported threats and potential infections (figure 24). For instance, SMTP was used by 78 percent of reported malicious code threats that propagate; however, based on potential infections, only 47 percent of the top 50 samples that propagate did so through SMTP. This difference may be due to the higher likelihood of mass-mailing worms being reported, which is because of their highly prolific nature. Since a single mass-mailing worm can generate a high volume of messages, it is likely to cause a larger number of reports than actual infections, as was discussed previously in the "Malicious code types" section.

The P2P propagation vector has been steadily climbing over the last 18 months. Malicious code reports using this vector to propagate rose from 23 percent of all propagating malicious code in the first six months of 2006 to 29 percent in the last half of the year. Based on potential infections, P2P was used by 47 percent of malicious code that propagated during the second half of 2006.

The use of P2P as a propagation mechanism is likely to continue to grow in the foreseeable future. P2P networks are effective mechanisms for propagation since there are an immense number of files—possibly including pirated software and programs to bypass copy protection on software—present at any time and because they are largely unregulated. There is little, if any, security in place between computers on a P2P network. Furthermore, social engineering attacks are easy to carry out through P2P. Attackers can simply give a malicious code sample the same name as a popular download and make it available over P2P. Many users will inherently trust the malicious file and download it.

Enterprises should take measures to prevent P2P clients from being installed on any computers on their networks. They should also block any ports used by these applications at the network boundary. End users who download files from P2P networks should scan all such files with a regularly updated antivirus product.

Malicious code that exploits vulnerabilities

The exploitation of vulnerabilities as a means of malicious code propagation is an ongoing concern for enterprises. This section of the *Internet Security Threat Report* will examine the relationship between vulnerabilities and malicious code by assessing the proportion of malicious code that exploits vulnerabilities.

During the second half of 2006, 23 percent of the 1,318 documented malicious code instances exploited vulnerabilities (figure 25).[83] This is higher than the 17 percent proportion of the 1,249 malicious code instances documented in the first half of 2006. In the second half of 2005, 22 percent of the 1,077 documented malicious code instances exploited vulnerabilities.

While the majority of malicious code uses vectors other than vulnerabilities as a means to spread, the proportion that does employ vulnerabilities is significant. During the current reporting period, there have been a number of noteworthy malicious code events that do so, such as the use of client-side Microsoft Office vulnerabilities by Trojans. One specific instance of this was Trojan.PPDropper.F.[84] This Trojan infected targeted computers by exploiting the Microsoft PowerPoint Unspecified Remote Code Execution Vulnerability.[85] The Ginwui back door also exploited a Microsoft Word zero-day vulnerability.[86]

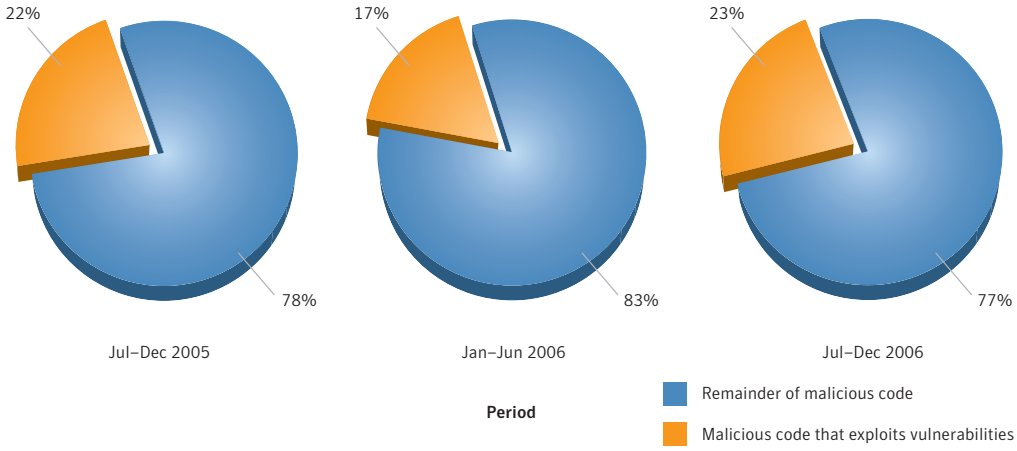

Figure 25. Malicious code that exploits vulnerabilities
Source: Symantec Corporation

In the second half of 2006, five zero-day exploits were released for vulnerabilities in Microsoft Office. This accounts for a significant proportion of malicious code that exploits vulnerabilities during the second half of 2006. Zero-day vulnerabilities present attackers with an opportunity to evade detection when compromising computers. In the context of malicious code, this will also increase the success rate when compromising computers, as the malicious code will appear to spread through an unknown vector until it has been discovered, analyzed, and mitigated by security and antivirus vendors.

[83] It should be noted that the number of documented malicious code instances differs from the number of malicious code submissions. Documented malicious code instances are those that have been analyzed and documented within the Symantec malicious code database.
[84] http://www.symantec.com/security_response/writeup.jsp?docid=2006-092715-1534-99
[85] http://www.securityfocus.com/bid/20226/info
[86] http://www.symantec.com/security_response/writeup.jsp?docid=2006-051914-5151-99

Phishing, Spam, and Security Risks

Traditionally, the Symantec *Internet Security Threat Report* has broken security threats down into three general categories: attacks, vulnerabilities, and malicious code. However, as Internet-based services and applications have expanded and diversified, the potential for computer programs to introduce other types of security risks has increased. The emergence of new risks, particularly spam, phishing, spyware, adware, and misleading applications, has necessitated an expansion of the traditional security taxonomy.

Symantec has monitored these new concerns as they have developed. This section will examine developments in these risks over the last six months of 2006. In particular, it will consist of three subsections, which will discuss:

- Phishing
- Spam
- Security risks, particularly adware, spyware, and misleading applications

Phishing

Phishing is an attempt by a third party to solicit confidential information from an individual, group, or organization, usually for financial gain. Phishers are groups or individuals who attempt to trick users into disclosing personal data, such as credit card numbers, online banking credentials, and other sensitive information. They may then use the information to commit fraudulent acts. This section of the *Internet Security Threat Report* will discuss phishing activity that Symantec detected between July 1 and December 31, 2006.

The data provided in this section is based on statistics derived from the Symantec Probe Network, which consists of over two million decoy email accounts that attract email messages from 20 different countries around the world. The main purpose of the network is to attract spam, phishing, viruses, and other email-borne threats. It encompasses more than 600 participating enterprises around the world, attracting email that is representative of traffic that would be received by over 250 million mailboxes. The Probe Network consists of previously used email addresses as well as email accounts that have been generated solely to be used as probes.

In addition to the Probe Network, Symantec also gathers phishing information through the Symantec Phish Report Network, an extensive antifraud community of enterprises and consumers. Members of the Phish Report Network contribute and receive fraudulent Web site addresses for alerting and filtering across a broad range of solutions.

Phishing is assessed according to two indicators: phishing messages and phishing attempts. A phishing message is a single, unique message that is sent to targets with the intent of gaining confidential and/or personal information from computer users. Each phishing message has different content and each one will represent a different way of trying to fool a user into disclosing information. A phishing message can be considered the “lure” with which a phisher attempts to entice a phishing target to disclose confidential information.

A single phishing message can be used in numerous distinct phishing attempts, usually targeting different end users. A phishing attempt can be defined as an instance of a phishing message being sent to a single user. Extending the fishing analogy, a phishing attempt can be considered a single cast of the lure (the phishing message) to try to ensnare a target.

Phishing Highlights

The following section will offer a brief summary of some of the phishing trends that Symantec observed during this period based on data provided by the sources listed above. Following this overview, the *Internet Security Threat Report* will discuss selected metrics in greater depth, providing analysis and discussion of the trends indicated by the data.

- The Symantec Probe Network detected a total of 166,248 unique phishing messages, a six percent increase over the first six months of 2006. This equates to an average of 904 unique phishing messages per day for the second half of 2006.
- Symantec blocked over 1.5 billion phishing messages, an increase of 19 percent over the first half of 2006. This means that Symantec blocked an average of 8.48 million phishing emails per day over the last six months of 2006.
- Throughout 2006, Symantec detected an average of 27 percent fewer unique phishing messages on weekends than the weekday average of 961.
- On weekends, the number of blocked phishing attempts was seven percent lower than the weekday average of 7,958,323 attempts per day.
- Organizations in the financial services sector accounted for 84 percent of the unique brands that were phished during this period.
- Forty-six percent of all known phishing Web sites were located in the United States, a much higher proportion than in any other country.

Phishing Discussion

This section will discuss selected phishing metrics in greater depth, providing analysis and discussion of the trends indicated by the data. The following metrics will be discussed:

- Daily and seasonal variations in phishing activity
- Phishing activity by sector
- Top countries hosting phishing Web sites
- Phishing—prevention and mitigation

Daily and seasonal variations in phishing activity

For the first time, in this volume of the *Internet Security Threat Report*, Symantec is analyzing the effects that the day of the week and seasonal events may have on phishing activity. To that end, Symantec calculated the average number of blocked phishing attempts and unique phishing messages on each day of the week for the year 2006. On average, Symantec detected 961 unique phishing messages each weekday (Monday through Friday) (figure 26). Throughout 2006, Symantec detected an average of

27 percent fewer unique phishing messages on weekends than the weekday average of 961. On weekends, the number of blocked phishing attempts was, on average, seven percent lower than the average of 7,958,323 attempts per weekday (figure 27).

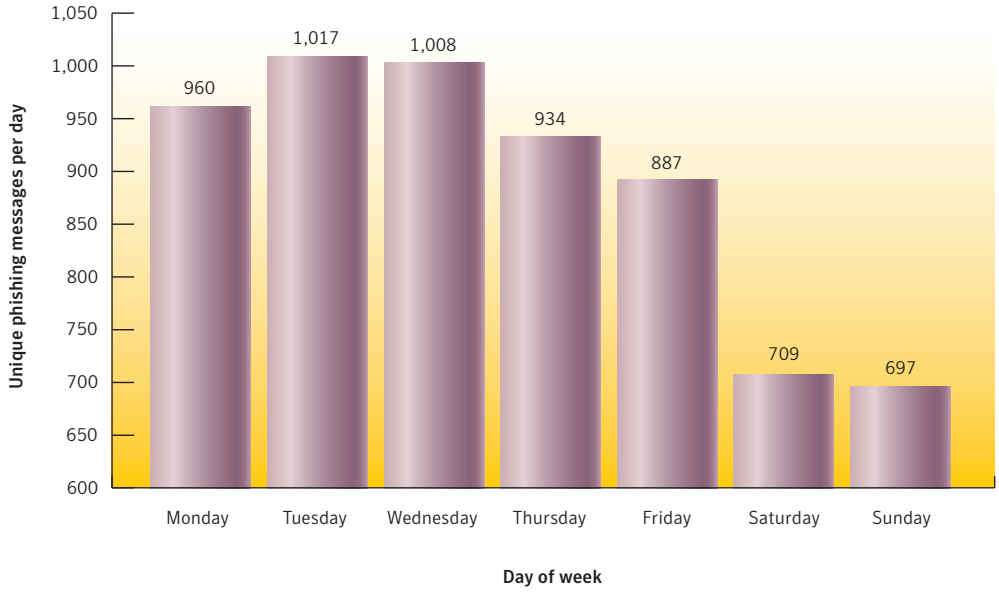

Figure 26. Unique phishing messages per day
Source: Symantec Corporation

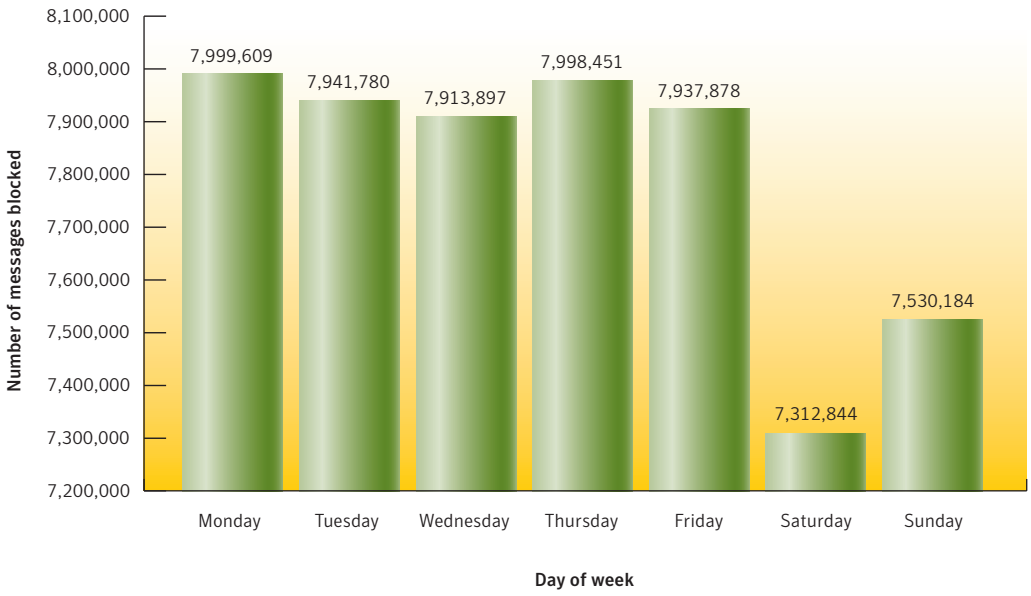

Figure 27. Blocked phishing messages per day
Source: Symantec Corporation

The decline in phishing activity on the weekends indicates that phishing activity mirrors the business week. Since legitimate companies are more likely to send email on weekdays, it may follow that attackers are emulating this pattern in their attacks. However, this pattern may also be due to the fact that phishing campaigns are generally short lived and, therefore, are most effective when people receive and read the phishing emails soon after they were sent, which may not be the case on weekends. Additionally, this could be a sign that phishing is beginning to follow a business model in which attackers work Monday through Friday.

Big events or holidays like Christmas and New Year appear to increase the amount of phishing activity. Attackers may find it easier to craft social engineering attacks around themes surrounding special events such as these. During the Christmas season of 2006, the number of blocked phishing messages climbed to a high of 29 percent above the average.

In 2006, Symantec observed a clear increase in blocked phishing attempts around the week of the Super Bowl final, which took place on February 5. Symantec blocked 33 percent more phishing messages during this period than on average (figure 28). Furthermore, during the week of the FIFA World Cup final (July 7, 2006) blocked phishing attempts were 40 percent higher than the average after having already been higher than normal for the beginning of the competition.

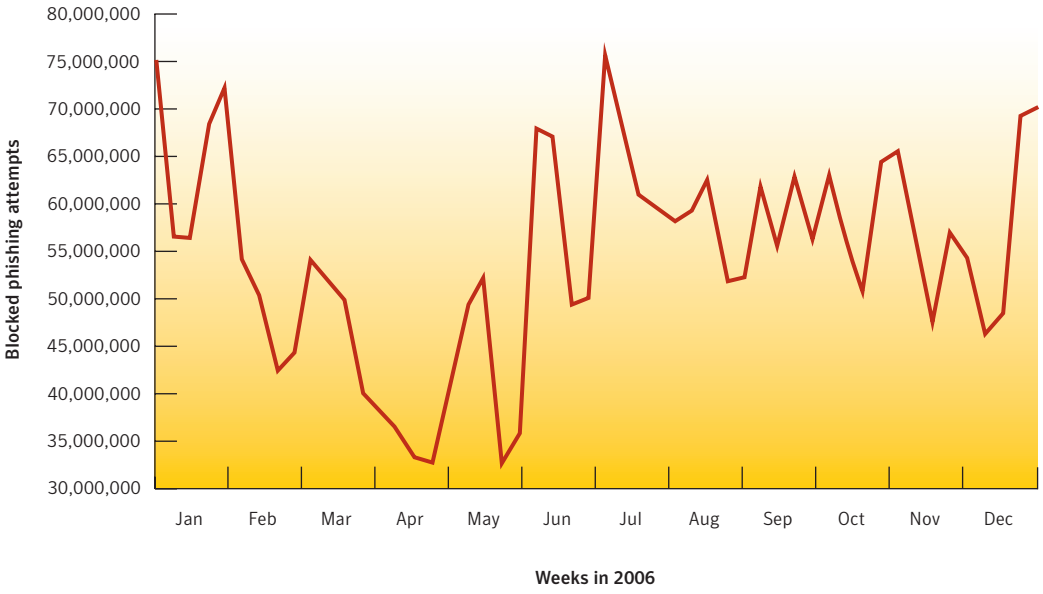

Figure 28. Blocked phishing messages per week
Source: Symantec Corporation

It should be noted that, in addition to the day of the week and the season, the amount of phishing attempts may be influenced by other factors, such as the release of new security products, the disclosure of vulnerabilities, and the availability of patches. Furthermore, the number of existing bot-infected computers could influence the number of phishing attacks, as they are often used to send phishing messages. If a number of bot-infected computers are disinfected or removed from the Internet at one time, the number of phishing attacks would likely drop.

Companies whose brands are frequently targeted by phishing attacks—that is, the company whose brand is spoofed in a phishing attack—should be aware of seasonal influences on phishing campaigns in order to counter them before they occur. Such organizations may want to warn customers of potential phishing activity prior to seasons or events that could be associated with increased phishing activity.

Phishing activity by sector

In the previous edition of the *Internet Security Threat Report*, Symantec began tracking the sectors of companies whose brands were being used in phishing attacks. Since that report, the Symantec Phish Report Network has grown substantially, which has had an effect upon the overall volume of phishing Web sites[87] tracked by Symantec this period. As a result, a broader range of phished industries was reported this period. This metric is important for enterprises because the use of an organization's brand can undermine consumer confidence and damage the organization's reputation. Furthermore, the company may be required to compensate victims of any phishing scams that use the company's brand.

Most of the unique brands phished in the last six months of 2006 were in the financial services sector. Organizations in that sector accounted for 84 percent of the brands that were phished during this period (figure 29). This is not surprising, as most phishing attacks are motivated by profit. A successful phishing attack that mimics the brand of a financial entity is most likely to yield data that could be used for financial gain.

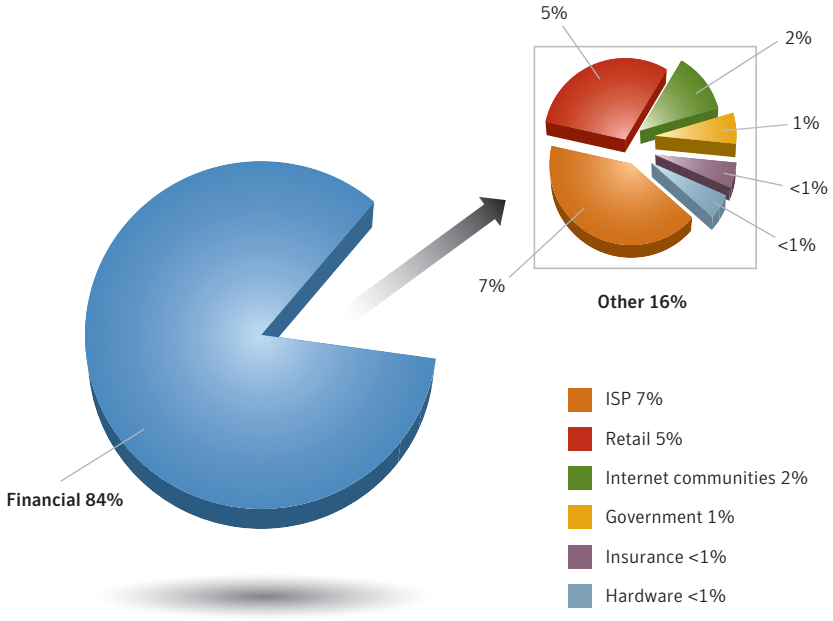

Figure 29. Brands used in phishing attacks by sector
Source: Symantec Corporation

[87] A phishing Web site is the site set up by the attacker to capture a victim's authentication information. In many cases, these sites are designed to mimic the actual site of the brand being spoofed.

Organizations in the Internet service provider (ISP) sector made up seven percent of the unique brands phished this period. ISP accounts can be valuable targets for phishers. While users may not think that their email accounts contain information that is of value to attackers, this may be misguided. In many cases, people reuse the same authentication credentials (such as usernames and passwords) for multiple accounts, including the email accounts. Additionally, most online banking and brokerage accounts have a utility to reset forgotten passwords for a user. If an attacker gains access to the email account used for this, they can submit a password reset request to the site in question and easily gain access to crucial accounts. Finally, free Web-hosting space that is often provided with these accounts can also be used to host phishing Web sites.

While the financial sector accounted for 84 percent of the unique brands being phished in the second half of 2006, it only made up 64 percent of the total phishing Web sites reported to Symantec (figure 30). Conversely, the retail sector accounted for only five percent of the unique brands phished, but 34 percent of the volume of phishing Web sites. The ISP sector accounted for approximately two percent of phishing Web sites, while the remaining sectors accounted for less than one percent.

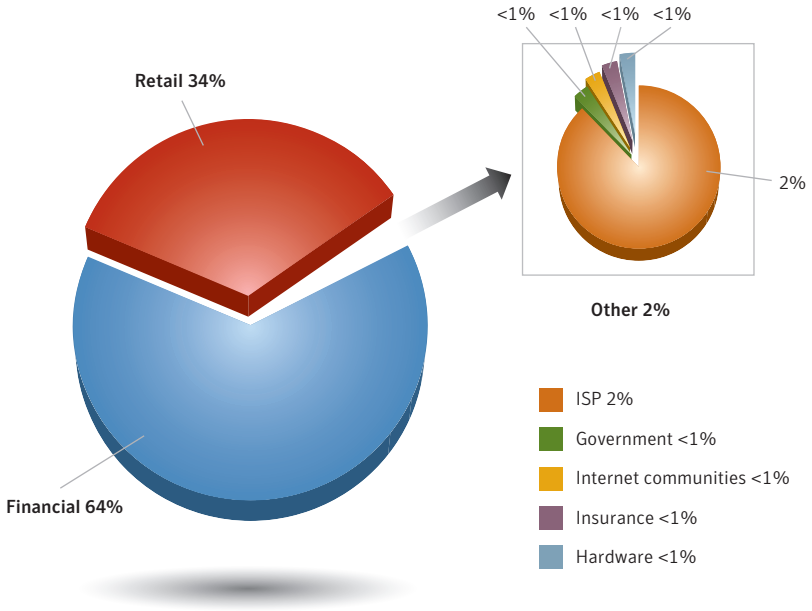

Figure 30. Sectors of unique brands being phished by volume
Source: Symantec Corporation

The high volume of phishing Web sites reported for a relatively small number of retail sector brands indicates that attackers are concentrating a large number of phishing attacks against a small number of brands. This suggests that attackers may feel that only a few retail brands are significant enough to provide an economic return or that they have experienced enough success attacking these few brands that they do not need to attack other retail brands. Some attackers may also be using ready-made phishing "kits" that are likely to focus on a small number of retail brands.[88]

[88] A phishing kit is a set of tools that an attacker can use to easily construct phishing email messages and Web sites based on a template.

If the revenue generated by phishing these brands diminishes, it is likely that attackers will move their efforts to other sectors and that the volume of phishing attacks in the retail sector will subsequently drop. On the other hand, the number of retail brands being phished could increase as attackers are forced to expand their efforts. This pattern has already been seen in the financial sector when some of the larger institutions began taking aggressive steps, such as two-factor authentication, to protect their brand and customers from phishing attacks.[89] As a result of this, phishers have begun to attack smaller brands, although larger brands are still used in the majority of phishing attacks.

Top countries hosting phishing Web sites

For the first time, in this edition of the *Internet Security Threat Report*, Symantec is assessing the countries in which the most phishing Web sites are hosted. This data is a snapshot in time, and does not offer insight into changes in the locations of certain phishing sites over the course of the reporting period. It should also be noted that the fact that a phishing site is hosted in a certain country does not necessarily mean that the attacker is located in that country.

In the second half of 2006, 46 percent of all known phishing sites were located in the United States (figure 31), a much higher proportion than in any other country. This is likely because a large number of Web-hosting providers—particularly free Web hosts— are located in the United States. Furthermore, the United States has the highest number of Internet users in the world and is home to a large number of Internet-connected organizations, both large and small.

A Web server belonging to a small company makes an ideal platform for phishers to use as a host. In many cases, these servers do not have full-time administrative or security staff. As a result, the security patch level of these computers may not be up to date, and other security measures may not have been fully implemented. An attacker could thus compromise the computer with less chance of the illicit Web site being discovered right away. Since the compromised computer already hosts a Web site, browser traffic destined to it will likely escape notice.

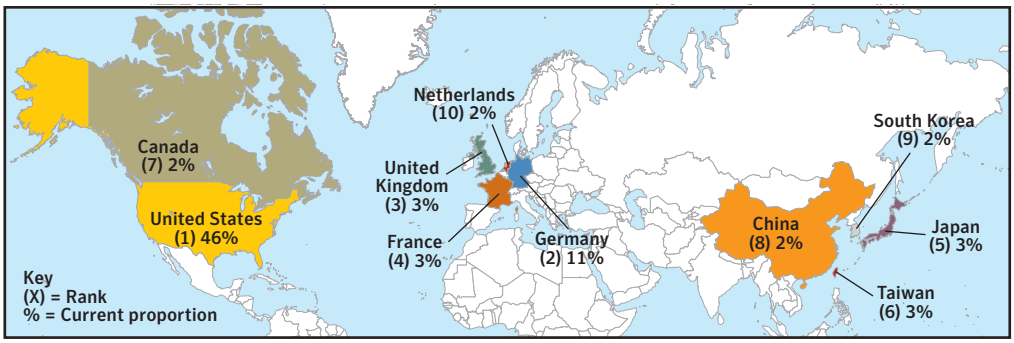

Figure 31. Top countries hosting phishing Web sites
Source: Symantec Corporation

[89] Two-factor authentication consists of using a password or PIN number generated by the user plus a physical device such as a one-time password list or physical token that generates random numbers.

Germany was home to the second highest percentage of phishing Web sites in the second half of 2006, with 11 percent of the worldwide total. It also has the largest number of Web-hosting providers in Europe.[90] By hosting their sites with large providers, phishers may be able to gain the advantage of obscurity. With so many sites hosted by a single provider, it may take days for the provider to discover a phishing site and shut it down.

Phishing—prevention and mitigation

Symantec recommends that enterprise users protect themselves against phishing threats by filtering email at the server level through the mail transfer agent (MTA). Although this will likely remain the primary point of filtering for phishing, organizations can also use IP-based filtering upstream, as well as HTTP filtering.

DNS block lists also offer protection against potential phishing emails.[91] Organizations could also consider using domain-level or email authentication in order to verify the actual origin of an email message. This can protect against phishers who are spoofing mail domains.[92]

To protect against potential phishing activity, administrators should always follow Symantec best practices as outlined in Appendix A of this report. Symantec also recommends that organizations educate their end users about phishing.[93] They should also keep their employees notified of the latest phishing attacks and how to avoid falling victim to them.[94]

Organizations can also employ Web server log monitoring to track if and when complete downloads of their Web sites are occurring. Such activity may indicate that someone is using the legitimate Web site to create an illegitimate Web site that could be used for phishing.

Organizations can detect phishing attacks that use spoofing by monitoring non-deliverable email addresses or bounced email that is returned to non-existent users. They should also monitor the purchasing of cousin domain names by other entities to identify purchases that could be used to spoof their corporate domains.[95] This can be done with the help of companies that specialize in domain monitoring; some registrars even provide this service.[96]

The use of antiphishing toolbars and components in Web browsers can also help protect users from phishing attacks. These measures notify the user if a Web page being visited does not appear to be legitimate. This way, even if a phishing email reaches a user's inbox, the user can still be alerted to the potential threat.

End users should follow best security practices, as outlined in Appendix A of this report. They should deploy an antiphishing solution. As some phishing attacks may use spyware and/or keystroke loggers, Symantec advises end users to use antivirus software, antispam software, firewalls, toolbar blockers, and other software detection methods. Symantec also advises end users to never disclose any confidential personal or financial information unless and until they can confirm that any request for such information is legitimate.

Users should review bank, credit card, and credit information frequently. This can provide information on any irregular activities. For further information, the Internet Fraud Complaint Center (IFCC) has also released a set of guidelines on how to avoid Internet-related scams.[97] Additionally, network administrators can review Web proxy logs to determine if any users have visited known phishing sites.

[90] http://www.webhosting.info/webhosts/tophosts/global
[91] A DNS block list (sometimes referred to as a black list) is simply a list of IP addresses that are known to send unwanted email traffic. It is used by email software to either allow or reject email coming from IP addresses on the list.
[92] Spoofing refers to instances where phishers forge the "From:" line of an email message using the domain of the entity they are targeting with the phishing attempt.
[93] For instance the United States Federal Trade Commission has published some basic guidelines on how to avoid phishing. They are available at: http://www.ftc.gov/bcp/edu/pubs/consumer/alerts/alt127.htm
[94] A good resource for information on the latest phishing threats can be found at: http://www.antiphishing.org
[95] "Cousin domains" refers to domain names that include some of the key words of an organization's domain or brand name; for example, for the corporate domain "bigbank.com" cousin domains could include "bigbank-alerts.com", "big-bank-security.com", and so on.
[96] See http://markmonitor.com/brandmanagement/index.html for instance.
[97] http://www.fbi.gov/majcases/fraud/internetschemes.htm

Spam

Spam is usually defined as junk or unsolicited email sent by a third party. While it is certainly an annoyance to users and administrators, spam is also a serious security concern as it can be used to deliver Trojans, viruses, and phishing attempts. It could also cause a loss of service or degradation in the performance of network resources and email gateways. This section of the *Internet Security Threat Report* will discuss developments in spam activity between July 1 and December 31, 2006.

The data used in this analysis is based on data returned from the Symantec Probe Network as well as data gathered from a statistical sampling of the Symantec Brightmail AntiSpam customer base. Specifically, statistics are gathered from enterprise customers' Symantec Brightmail AntiSpam servers that receive more than 1,000 email messages per day. This removes the smaller data samples (that is, smaller customers and test servers), thereby allowing for a more accurate representation of data.

The Symantec Probe Network consists of millions of decoy email addresses that are configured to attract a large stream of spam attacks. An attack can consist of one or more messages. The goal of the Probe Network is to simulate a wide variety of Internet email users, thereby attracting a stream of traffic that is representative of spam activity across the Internet as a whole. For this reason, the Probe Network is continuously optimized in order to attract new varieties of spam attacks. This is accomplished through internal production changes that are made to the network, which thus affect the number of new spam attacks it receives as a whole.

Spam Highlights

The following section will offer a brief summary of some of the spam trends that Symantec observed during this period based on data provided by the sources listed above. Following this overview, the *Internet Security Threat Report* will discuss selected metrics in greater depth, providing analysis and discussion of the trends indicated by the data.

- Between July 1 and December 31, 2006, spam made up 59 percent of all monitored email traffic. This is an increase over the first six months of 2006 when 54 percent of email was classified as spam.
- Sixty-five percent of all spam detected during this period was composed in English.
- In the last six months of 2006, 0.68 percent of all spam email contained malicious code. This means that one out of every 147 spam messages blocked by Symantec Brightmail AntiSpam contained malicious code.
- Spam related to financial services made up 30 percent of all spam during this period, the most of any category.
- During the last six months of 2006, 44 percent of all spam detected worldwide originated in the United States.
- The United States hosted the largest proportion of spam zombies, with 10 percent of the worldwide total.

Spam Discussion

This section will discuss selected spam metrics in greater depth, providing analysis and discussion of the trends indicated by the data. The following metrics will be discussed:

- Top spam categories
- Top countries of spam origin
- Distribution of spam zombies

Top spam categories

Spam categories are assigned by Symantec Email Security Group analysts based on spam activity that is detected by the Symantec Probe Network. While some of the categories may overlap, this data provides a general overview of the types of spam that are most commonly seen on the Internet today.

It is important to note that this data is restricted to spam attacks that are detected and processed by the Symantec Probe Network. Internal upstream processing may weed out particular spam attacks, such as those that are determined to be potential fraud attacks.

The most common type of spam detected in the latter half of 2006 was related to financial services (figure 32), which made up 30 percent of all spam during this period. Spam related to health services and products made up 23 percent of all spam, and spam related to commercial products made up 21 percent of the total during this period.

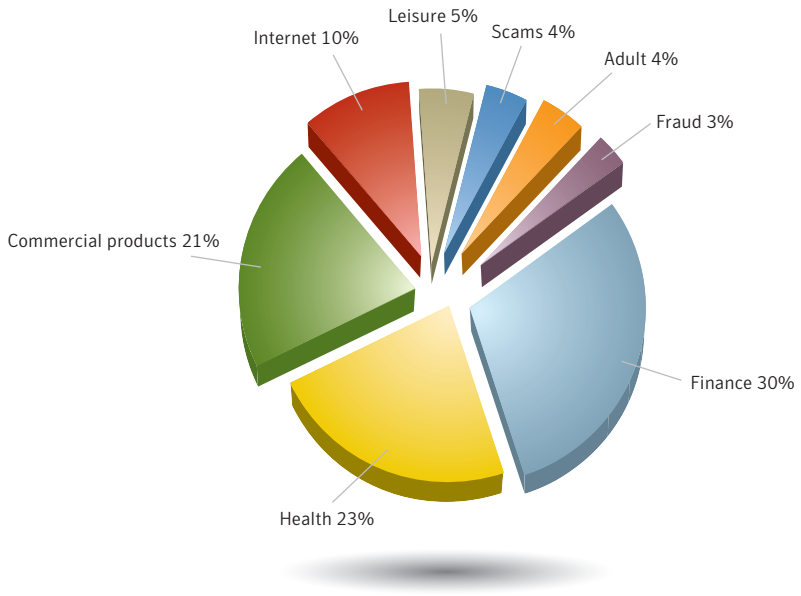

Figure 32. Spam categories
Source: Symantec Corporation

The rise in financially-related spam in the second half of 2006 was due mainly to a noticeable increase in stock market "pump and dump" spam. Pump and dump is the name given to schemes in which criminals profit by creating an artificial interest in a stock they own. They buy a penny stock when the price is low. They then artificially pump up demand for the stock by sending out spam that appears to be from a respected stock advisor, but that actually contains false predictions of high performance for the stock. Recipients of the message, trusting the spam content, buy the stock, creating demand for it and thereby raising the price. When the prices are high, the perpetrators of the scheme sell their stock for a profit.[98]

This type of spam has been proven to allow the spammers to generate revenue directly and almost immediately.[99] This factor alone is likely to make stock market spam more appealing than other types of spam.

The increase in financial services spam was almost in direct proportion to the decrease in adult-related spam this period. In the previous edition of the *Internet Security Threat Report*, adult-related spam accounted for 22 percent of spam on the Internet. However, it dropped sharply over the last half of 2006, making up only four percent of all spam during this period. This is likely because a large portion of profit-driven spammers shifted their efforts towards the more lucrative stock market spam.

The second most common type of spam detected in the last six months of 2006 was related to health services and products, which accounted for 23 percent of all spam. It is not surprising that health-related spam makes up such a high proportion of the total. This category traditionally has one of the highest "click-through" rates, as it tends to be more difficult to market through more legitimate and traditional means. A click-through is a link that is embedded in a spam message. The link contains uniquely identifiable information about its originator. Each time a user clicks on the link, it is considered a click-through. Typically, the originator receives financial compensation for each click-through. Spammers have an economic incentive to have a high click-through rate in order to increase their return on investment. Therefore, it is reasonable to conclude that they would use spam content that has a high click-through rate.

Top countries of spam origin

This section will discuss the top ten countries of spam origin. The nature of spam and its distribution on the Internet presents challenges in identifying the location of people who are sending spam. Many spammers try to redirect attention away from their actual geographic location. In an attempt to bypass DNS block lists, they build coordinated networks of compromised computers known as bot networks, which allow them to send spam from sites that are distant from their physical location. In doing so, they will likely focus on compromised computers in those regions with the largest bandwidth capabilities (for a more in-depth discussion of this, please refer to the "Attack Trends" report of this report). Following this logic, the region from which the spam originates may not correspond with the region in which the spammers are located.

[98] http://www.sec.gov/answers/pumpdump.htm
[99] http://papers.ssrn.com/sol3/papers.cfm?abstract_id=920553

During the last six months of 2006, 44 percent of all spam detected worldwide originated in the United States (table 6), more than in any other country during this period. This is likely due to the high number of broadband users in that country and the high percentage of bot-infected computers located there, as was discussed in the "Attack Trends" section of this report. The United States was also the top country of spam origin in the first half of 2006, when 49 percent of spam originated there.

The second highest source of spam this period was a group of undetermined European Union countries. In this group, the specific source countries cannot be definitively identified because the ISPs through whose networks the spam was sent operate in more than one EU country.

China was the third highest country of spam origin in the second half of 2006. Six percent of spam detected by Symantec during this period originated there, compared to 11 percent in the first half of the year. Symantec believes the drop since last period may have occurred because some companies that do not do business in China automatically block all email originating there.

Country	Jul–Dec 2006	Jan–Jun 2006
United States	44%	49%
Undetermined EU Countries	7%	4%
China	6%	11%
Canada	4%	5%
United Kingdom	3%	4%
South Korea	3%	5%
Japan	3%	2%
France	3%	2%
Spain	3%	2%
Poland	3%	2%

Table 6. Top ten countries of spam origin
Source: Symantec Corporation

Distribution of spam zombies

For the first time, in this volume of the *Internet Security Threat Report*, Symantec is tracking countries that host spam zombies. A spam zombie is a computer infected with a bot or some other malicious code that allows email messages to be relayed through it.

It should be noted that the data on which the "Countries of spam origin" discussion was based includes spam messages that may also be sent from legitimate email servers as well as those that were sent through spam zombies. Since spam zombies are the result of an infection by a bot, worm, or Trojan, there is a wider distribution among the top countries for spam zombies than is evident in the "Countries of spam origin" discussion.

Between July 1 and December 31, 2006, the United States hosted the largest proportion of spam zombies, but only by a small margin. Ten percent of spam zombies were located there (figure 33). This is drastically different from the top countries of spam origin, in which the United States accounted for nearly half the total volume. During this period, the United States was one of the top reporting countries of bots such as Spybot and Mytob, both of which can be used to send spam.

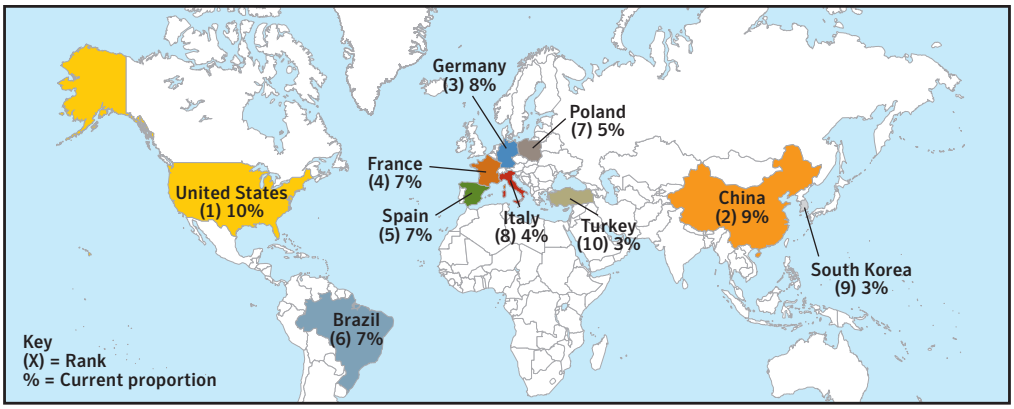

Figure 33. Distribution of spam zombies
Source: Symantec Corporation

China and Germany were the second and third highest countries for spam zombies, hosting nine and eight percent respectively. France, Spain, and Brazil followed closely at seven percent each. The small variance between the top countries hosting spam zombies is quite different from the distribution of bots during this period (as was discussed in the "Bot-infected computers by country" metric in the "Attack Trends" section of this report). This indicates that not all spam zombies are necessarily bots and that not all bots are used to send spam.

Security Risks

Symantec uses the term "security risk" to describe adware, spyware, misleading applications, and other programs that users may not want on their system. While these risks are not categorized as malicious code, Symantec monitors them with many of the same methods employed in tracking malicious code. This involves an ongoing analysis of reports and data delivered from over 120 million client server and gateway email systems deploying Symantec antivirus security solutions, as well as filtration of 25 million email messages per day by Symantec Brightmail AntiSpam antifraud filters. Symantec then compiles the most common reports and analyzes them to determine the appropriate categorization. Steps for the protection against and mitigation of these security risks are presented at the end of the "Security Risks" section.

Security Risk Highlights

The following section will offer a brief summary of some of the security risks that Symantec observed during this period based on data provided by the sources listed above. Following this overview, the *Internet Security Threat Report* will discuss selected metrics in greater depth, providing analysis and discussion of the trends indicated by the data.

- The most commonly reported security risk was an adware program named ZangoSearch.
- All of the top ten security risks reported in the last six months of 2006 employ at least one anti-removal technique compared to only five of the top ten security risks in the last reporting period.
- All of the top ten security risks reported during this period employ self-updating.
- Potentially unwanted applications accounted for 41 percent of reports in the top ten new security risks in the second half of 2006.
- Misleading application detections increased by 40 percent in the second half of 2006.

Security Risks Discussion

This section will discuss selected security risks metrics in greater depth, providing analysis and discussion of the trends indicated by the data. The following metrics will be discussed:

- Top ten reported security risks
- Top ten new security risks

Top ten reported security risks

Between July 1 and December 31, 2006, the most commonly reported security risk was ZangoSearch (table 7).[100] This is an adware program that accounted for 13 percent of the top ten reported security risks. A new entry in the top ten, ZangoSearch monitors the contents of Web browser windows. When certain keywords are detected in Internet search or shopping browser windows, ZangoSearch opens the Web sites of companies whose products ZangoSearch has agreed to promote.

It has been reported that ZangoSearch uses questionable methods to install itself on users' computers.[101] By clicking on misleading video links on certain MySpace sites, users would inadvertently visit a fake YouTube site, which would then download a Zango Cash toolbar on the unsuspecting user's computer.

Earlier in 2006, Zango merged with another adware toolbar provider named Hotbar. Hotbar was the most prevalent security risk in the first six months of 2006, accounting for 24 percent of the top ten security risk reports during that period. In the current reporting period, however, Hotbar was the fifth most common security risk, accounting for 11 percent of the top ten submissions. This decrease may be the result of the merge and the promotion of one unified product.

[100] http://www.symantec.com/security_response/writeup.jsp?docid=2005-050416-3519-99&tabid=1
[101] http://www.eweek.com/article2/0,1895,2052998,00.asp

Rank	Risk Name	Risk Type
1	ZangoSearch	Adware
2	SpySheriff	Misleading application
3	Purityscan	Adware
4	Websearch	Adware
5	Hotbar	Adware
6	Borlan	Adware
7	QoolAid	Adware
8	Look2Me	Adware
9	NDotNet	Adware
10	PigSearch	Adware

Table 7. Top ten reported security risks
Source: Symantec Corporation

The second most common security risk in the second half of 2006 was SpySheriff,[102] a misleading application. A new entry to the top ten, SpySheriff was first discovered in December 2005 and accounted for 12 percent of security risks in the top ten during this period.

SpySheriff purportedly detects and removes programs such as keystroke loggers, Trojan horses, and password-stealing applications. Consumers can install SpySheriff from the company's Web site, but many consumers encounter SpySheriff through misleading banner advertisements, full-screen pop-up windows, and misleading Web sites. Once installed, the program reports on false security risks. To remove these false security risks the end user is asked to register the program and pay for its usage.

The third most commonly reported security risk over the last six months of 2006 was Purityscan,[103] another new entry into the top ten. The first variant of this adware was detected in September 2003. In the second half of 2006, it accounted for 12 percent of the submissions in the top ten. Purityscan is an adware program that downloads and displays advertisements on a computer. Once it is installed on a user's computer, it is also capable of downloading and installing programs automatically without user consent.

[102] http://www.symantec.com/security_response/writeup.jsp?docid=2005-122910-4625-99
[103] http://www.symantec.com/security_response/writeup.jsp?docid=2003-090516-2325-99

Top ten new security risks

Between July 1 and December 31, 2006, Symantec saw a slight drop in the detection of new security risks. This may be an indication that security risk developers are trying to create alternative sources of revenue.

Rank	Risk Name	Risk Type
1	Movieland	Potentially unwanted application
2	Searchnet	Adware
3	VirusBurst	Misleading application
4	Roogoo	Adware
5	Trustyhound	Spyware
6	Baigoo	Trackware
7	VirusBlast	Misleading application
8	2AntiSpyware	Spyware
9	DoctorAdwarePro	Adware
10	Netmedia	Adware

Table 8. Top ten new security risks
Source: Symantec Corporation

The top new security risk falls under the category of a potentially unwanted application (table 8). This is a new category that Symantec recently introduced to allow for the detection of applications that have an impact on security, privacy, resource consumption, or are associated with other security risks. Potentially unwanted applications accounted for 41 percent of reports in the top ten new security risks in the second half of 2006 (table 9).

The most common new security risk during this period was Movieland,[104] a potentially unwanted application that accounted for 41 percent of reports. Movieland is installed on a user's computer surreptitiously, causing unwanted pop-up advertisements to appear on the computer. It is very difficult to remove this application. In January 2007, Movieland entered into stipulated interim agreements with the U.S. Federal Trade Commission to provide disclosures relating to their practices.

SearchNet was the second most common new security risk during the second half of 2006,[105] accounting for 21 percent of reports. An adware program, SearchNet is a Browser Helper Object (BHO) that replaces the default search page in Internet Explorer.[106]

[104] http://www.symantec.com/smb/security_response/writeup.jsp?docid=2006-091511-1921-99
[105] http://www.symantec.com/security_response/writeup.jsp?docid=2006-071912-4640-99
[106] Browser helper objects (BHOs) are add on programs that can add legitimate features to a user's browser (IE 4.x and up). For example, document readers used to read programs within the browser do so via BHOs. BHOs can also be used to install security risks on a user's Web browser using ActiveX controls.

The third most common new security risk in the second half of 2006 was VirusBurst.[107] This is a misleading application that accounts for 16 percent of reports. VirusBurst gives exaggerated reports of threats on the computer. It is installed surreptitiously by a Trojan, which it later detects as a threat. The program then prompts the user to purchase a registered version of the software in order to remove the reported threats.

Risk Type	Percent of New Risks
Potentially unwanted applications	41%
Adware	35%
Misleading applications	18%
Dialers	0%
Security assessment tools	0%
Spyware	5%
Security risk	0%
Trackware	4%

Table 9. New security risks by category
Source: Symantec Corporation

Towards the end of 2006, Symantec was seeing a lot of new clones of misleading applications. A clone refers to the same basic program with a new name and graphical user interface (GUI). Clones are often used in an attempt to avoid antispyware detection as a misleading application. For example, a company named KlikSoftware.com appears to have been responsible for a number of clones of rogue antispyware.[108] Some of their programs, such as Remedy AntiSpy, Adware Bazooka, and HitVirus, are clones of a security risk known as Punisher.[109]

[107] http://www.symantec.com/security_response/writeup.jsp?docid=2006-090516-0204-99
[108] http://spywarewarrior.com/rogue_anti-spyware.htm
[109] http://www.symantec.com/enterprise/security_response/weblog/2006/06/rogue_antispyware_in_action.html

Appendix A—Symantec Best Practices

Enterprise Best Practices

1. Employ defense-in-depth strategies, which emphasize multiple, overlapping, and mutually supportive defensive systems to guard against single-point failures in any specific technology or protection method. This should include the deployment of regularly updated antivirus, firewalls, intrusion detection, and intrusion protection systems on client systems.
2. Turn off and remove services that are not needed.
3. If malicious code or some other threat exploits one or more network services, disable or block access to those services until a patch is applied.
4. Always keep patch levels up to date, especially on computers that host public services and are accessible through the firewall, such as HTTP, FTP, mail, and DNS services.
5. Consider implementing network compliance solutions that will help keep infected mobile users out of the network (and disinfect them before rejoining the network).
6. Enforce an effective password policy.
7. Configure mail servers to block or remove email that contains file attachments that are commonly used to spread viruses, such as .VBS, .BAT, .EXE, .PIF, and .SCR files.
8. Isolate infected computers quickly to prevent the risk of further infection within the organization. Perform a forensic analysis and restore the computers using trusted media.
9. Train employees to not open attachments unless they are expected and come from a known and trusted source, and to not execute software that is downloaded from the Internet unless it has been scanned for viruses.
10. Ensure that emergency response procedures are in place. This includes having a backup-and-restore solution in place in order to restore lost or compromised data in the event of successful attack or catastrophic data loss.
11. Educate management on security budgeting needs.
12. Test security to ensure that adequate controls are in place.
13. Be aware that security risks may be automatically installed on computers with the installation of file-sharing programs, free downloads, and freeware and shareware versions of software. Clicking on links and/or attachments in email messages (or IM messages) may also expose computers to unnecessary risks. Ensure that only applications approved by the organization are deployed on desktop computers.

Consumer Best Practices

1. Consumers should use an Internet security solution that combines antivirus, firewall, intrusion detection, and vulnerability management for maximum protection against malicious code and other threats.
2. Consumers should ensure that security patches are up to date and that they are applied to all vulnerable applications in a timely manner.
3. Consumers should ensure that passwords are a mix of letters and numbers, and should change them often. Passwords should not consist of words from the dictionary.
4. Consumers should never view, open, or execute any email attachment unless the attachment is expected and the purpose of the attachment is known.
5. Consumers should keep virus definitions updated regularly. By deploying the latest virus definitions, consumers can protect their computers against the latest viruses known to be spreading "in the wild."
6. Consumers should routinely check to see if their PC or Macintosh system is vulnerable to threats by using Symantec Security Check at www.symantec.com/securitycheck.
7. Consumers should deploy an antiphishing solution. They should never disclose any confidential personal or financial information unless and until they can confirm that any request for such information is legitimate.
8. Consumers can get involved in fighting cybercrime by tracking and reporting intruders. With Symantec Security Check's tracing service, users can quickly identify the location of potential hackers and forward the information to the attacker's ISP or local police.
9. Consumers should be aware that security risks may be automatically installed on computers with the installation of file-sharing programs, free downloads, and freeware and shareware versions of software. Clicking on links and/or attachments in email messages (or IM messages) may also expose computers to unnecessary risks. Ensure that only applications approved by the organization are deployed on desktop computers.
10. Some spyware and adware applications can be installed after an end user has accepted the end-user license agreement (EULA), or as a consequence of that acceptance. Consumers should read EULAs carefully and understand all terms before agreeing to them.
11. Consumers should beware of programs that flash ads in the user interface. Many spyware programs track how users respond to these ads, and their presence is a red flag. When users see ads in a program's user interface, they may be looking at a piece of spyware.

Appendix B—Attack Trends Methodology

Attack trends in this report are based on the analysis of data derived from the Symantec™ Global Intelligence Network, which includes the Symantec DeepSight™ Threat Management System, Symantec™ Managed Security Services, and the Symantec Honeypot Network. Symantec combines data derived from these sources for analysis.

Attack definitions

In order to avoid ambiguity with the findings presented in this discussion, Symantec's methodology for identifying various forms of attack activity is outlined clearly below. This methodology is applied consistently throughout our monitoring and analysis. The first step in analyzing attack activity is to define precisely what an attack is. Attacks are individual instances of malicious network activity. Attacks consist of one IDS or firewall alert that is indicative of a single attack action.

Explanation of research inquiries

This section will provide more detail on specific methodologies used to gather and analyze the data and statistics in this report. While most methodologies are adequately explained in the analysis section of the report, the following investigations warranted additional detail.

Targeted Web browsers

Symantec identifies attacks that are detected being carried out against Web browsers across the Symantec™ Global Intelligence Network, assesses which of these attacks target Web browsers, and determines which specific Web browser(s) is targeted by the attack. The distribution of targeted Web browsers is derived by determining what proportion of the source IP addresses of Web browser attacks is targeting each of the specific Web browsers.

Denial of service attacks

Although there are numerous methods for carrying out denial of service (DoS) attacks, Symantec derives this metric by measuring DoS attacks that are carried out by flooding a target with SYN requests. These are often referred to as SYN flood attacks. This type of attack works by overwhelming a target with SYN requests and not completing the initial request, which thus prevents other valid requests from being processed.

In many cases, SYN requests with forged IP addresses are sent to a target, allowing a single attacking computer to initiate multiple connections, resulting in unsolicited traffic, known as backscatter, being sent to other computers on the Internet. This backscatter is used to derive the number of DoS attacks observed throughout the reporting period. Although the values Symantec derives from this metric will not identify all DoS attacks carried out, it will highlight DoS attack trends.

To determine the countries targeted by DoS attacks, Symantec cross-references the target IP addresses of every attack with several third-party, subscription-based databases that link the geographic location of systems to source IP addresses. While these databases are generally reliable, there is a small margin of error.

Sectors targeted by DoS attacks were identified using the same methodology as targeted countries. However, in this case, attackers who were considered were those carrying out a set of DoS attacks that were detected by IDS and IPS software.

Identity theft data breaches

Symantec identifies the proportional distribution of cause and sector for data breaches that may facilitate identity theft based on data provided by the Privacy Rights Clearinghouse,[110] who in turn derived the data from Attrition.org.[111] The sector that experienced the loss along with the cause of loss that occurred is determined through analysis of the organization reporting the loss and the method that facilitated the loss.

Bot-infected computers

Symantec identifies bots based on coordinated scanning and attack behavior observed in network traffic. For an attacking computer to be considered to be participating in this coordinated scanning and attacking, it must fit into that pattern to the exclusion of any other activity. This behavioral matching will not catch every bot network computer, and may identify other malicious code or individual attackers behaving in a similarly coordinated way as a bot network. This behavioral matching will, however, identify many of the most coordinated and aggressive bot-infected computers and will give insight into the population trends of bot network computers, including those that are considered to be actively working in a well coordinated and aggressive fashion at some point in time during the reporting period.

This metric explores the number of active bot-infected computers that the Symantec™ Global Intelligence Network has detected and identified during the last six months of 2006. Identification is carried out on an individual basis by analyzing attack and scanning patterns. Computers generating attack patterns that show a high degree of coordination are considered to be bot-infected computers.

[110] http://www.privacyrights.org
[111] http://www.attrition.org

As a consequence of this, Symantec does not identify all bot-infected computers, but only those that are actively working in a well coordinated and aggressive fashion. Given Symantec's extensive and globally distributed sensor base, it is reasonable to assume that the bot activities discussed here are representative of worldwide bot trends, and can thus provide an understanding of current bot activity across the Internet as a whole.

Bot-infected computers by countries and cities

This metric is based on the same data as the "Active bot-infected computers" discussion of the "Attacks Trends" section of the report. Symantec cross-references the IP addresses of every identified bot-infected computer with several third-party subscription-based databases that link the geographic location of systems to IP addresses. While these databases are generally reliable, there is a small margin of error. Only cities that can be determined with a confidence rating of at least four out of five are included for consideration. The data produced is then used to determine the global distribution of bot-infected computers.

Top originating countries

Symantec identifies the national sources of attacks by automatically cross-referencing source IP addresses of every attacking IP with several third-party, subscription-based databases that link the geographic location of systems to source IP addresses. While these databases are generally reliable, there is a small margin of error.

Top targeted sectors

For the purposes of the *Internet Security Threat Report*, a targeted attacker is defined as one that is detected attacking at least three users or organizations in a specific sector, to the exclusion of all other sectors. The targeted sector attack rate is a measure of the percentage of all attackers that target only organizations or users in a specific sector and is represented as a proportion of all targeted attacks. Figure 34 represents the proportional sensor distribution for each sector. Sectors with less than ten sensors have been excluded from the resulting totals.

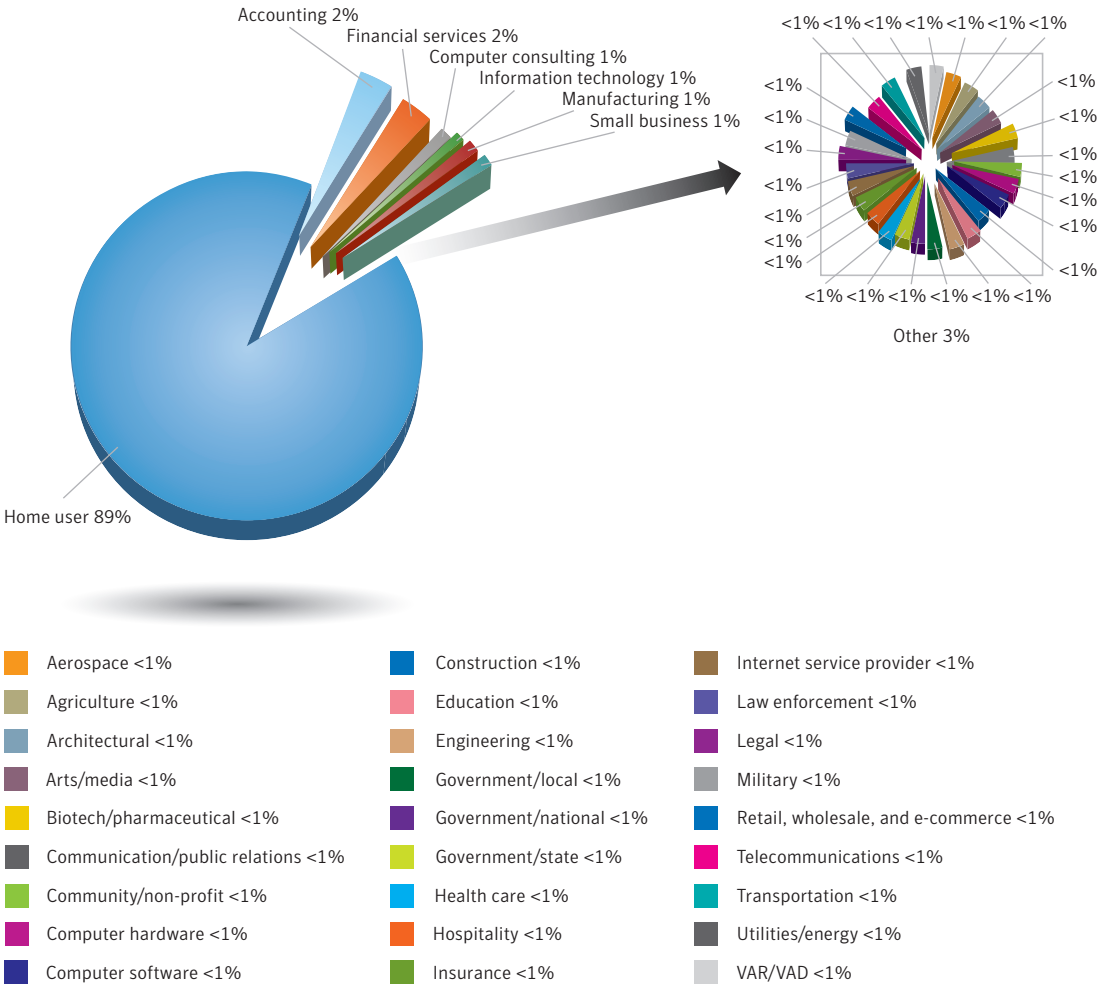

Figure 34. Distribution of sensors by sector
Source: Symantec Corporation

Malicious activity by country

To determine the top countries for the "Malicious activity by country" metric, Symantec compiled geographical data on each type of malicious activity to be considered, which included: bot network computers, bot command-and-control servers, phishing Web sites, malicious code submissions, spam relay hosts, and Internet attacks. The proportion of each activity originating in each country was determined. The mean of the proportions of each malicious activity that originated in each country was calculated. This average determined the proportion of overall malicious activity that originated from the country in question and was used to rank each country.

Symantec also evaluated the top 25 of these countries according to the percentage of worldwide Internet users located there. Symantec determined the top 25 countries for malicious activity per Internet users by employing the same data as above. This measure is meant to remove the bias of high Internet users from the consideration of the "Malicious activity by country" metric. In order to determine this, Symantec divided the amount of malicious activity originating in each of the top 25 countries for malicious activity by the percentage of worldwide Internet users who are located in that country.

The proportion assigned to each country in the discussion thus corresponds to the proportion of malicious activity that could be attributed to a single (average) Internet user in that country. That is, Symantec estimates the amount of malicious activity that could be attributed to the average Internet user from each of the top 25 countries. The proportion of malicious activity that would be carried out by each person is the proportion assigned to each country.

Underground economy servers

This metric is based on data gathered by proprietary Symantec technologies. These technologies monitor activity and collect data on underground economy servers. Underground economy servers are typically chat servers where stolen data, such as identities, credit card numbers, access to compromised computers, and email accounts are bought and sold. Symantec monitors this activity by recording communications that take place on these chat servers, which typically includes advertisements for stolen data. This data was used to derive the data presented in this metric.

Appendix C—Vulnerability Trends Methodology

The "Vulnerability Trends" section of the Symantec *Internet Security Threat Report* discusses developments in the discovery and exploitation of vulnerabilities over the past six months and compares that activity to activity observed in the two previous six-month periods. This methodology section will discuss how the data was gathered and how it was analyzed to come to the conclusions that are presented in the "Vulnerability Trends" section.

Symantec maintains one of the world's most comprehensive databases of security vulnerabilities, consisting of over 20,000 distinct entries. Each distinct entry is created and maintained by Symantec threat analysts who vet the content for accuracy, veracity, and the applicability of its inclusion in the vulnerability database based on available information. The following metrics discussed in the "Vulnerability Trends" report are based on the analysis of that data by Symantec researchers:

- Total number of vulnerabilities disclosed
- Severity of vulnerabilities
- Web application vulnerabilities
- Easily exploitable vulnerabilities (Total, and breakdown by type)
- Web browser vulnerabilities

The ways in which the data for the remaining metrics is gathered and analyzed will be discussed in the remainder of this methodology.

Vulnerability classifications

Following the discovery and/or announcement of a new vulnerability, Symantec analysts gather all relevant characteristics of the new vulnerability and create an alert. This alert describes important traits of the vulnerability, such as the severity, ease of exploitation, and a list of affected products. These traits are subsequently used both directly and indirectly for this analysis.

Vulnerability type

After discovering a new vulnerability, Symantec threat analysts classify the vulnerability into one of 12 possible categories based on the available information. These categories focus on defining the core cause of the vulnerability, as opposed to classifying the vulnerability merely by its effect.

The classification system is derived from the academic taxonomy presented by Taimur Aslam et al (1996) to define classifications of vulnerabilities.[112] Possible values are indicated below; the previously mentioned white paper provides a full description of the meaning behind each classification:

- Boundary condition error
- Access validation error
- Origin validation error
- Input validation error
- Failure to handle exceptional conditions

[112] "Use of a Taxonomy of Security Faults" http://ftp.cerias.purdue.edu/pub/papers/taimur-aslam/aslam-krsul-spaf-taxonomy.pdf

- Race condition error
- Serialization error
- Atomicity error
- Environment error
- Configuration error
- Design error

Severity of vulnerabilities

Severity of vulnerabilities has been discussed in previous versions of the Symantec *Internet Security Threat Report*; however, it was omitted in Volume X of the report (September 2006) to account for Symantec's adoption of the Common Vulnerability Scoring System (CVSS).[113]

The "Severity of vulnerabilities" metric that has been included in this report corresponds to the base score field of the CVSS. The base score is representative of the inherent properties of a vulnerability, such as: the degree of confidentiality, integrity, or availability of data that may be affected by the vulnerability; local versus remote exploitability; whether or not authentication is required for exploitation; and/or if there are additional factors that may complicate exploitation of the vulnerability.

These values are not adjusted for temporal factors such as the availability of exploit code. The base score is meant to be a static value that should only change if additional information is made available that changes the inherent characteristics of the vulnerability. The base score can have a value of zero to 10. For the sake of categorizing vulnerabilities by their respective severities, the following standard is used:

- **Low severity (base score of 0–3).** Successful exploitation of these vulnerabilities will have a minimal impact on the confidentiality, integrity, and availability of data stored upon or transmitted over systems on which the vulnerability may be found. These vulnerabilities also tend to be local in nature, have a high degree of access complexity, and may require authentication to be exploited successfully.
- **Medium severity (base score of 4–7).** Successful exploitation of these vulnerabilities may allow a "partial" compromise of the confidentiality, integrity, and availability of data stored upon or transmitted over systems on which the vulnerability may be found, although this may not always be the case. These vulnerabilities can be exploited remotely over a network and may have a lower access complexity or may or may not require authentication to successfully exploit.
- **High severity (base score of 8–10).** These vulnerabilities have innate characteristics that present the highest threat profile. Successful exploitation often allows a "complete" compromise of the confidentiality, integrity, and availability of data stored upon or transmitted over systems on which the vulnerability may be found. These vulnerabilities are exploited remotely across a network, have a low degree of access complexity, and usually do not require authentication prior to successful exploitation.

Base scores are computed from related fields in the Symantec Vulnerability Database. They are then categorized into low, medium, and high, as described above, and broken out by reporting period.

[113] http://www.first.org/cvss/cvss-guide.html

Easily exploitable vulnerabilities

The "Easily exploitable vulnerabilities" metric covers vulnerabilities that attackers can exploit with little effort based on publicly available information. The vulnerability analyst assigns an exploit availability rating after thoroughly researching the need for and availability of exploits for the vulnerability.

The "Easily exploitable vulnerabilities" metric replaces the "Ease of exploitation" metric, which was included in previous versions of the *Internet Security Threat Report*. This change was made to accommodate Symantec's adoption of the exploitability rating in the Common Vulnerability Scoring System (CVSS).[114]

All vulnerabilities are classified into one of four possible categories defined by the CVSS, as described below:

- **Unconfirmed.** Would-be attackers must use exploit code to make use of the vulnerability; however, no such exploit code is publicly available.
- **Proof-of-concept.** Would-be attacks must use exploit code to make use of the vulnerability; however, there is only proof-of-concept exploit available that is not functional enough to fully exploit the vulnerability.
- **Functional.** This rating is used under the following circumstances:

  1. Exploit code to enable the exploitation of the vulnerability is publicly available to all would-be attackers; and/or,

  2. Would-be attackers can exploit the vulnerability without having to use any form of exploit code.

  In other words, the attacker does not need to create or use complex scripts or tools to exploit the vulnerability.

- **High.** The vulnerability is reliably exploitable and there have been instances of self-propagating malicious code exploiting the vulnerability in the wild.

For the purposes of this report, the last two categories of vulnerabilities are considered "easily exploitable" because the attacker requires only limited sophistication to exploit the vulnerability. The first two categories of vulnerability are considered more difficult to exploit because attackers must develop their own exploit code or improve an existing proof-of-concept to make use of the vulnerability.

Easily exploitable vulnerabilities by type

This version of the *Internet Security Threat Report* includes an analysis of easily exploitable vulnerabilities by type. To provide further insight into the types of vulnerabilities that are considered easily exploitable, Symantec has categorized them into several categories. They are as follows:

- **Browser vulnerabilities.** These vulnerabilities threaten Web browser applications through remote attack vectors.
- **Client-side vulnerabilities.** These vulnerabilities threaten network client applications or non-networked applications that process malicious data that may arrive through another networked application. Remote

[114] http://www.first.org/cvss

attack vectors may exist, but client-side vulnerabilities usually require some amount of user interaction on the part of the victim to be exploited.

- **Local vulnerabilities.** These are vulnerabilities that require local access in order to be successfully exploited. Local attacks may affect a large variety of applications that may or may not include network capabilities. The differentiator is that these vulnerabilities are not exploitable by remote attackers unless they can log on to the system and interactively run commands as an unprivileged user.
- **Server vulnerabilities.** These are vulnerabilities that affect server applications. Server applications are typically defined as applications that are accessible to remote clients via connections on a range of TCP/UDP ports. Server vulnerabilities generally do not require user interaction on the part of the victim beyond enabling and starting the service so that it listens for incoming requests.
- **Web application vulnerabilities.** These vulnerabilities affect applications that use a browser for their user interface, rely on HTTP as the transport protocol, and reside on Web servers. Such applications are usually implemented in a server-side scripting language such as PHP or ASP.NET and are accessed through the HTTP/HTTPS protocols.
- **Other.** These are vulnerabilities that do not fall discretely into any of the previous categories. They can include applications for which the distinction is blurred between server and client, or hardware platforms in which the affected component cannot be described by any of the other categories.

These categories are generally defined by the attack vector and by the type of application that is affected. The specific categories were devised so that the majority of vulnerabilities could easily be classified within them, with little overlap between categories, so that the total percentage of all categories would equal 100 percent.

Window of exposure for enterprise vendors

Symantec records the time lapse between the publication of an initial vulnerability report and the appearance of third-party exploit code; this is known as the exploit development time. The time period between the disclosure date of a vulnerability and the release date of an associated patch is known as the patch development time.[115] The time lapse between the public release of exploit code and the time that the affected vendor releases a patch for the affected vulnerability is known as the window of exposure.

The window of exposure is calculated as the difference in days between the average exploit development time and the average patch development time. (Explanations of the exploit development time average and the patch development time average are included below.) During this time, the computer or system on which the affected application is deployed may be susceptible to attack, as administrators have no official recourse against the vulnerability and must resort to best practices and workarounds to reduce the risk of successful exploitation.

It is also important to note that the set of vulnerabilities included in this metric is limited and does not represent all software from all possible vendors. Instead, it only includes vendors who are classified as enterprise vendors. The purpose is to illustrate the window of exposure for widely deployed mission-critical software. Because of the large number of vendors with technologies that have a very low deployment

[115] This statistic only considers specific file-based patches or upgrades, and not general solutions. Instances in which the vendor provides a workaround or manual fix steps, for example, are not included.

(which form the majority), only exploits for technologies from enterprise vendors (that is, those that generally have widespread deployment) are included. Vulnerabilities in those vendors' products will likely affect more enterprises than those in less widely deployed technologies. Those vendors are:

- CA™ (Computer Associates)
- Cisco®
- EMC
- HP®
- IBM®
- McAfee®
- Microsoft
- Oracle®
- Sun™
- Symantec

Patch development time for enterprise vendors

The patch development time is the time period between the disclosure date of a vulnerability and the release date of an associated patch. Only those patches that are independent objects (such as fixes, upgrades, etc.) are included in this analysis. Other remediation solutions—such as workaround steps, for instance—are excluded.

For each individual patch from these vendors, the time lapse between the patch release date and the publish date of the vulnerability is computed. The mean average is calculated from the aggregate of these. As some vendors may release more patches than others for a particular vulnerability, Symantec considers only the first instance of a single patch for each vulnerability. This metric is incorporated when computing the window of exposure, which is calculated as the difference between the average patch development time and the average exploit development time.

Exploit code development time for enterprise vendors

The ability to measure exploit code development time is limited and applies only to vulnerabilities that would normally require exploit code. Therefore, the metric is based on vulnerabilities that Symantec considers to be of sufficient complexity, and for which functional exploit code was not available until it was created by a third party. This consideration, therefore, excludes the following:

- Vulnerabilities that do not require exploit code (unconfirmed exploitability)
- Vulnerabilities associated with non-functional proof-of-concept code (proof-of-concept exploitability)

The date of vulnerability disclosure is based on the date of the first publicly available reference (such as a mailing list post). The date of exploit code publication is the date of the first publicly known reference to the exploit code. Because the purpose of this metric is to estimate the time it takes for exploit code to materialize as a result of active development, exploit code publication dates that fall outside of the 30-day range from initial vulnerability publication are excluded from this metric. It is assumed that exploit code that was published after this period was not actively developed from the initial announcement of the vulnerability.

Since this metric only considers the appearance of the first functional exploit, it is possible that reliable exploits that improve upon the initial exploit may appear later. These exploits may take much longer to develop, but are not considered because the window of exposure begins as soon as the first functional exploit surfaces.

The time lapse between the disclosure of a vulnerability and the appearance of exploit code for that vulnerability is determined. The aggregate time for all vulnerabilities is determined and the average time is calculated. This metric is incorporated when computing the window of exposure, which is the difference between the average patch development time and the average exploit development time.

Operating system patch development time

This metric has a similar methodology to the "Patch development time for enterprise vendors" metric, which was explained earlier in this methodology. However, instead of applying it to enterprise-scale vendors, the patch development time average is calculated from patched vulnerabilities for the following operating systems:

- Apple Mac OS X
- Hewlett-Packard HP-UX
- Microsoft Windows
- Red Hat Linux (including enterprise versions and Red Hat Fedora)
- Sun Microsystems Solaris

An average is calculated from the patch release times for each vulnerability in the reporting period per operating system. The patch development time average for each operating system is then compared.

Window of exposure for Web browsers

This metric has a similar methodology to the "Window of exposure for enterprise vendors" metric. However, instead of applying it to enterprise-scale vendors, the window of exposure is calculated for vulnerabilities associated with the following Web browsers:

- Apple Safari
- Microsoft Internet Explorer
- Mozilla Firefox and Mozilla browsers
- Opera

Symantec records the window of time between the publication of an initial vulnerability report and the appearance of third-party exploit code; this is known as the exploit code development time. The time period between the disclosure date of a vulnerability and the release date of an associated patch is known as the patch development time.[116] The time lapse between the public release of exploit code and the time that the affected vendor releases a patch for the affected vulnerability is known as the window of exposure.

The window of exposure is calculated as the difference in days between the average patch development time and the average exploit code development time. During this time, the computer or system on which the affected application is deployed may be susceptible to attack, as administrators may have no official

[116] This statistic only considers specific file-based patches or upgrades, and not general solutions. Instances in which the vendor provides a workaround or manual repair steps, for example, are not included.

recourse against a vulnerability and must resort to best practices and workarounds to reduce the risk of attacks. Explanations of the average exploit development time and the average patch development time are included below.

Patch development time for Web browsers

The cumulative patch development time for vulnerabilities affecting each browser is calculated. Each cumulative time is then divided by the number of vulnerabilities affecting that browser to determine the average patch development time for that browser. The patch development time average for each browser is then compared. This metric is used to compute the window of exposure for Web browsers, which amounts to the difference between the average patch development time and the average exploit code development time.

Exploit code development time for Web browsers

The cumulative exploit code development time for each vulnerability affecting a Web browser is calculated. Each cumulative time is then divided by the number of vulnerabilities affecting that browser to determine the average exploit code development time for that browser. The exploit development time average for each browser is then compared. This metric is used to compute the window of exposure, which amounts to the difference between the average patch development time and the average exploit code development time.

Web browser vulnerabilities

This metric will offer a comparison of vulnerability data for numerous Web browsers, namely: Microsoft Internet Explorer, the Mozilla browsers (which includes Firefox), Opera, and Safari. However, in assessing the comparative data, the following important caveats should be kept in mind before making any conclusions:

- The total number of vulnerabilities in the aforementioned Web browsers was computed for this report. This includes vulnerabilities that have been confirmed by the vendor and those that are not vendor confirmed.

  Previous versions of the *Internet Security Threat Report* have discussed vulnerabilities according to whether they were vendor confirmed or non-vendor confirmed in that vulnerabilities that are not confirmed are also included in the data. This differentiation was important, especially given the disparity in patch times between vendors. However, starting with Volume X of the *Internet Security Threat Report*, this convention was no longer followed. This version of the report does not differentiate between vendor-confirmed vulnerabilities and non-vendor-confirmed vulnerabilities when calculating the total number of vulnerabilities.

- Individual browser vulnerabilities are notoriously difficult to pinpoint and identify precisely. A reported attack may be a combination of several conditions, each of which could be considered a vulnerability in its own right. This may distort the total vulnerability count. Some browser issues have also been

improperly identified as operating system vulnerabilities or vice versa. This is, in part, due to increasing operating system integration that makes it difficult to correctly identify the affected component in many cases.

Many vulnerabilities in shared operating system components can potentially be exposed to attacks through the browser. This report, where sufficient information is available to make the distinction, enumerates only those vulnerabilities that are known to affect the browser itself.

- Not every vulnerability that is discovered is exploited. As of this writing, there has been no widespread exploitation of any browser except Microsoft Internet Explorer. This is expected to change as other browsers become more widely deployed.

Exploit code release period

This metric provides a breakdown of the number of exploits according to the length of time that has elapsed between the publication of a vulnerability and the release of the exploit code. The exploit code release periods have been categorized in the following increments: less than one day, one to six days, seven to 30 days, 31 to 100 days, and more than 100 days. This is computed by comparing the vulnerability publication date against the date that an instance of exploit code was published, and then categorizing it in the appropriate time period.

Unlike the "Exploit code development time" metrics that are described previously in this methodology, this metric does include multiple instances of exploits for a single vulnerability. Additional exploit code of varying quality and reliability may be released after the initial appearance of a first exploit for a vulnerability. Some exploit code may not be developed until well after the release of a vulnerability for several reasons. These could include:

- The vulnerability is particularly difficult to exploit.
- The exploit code is advanced and improves upon previous exploit code.
- The vulnerability is considered to be a lower priority for attackers and thus have not received the concerted exploit development effort associated with high profile vulnerabilities.

There have also been instances where exploit code surfaces in the wild much after the initial publication of a vulnerability.

Zero-day vulnerabilities

This metric quantifies the number of zero-day vulnerabilities that have been documented during the relevant reporting periods of the current *Internet Security Threat Report*.

For the purpose of this metric, a zero-day vulnerability is one for which there is sufficient public evidence to indicate that the vulnerability has been exploited in the wild prior to being publicly known. It may not have been known to the vendor prior to exploitation, and the vendor had not released a patch at the time of the exploit activity.

This metric is derived from a mix of data from public sources and the Symantec vulnerability database. This metric is meant to calculate the number of high-profile, publicly documented zero-day vulnerability instances during the relevant reporting periods.

Vendor responsiveness

This metric quantifies the number of vulnerabilities that have not been confirmed by vendors over the relevant reporting periods. The metric is calculated by determining the number of vulnerabilities that are not considered to be "vendor confirmed," or confirmed and patched by the vendor, and comparing them to the total number of vulnerabilities documented in the period. This gives insight into the number of vulnerabilities that remain unconfirmed and unpatched over time.

Database vulnerabilities

This metric offers a comparison of the vulnerabilities across multiple database vendors and implementations. For the purpose of this report, databases to be assessed were chosen to reflect the most widely deployed database implementations and to compare commercial and open source vendors.[117] To this end, the following five database implementations are discussed:

- IBM® DB2
- Microsoft® SQL Server
- MySQL
- Oracle®
- PostgreSQL

The volume of database vulnerabilities is determined by querying the vulnerability database for vulnerabilities that affect the aforementioned database implementations. The results are broken out by implementation and reporting period.

[117] Oracle, DB2, and Microsoft SQL Server are the three most widely deployed commercial database implementations (http://databases.about.com/b/a/016881.htm). MySQL and PostgreSQL are the two most popular open-source databases (http://www.mysql.com/why-mysql/marketshare).

Appendix D—Malicious Code Trends Methodology

The trends in the "Malicious Code Trends" section are based on statistics from malicious code samples reported to Symantec for analysis. Symantec gathers data from over 120 million client, server, and gateway systems that have deployed Symantec's antivirus products in both consumer and corporate environments. The Symantec Digital Immune System and Scan and Deliver technologies allow customers to automate this reporting process.

Observations in the "Malicious Code Trends" section are based on empirical data and expert analysis of this data. The data and analysis draw primarily from two databases described below.

Infection database

To help detect and eradicate computer viruses, Symantec developed the Symantec AntiVirus™ Research Automation (SARA) technology. Symantec uses this technology to analyze, replicate, and define a large subset of the most common computer viruses that are quarantined by Symantec Antivirus customers.

On average, SARA receives hundreds of thousands of suspect files daily from both enterprise and individual consumers located throughout the world. Symantec then analyzes these suspect files, matching them with virus definitions. An analysis of this aggregate data set provides statistics on infection rates for different types of malicious code.

Malicious code database

In addition to infection data, Symantec Security Response analyzes and documents attributes for each new form of malicious code that emerges both in the wild and in a "zoo" (or controlled laboratory) environment. Descriptive records of new forms of malicious code are then entered into a database for future reference. For this report, a historical trend analysis was performed on this database to identify, assess, and discuss any possible trends, such as the use of different infection vectors and the frequency of various types of payloads.

In some cases, Symantec antivirus products may initially detect new malicious code heuristically or by generic signatures. These may later be reclassified and given unique detections. Because of this, there may be slight variance in the presentation of the same data set from one volume of the *Internet Security Threat Report* to the next.

Previously unseen malicious code threats

This metric derives its data from the Symantec Honeypot Network. Computers compromised on the honeypot network track and analyze each piece of malicious code that is installed by the attacker. Symantec defines previously unseen malicious threats as those that have not been installed by attackers on the Symantec Honeypot Network. The proportion of previously unseen malicious code threats is derived by comparison with the total number of distinct malicious code threats observed.

Percentage of malicious code that exploits vulnerabilities

Symantec maintains a malicious code database to analyze and document individual instances of malicious code. This database contains 8,000 distinct entries, with the earliest discovery dating back to 1998. The database includes metadata for classifying malicious code by type, discovery date, and by threat profile, in addition to providing mitigating factors and manual removal steps. Where applicable, this database includes correlations between malicious code instances and vulnerabilities from the Symantec vulnerability database. This capability was used as a basis for the data in this metric. Symantec examined the means by which the malicious code propagated, and counted those that propagate by exploiting vulnerabilities.

Appendix E—Phishing, Spam, and Security Risks Methodology

Traditionally, the Symantec *Internet Security Threat Report*, has broken security threats down into three general categories: attacks, vulnerabilities, and malicious code. However, as Internet-based services and applications have expanded and diversified, the potential for computer programs to introduce other types of security risks has increased. The emergence of new risks, particularly spam, phishing, spyware, adware, and misleading applications has necessitated an expansion of the traditional security taxonomy.

Symantec has monitored these new concerns as they have developed. In particular, the *Internet Security Threat Report* assess these risks according to three categories:

- Phishing
- Spam
- Security risks, particularly adware, spyware, and misleading applications

The methodology for each of these discussions will be discussed in the sections below.

Phishing

Phishing attack trends in this report are based on the analysis of data derived from the Symantec Probe Network. Symantec Brightmail AntiSpam data is assessed to gauge the growth in phishing attempts as well as the percentage of Internet mail that is determined to be phishing attempts. Symantec Brightmail AntiSpam field data consists of statistics reported back from customer installations that provide feedback about the detection behaviors of antifraud filters as well as the overall volume of mail being processed.

It should be noted that different monitoring organizations use different methods to track phishing attempts. Some groups may identify and count unique phishing messages based solely on specific content items such as subject headers or URLs. These varied methods can often lead to differences in the number of phishing attempts reported by different organizations.

Phishing attempt definition

The Symantec Probe Network is a system of over two million decoy accounts that attract email messages from 20 different countries around the world. It encompasses more than 600 participating enterprises and attracts email samples that are representative of traffic that would be received by over 250 million mailboxes. The Probe Network covers countries in the Americas, Europe, Asia, Africa, and Australia/Oceania.

The Symantec Probe Network data is used to track the growth in new phishing activity. A phishing attempt is a group of email messages with similar properties, such as headers and content, that are sent to unique users. The messages attempt to gain confidential and personal information from online users.

Symantec Brightmail AntiSpam software reports statistics to Symantec Security Response that indicate messages processed, messages filtered, and filter-specific data. Symantec has classified different filters so that spam statistics and phishing statistics can be determined separately. Symantec Brightmail AntiSpam field data is used to identify general trends in phishing email messages.

Explanation of research inquiries

This section will provide more detail on specific methodologies used to produce the data and statistics in this report. While most methodologies are adequately explained in the analysis section of the report, the following investigations warrant additional detail.

Daily and seasonal variations in phishing activity

The data for this section is determined by the number of email messages that trigger antifraud filters in the field versus the total number of email messages scanned. These filters are distributed across the Symantec Brightmail AntiSpam customer base. The data for this section is based on daily totals for each day of the week.

Unique phishing messages

Symantec maintains automated systems to identify new unique phishing messages received by the Symantec Probe Network. Messages are grouped into attacks based on similarities in the message bodies and headers. Sample messages are then passed through general fraud heuristics to identify messages as potential phishing attempts. Symantec reviews events that are identified as phishing attempts for the purposes of confirmation and to develop filters.

The data presented in this section is based on monthly totals in the number of new unique phishing messages discovered and ruled upon by Symantec Security Response. Security Response addresses only those phishing messages not caught by existing antispam and antifraud filters. Existing filters refer only to those antispam and antifraud filters used across the Symantec Brightmail AntiSpam customer base.

Some phishing messages will be captured in the field based upon predictive filters (heuristics); however, not all of Symantec's customers utilize this technology or have upgraded to this technology. Therefore, the messages are still reviewed by Security Response for development of filters that are more widely dispersed.

Blocked phishing attempts

The number of blocked phishing attempts is calculated from the total number of phishing email messages that were blocked in the field by Symantec Brightmail AntiSpam antifraud filters. The data for this section is based on monthly totals.

Phishing activity by sector

The Symantec Phish Report Network is an extensive antifraud community where members contribute and receive fraudulent Web site addresses for alerting and filtering across a broad range of solutions.[118] These sites are categorized according to the brand being phished and the industry to which it belongs. The Phish Report Network has members and contributors that send in phishing attacks from many different sources. This includes a client detection network that detects phishing Web sites as the clients visit various Web sites on the Internet. It also includes server detection from spam emails.

The sender confirms all spoofed Web sites before sending the address of the Web site into the Phish Report Network. After the spoofed site is sent into the Phish Report Network, Symantec spoof detection technology is used to verify that the Web site is a spoof site. Research analysts manage the Phish Report Network Console 24 hours a day, 365 days of the year, and manually review all spoof sites sent into the Phish Report Network to eliminate false positives.

Top countries hosting phishing sites

The data for this section is determined by gathering links in phishing email messages and cross-referencing the addresses with several third-party subscription-based databases that link the geographic locations of systems to IP addresses. While these databases are generally reliable, there is a small margin of error. The data produced is then used to determine the global distribution of phishing Web sites.

Spam

The Symantec Probe Network is a system of over two million decoy accounts that attract email messages from 20 different countries around the world. It encompasses more than 600 participating enterprises and attracts email samples that are representative of traffic that would be received by over 250 million mailboxes. The Probe Network includes accounts in countries in the Americas, Europe, Asia, Africa, and Australia/Oceania.

Spam trends in this report are based on the analysis of data derived from both the Symantec Probe Network as well as Symantec Brightmail AntiSpam field data. Symantec Brightmail AntiSpam software reports statistics to the Brightmail Logistical Operations Center (BLOC) indicating messages processed, messages filtered, and filter-specific data.

Symantec has classified different filters so that spam statistics and phishing statistics can be determined separately. Symantec Brightmail AntiSpam field data includes data reported back from customer installations providing feedback from antispam filters as well as overall mail volume being processed.

Symantec Brightmail AntiSpam only gathers data at the SMTP layer and not the network layer, where DNS block lists typically operate. This is because SMTP-layer spam filtering is more robust than network-layer filtering and is able to block spam missed at the network layer. Network layer-filtering takes place before email reaches the enterprise mail server. As a result, data from the SMTP layer is a more accurate reflection of the impact of spam on the mail server itself.

[118] http://www.phishreport.net

Sample set normalization

Due to the numerous variables influencing a company's spam activity, Symantec focuses on identifying spam activity and growth projections with Symantec Brightmail AntiSpam field data from enterprise customer installations having more than 1,000 total messages per day. This normalization yields a more accurate summary of Internet spam trends by ruling out problematic and laboratory test servers that produce smaller sample sets.

Explanation of research inquiries

This section will provide more detail on specific methodologies used to produce the data and statistics in this report. While most methodologies are adequately explained in the analysis section of the report, the following investigations warranted additional detail.

Spam as a percentage of email scanned

The data for this section is determined by dividing the number of email messages that trigger antispam filters in the field by the total number of email messages scanned. These filters are distributed across the Symantec Brightmail AntiSpam customer base. The data for this section is based on monthly totals.

Top ten countries of spam origin

The data for this section is determined by calculating the frequency of originating server IP addresses in email messages that trigger antispam filters in the field. The IP addresses are mapped to their host country of origin and the data is summarized by country based on monthly totals. The percentage of spam per country is calculated from the total spam detected in the field.

It should be noted that the location of the computer from which spam is detected being sent is not necessarily the location of the spammer. Spammers can build networks of compromised computers globally and thereby use computers that are geographically separate from their location.

Top countries by spam zombies

The data for this section is determined by examining the IP addresses in spam messages received by the Symantec Probe Network. Only IP addresses that are dynamically assigned are examined. If the computers at those IP addresses do not appear to be email servers—for example, if they do not respond to requests on TCP port 25—they are classified as spam zombies. Symantec then cross-references the addresses with several third-party subscription-based databases that link the geographic locations of systems to IP addresses. While these databases are generally reliable, there is a small margin of error. The data produced is then used to determine the global distribution of spam zombies.

Security Risks

Symantec products not only help users to protect their data from the threat of viruses, worms, and Trojan horses, but to evaluate potential security risks from the introduction of other programs as well. Symantec AntiVirus classifies these other programs as additional security risks. Security risks include programs that may be categorized, based upon functional criteria, as adware, spyware, and misleading applications. Symantec classifies these programs based on a number of characteristics. Once categorized, they can be detected, allowing users to choose whether to keep or remove them based on their personal needs and security policies.

General criteria for security risks

A program classified as an additional security risk is an application or software-based executable that is either independent or interdependent on another software program and meets the following criteria:

- It is considered to be non-viral in nature;
- It meets criteria for programmatic functionality having potential to impact security; and/or,
- It has been reported to Symantec by a critical number of either corporate or individual users within a given time frame. The time frame and number may vary by category or risk.

Symantec further classifies programs based upon functional criteria related to the result of the program's introduction to a computer system. The criteria take into consideration functionality that includes stealth, privacy, performance impact, damage, and removal.

Adware, spyware, and misleading applications

Adware programs are those that facilitate the delivery and display of advertising content onto the user's display device. This may be done without the user's prior consent or explicit knowledge. The advertising is often, but not always, presented in the form of pop-up windows or bars that appear on the screen. In some cases, these programs may gather information from the user's computer, including information related to Internet browser usage or other computing habits, and relay this information back to a remote computer.

Spyware programs are stand-alone programs that can unobtrusively monitor system activity and either relay the information back to another computer or hold it for subsequent retrieval. In some cases, spyware programs may be used by corporations to monitor employee Internet usage or by parents to monitor their children's Internet usage.

Spyware programs can be surreptitiously placed on users' systems in order to gather confidential information such as passwords, login details, and credit card details. This can be done through keystroke logging and by capturing email and instant messaging traffic.

Misleading applications are programs that intentionally misrepresent the security status of a computer by informing the user that a threat, usually nonexistent or fake, is on the user's computer. This is usually done in order to persuade the user to pay money to upgrade to a paid-for version of the software that will remove the "threats" that are claimed to be found.

The potential security risks introduced by adware, spyware, and misleading applications are discussed according to samples, or individual cases of each security risk, reported to Symantec by customers deploying Symantec AntiVirus. While security risks are not categorized as malicious code, Symantec monitors them using many of the same types of methods used for tracking malicious code development and proliferation. This involves an ongoing analysis of reports and data delivered from over 120 million client, server, and gateway email systems, as well as filtration of 25 million email messages per day. Symantec then compiles the most common reports and analyzes them to determine the appropriate categorization. The discussion included in the "Security Risks" section is based on Symantec's analysis of these reports.

About Symantec

Symantec is a global leader in infrastructure software, enabling businesses and consumers to have confidence in a connected world. The company helps customers protect their infrastructure, information, and interactions by delivering software and services that address risks to security, availability, compliance, and performance. Headquartered in Cupertino, Calif., Symantec has operations in 40 countries. More information is available at www.symantec.com.

For specific country offices and contact numbers, please visit our Web site. For product information in the U.S., call toll-free 1 (800) 745 6054.

Symantec Corporation
World Headquarters
20330 Stevens Creek Blvd.
Cupertino, CA 95014 USA
+1 (408) 517 8000
1 (800) 721 3934
www.symantec.com

03/07 12078591