SYMANTEC ENTERPRISE SECURITY

Symantec Internet Security Threat Report

Trends for January 06–June 06

Volume X, Published September 2006

Dean Turner
Executive Editor
Symantec Security Response

Stephen Entwisle
Editor
Symantec Security Response

Marc Fossi
Analyst—DeepSight Threat Analyst
Symantec Security Response

Joseph Blackbird
Analyst—Assoc. Software Engineer
Symantec Security Response

David McKinney
Analyst—Manager Software Engineering
Symantec Security Response

Tony Conneff
Analyst—Development Manager
Symantec Security Response

Ollie Whitehouse
Technical Advisor—Security Architect
Symantec Security Response

Contributors

Dave Cole
Director, Product Management
Symantec Security Response

Peter Szor
Security Architect
Symantec Security Response

Peter Ferrie
Sr. Principal Software Engineer
Symantec Security Response

David Cowings
Sr. Business Intelligence Manager
Symantec Business Intelligence

Dylan Morss
Principal Business Intelligence Manager
Symantec Business Intelligence

Scott Carlton
Manager, IT Operations
Product Operations

Igor Moochnick
Sr. Software Engineer
Instant Messaging Security

September 25, 2006

A message from the Executive Editor

On January 28, 2002, the first *Internet Security Threat Report* was published by Riptech, a Managed Security Services company that was acquired by Symantec in July 2002. At a little over 33 pages, the initial *Internet Security Threat Report* was one of the first reports to summarize and analyze network attack trends in a single, comprehensive document. The premiere issue was based on data captured by Riptech's firewall and intrusion detection systems, which the company's analysts used to produce a first-of-its kind report on attack trends. In that first issue, Code Red and Nimda dominated the threat landscape and large blended threats and perimeter attacks were the attackers' modus operandi.

Since that first report, much has changed. Large Internet worms targeting everything and everyone have given way to smaller, more targeted attacks focusing on fraud, data theft, and criminal activity. The days of Web site defacements and low-level information gathering attacks are behind us. Today we are seeing encrypted bot networks, remotely initiated database breaches, sophisticated phishing scams, and customized malicious code targeting specific companies. As threats have evolved, so too has the job of tracking and reporting on them.

Over the past four years, the Symantec™ Global Intelligence Network has expanded to include data from millions of antivirus products and thousands of intrusion detection sensors deployed around the globe, as well as data gathered from Symantec antifraud solutions. This exponential growth in data collection has allowed us to produce one of the most thorough and complete analyses of current Internet threat activity in the world.

Utilizing a team of more than 1,600 dedicated security analysts around the globe, the *Internet Security Threat Report* has become much more than just a collection of facts and figures. It has become an invaluable tool in helping enterprise organizations, small businesses, and consumers to make sense of the ever-changing threat landscape and secure their systems accordingly.

Today, Symantec is pleased to announce the latest *Internet Security Threat Report*, Volume X. Four years and nine reports after Riptech's innovative first effort, this edition incorporates a number of changes to both the look and feel of the report, as well as new metrics analyzing and discussing emerging threat trends. The dedicated team of individuals who write, compile, and edit the report have spent hundreds of hours analyzing data and trends to bring you what we hope is the most comprehensive and thought-provoking report to date.

On behalf of the entire Symantec team, I hope you find this report as interesting and informative to read as we have found it to research, develop, and publish.

Sincerely,

Dean Turner
Executive Editor

Volume X, Published September 2006

Symantec Internet Security Threat Report

Contents

Internet Security Threat Report Volume X Executive Summary

Previous editions of the Symantec *Internet Security Threat Report* have discussed a shift in the threat landscape. In contrast to previously observed widespread, network-based attacks, attackers today tend to be more focused, often targeting client-side applications. This has had numerous effects on security issues.

As vendors and enterprises have adapted to the changing threat environment by implementing best security practices and defense in-depth strategies, attackers have begun to adopt new techniques. In part, this has resulted in more targeted malicious code and targeted attacks aimed at client-side applications, such as Web browsers, email clients, and other applications. These applications are used to communicate over networks and interact with Web-based services and applications and Web sites. They may also include programs such as word processing or spreadsheet programs, which can open untrusted content that is downloaded or received by a network client.

As security technologies have matured to address the types of flaws typically exploited by traditional attacks, attackers have shifted their focus to new attack vectors. Further, as technological solutions are proving increasingly more effective, attackers are reverting to older, non-technical means of compromise, such as social engineering, in order to launch successful attacks.[1] Attackers are thus shifting attack activity away from network infrastructures and operating system services toward attacks that focus on the end user as the weakest link in the security chain.

The current threat landscape is populated by lower profile, more targeted attacks, attacks that propagate at a slower rate in order to avoid detection and thereby increase the likelihood of successful compromise. Instead of exploiting vulnerabilities in servers, as traditional attacks often did, these threats tend to exploit vulnerabilities in client-side applications that require a degree of user interaction, such as word processing and spreadsheet programs. A number of these have been zero-day vulnerabilities.[2] These types of threats also attempt to escape detection in order to remain on host systems for longer periods so that they can steal information or provide remote access.

Previous editions of the *Internet Security Threat Report* have also remarked that attack activity has shifted from being motivated by status for technical prowess to being motivated by financial gain. Many of today's threats are designed to gather information that has some value to the attacker. This may consist of personal information that can be used for the purpose of identity theft or fraud, or it may have the potential to be used for corporate espionage, as in the case of the Hotword Trojan[3] and more recently the Ginwui[4] and PPDropper Trojans.[5] In the enterprise environment, such targeted threats could be used to gain unauthorized access to privileged, proprietary information, thereby threatening the intellectual property of the organization.

This volume of the *Internet Security Threat Report* will offer an analysis and discussion of threat activity that took place between January 1 and June 30, 2006. This brief summary will offer a synopsis of the data and trends discussed 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.

[1] Social engineering refers to the use of persuasive techniques, manipulation, and/or deception to persuade or fool a computer user into disclosing confidential information that can then be used to gain unauthorized access to computers or information stored upon computers, usually for malicious purposes.
[2] A zero-day vulnerability is one that has not yet been disclosed publicly and that may not yet be known of by the vendor of the affected technology.
[3] http://www.securityfocus.com/news/11209
[4] http://www.symantec.com/security_response/writeup.jsp?docid=2006-051914-5151-99
[5] http://www.eweek.com/article2/0,1895,1992128,00.asp

Internet Security Threat Report Highlights

Attack Trend Highlights

- Microsoft Internet Explorer was the most frequently targeted Web browser, accounting for 47% of all Web browser attacks.
- Symantec observed an average of 6,110 DoS attacks per day.
- The United States was the target of the most DoS attacks, accounting for 54% of the worldwide total.
- The Internet service provider (ISP) sector was the most frequently targeted by DoS attacks.
- China had the highest number of bot-infected computers during the first half of 2006, accounting for 20% of the worldwide total.
- The United States had the highest percentage of bot command-and-control servers with 42%.
- Beijing was the city with the most bot-infected computers in the world.
- The United States ranked as the top country of attack origin, accounting for 37% of the worldwide total.
- The home user sector was the most highly targeted sector, accounting for 86% of all targeted attacks.

Vulnerability Trend Highlights

- Symantec documented 2,249 new vulnerabilities, up 18% over the second half of 2005. This is the highest number ever recorded for a six-month period.
- Web application vulnerabilities made up 69% of all vulnerabilities this period.
- Mozilla browsers had the most vulnerabilities, 47, compared to 38 in Microsoft Internet Explorer.
- In the first six months of 2006, 80% of vulnerabilities were considered easily exploitable, up from 79%.
- Seventy-eight percent of easily exploitable vulnerabilities affected Web applications.
- The window of exposure for enterprise vulnerabilities was 28 days.
- Internet Explorer had an average window of exposure of nine days, the largest of any Web browser. Apple Safari averaged five days, followed by Opera with two days and Mozilla with one day.
- In the first half of 2006, Sun operating systems had the highest average patch development time, with 89 days, followed by Hewlett Packard with 53 days, Apple with 37 days and Microsoft and Red Hat with 13 days.

Internet Security Threat Report Highlights *continued*

Malicious Code Trend Highlights

- Eighteen percent of all distinct malicious code samples detected by Symantec honeypots were new.
- Five of the top ten new malicious code families reported were Trojan horse programs.
- The most prevalent new malicious code family this period was that of the Polip virus.
- Worms made up 38 of the top 50 malicious code samples.
- Worms made up 75% of the volume of top 50 malicious code reports.
- Symantec documented 6,784 new Win32 viruses and worms.
- Bots accounted for 22% of the top 50 malicious code reports, up slightly from the 20% reported in the last period.
- Thirty of the top 50 malicious code samples exposed confidential information.
- Modular malicious code accounted for 79% of the volume of top 50 malicious code, down from 88% in the second half of 2005.

Phishing, Spam and Security Risks

- The Symantec Probe Network detected 157,477 unique phishing messages, an increase of 81%.
- Financial services was the most heavily phished sector, accounting for 84% of phishing activity.
- Spam made up 54% of all monitored email traffic, up from 50% in the last period.
- The most common type of spam detected in the first six months of 2006 was related to health services and products.
- Fifty-eight percent of all spam detected worldwide originated in the United States
- Eight of the top ten reported security risks were adware programs.
- Three of the top ten new security risks are what Symantec calls “misleading applications.”

Internet Security Threat Report Overview

The Symantec *Internet Security Threat Report* provides a six-month update of Internet threat activity. It includes an analysis of network-based attacks, disclosed vulnerabilities, malicious code reports, and security risks. This summary of the most recent *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 January 1 to June 30, 2006.

Symantec has established some of the most comprehensive sources of Internet threat data in the world. The Symantec™ Global Intelligence Network comprehensively tracks attack activity across the entire Internet. The Global Intelligence Network, which includes the Symantec DeepSight™ Threat Management System and Symantec™ Managed Security Services, 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 also maintains one of the world's most comprehensive databases of security vulnerabilities, covering over 18,000 vulnerabilities affecting more than 30,000 technologies from over 4,000 vendors. In addition to the vulnerability database, Symantec operates BugTraq™, one of the most popular forums for the disclosure and discussion of vulnerabilities on the Internet. 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. Finally, the Symantec Phish Report Network is an extensive antifraud community in which members contribute and receive fraudulent Web site addresses for alerting and filtering across a broad range of solutions. These resources give Symantec analysts unparalleled sources of data with which to identify emerging trends in attacks and malicious code activity.

The Symantec *Internet Security Threat Report* is grounded principally on the expert analysis of this data. 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 *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.

Distribution of attacks targeting Web browsers

In the first six months of 2006, Microsoft Internet Explorer was the most frequently targeted Web browser. Attacks targeting it accounted for 47% of all attacking computers targeting Web browsers (figure 1). The prominence of Microsoft Internet Explorer is not surprising, as it is the most widely deployed browser worldwide. Furthermore, it had the second highest number of vulnerabilities of all Web browsers during this period.

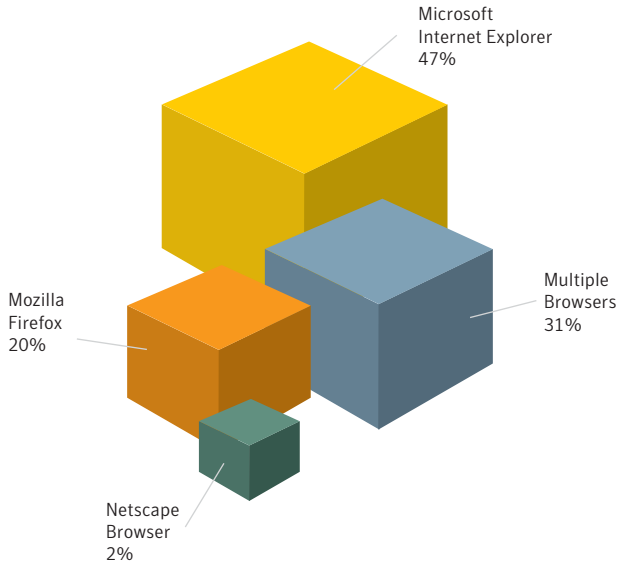

Figure 1. Distribution of attacks targeting Web browsers
Source: Symantec Corporation

Some attacks target vulnerabilities that are present in more than one Web browser. These vulnerabilities, which are referred to here as "multiple browser vulnerabilities," are typically present in numerous browsers because of shared source code, although this is not always the case. Multiple browser vulnerabilities may affect Apple Safari, KDE Konqueror, the Mozilla Browser family, Netscape, Opera, and/or Internet Explorer. Attacks targeting multiple browsers were the second most common Web browser attacks during the first half of 2006, accounting for 31% of all attacks targeting Web browsers.

Mozilla Firefox was targeted by the third highest number of detected Web browser attacks during the first half of 2006. Twenty percent of all attacking IP addresses targeted the Firefox browser during this period.

In order to protect against Web browser attacks, Symantec advises users and administrators to upgrade all browsers to the latest, patched versions. To reduce exposure to attacks, Symantec recommends that organizations educate users to be extremely cautious about visiting unknown or untrusted Web sites and/or viewing or following links in unsolicited emails.

Top targeted sectors

Although many attackers choose targets randomly, some target computers within a specific sector, industry, or organization. Symantec refers to these as "targeted attacks." For the purposes of this metric, a targeted attack is identified as an IP address that has attacked at least three sensors in a given industry to the exclusion of all other industries during the reporting period.

Current Rank	Previous Rank	Sector	Current Proportion of Targeted attacks	Previous Proportion of Targeted attacks
1	1	Home user	86%	93%
2	2	Financial Services	14%	4%
3	6	Government	<1%	<1%
4	3	Education	<1%	2%
5	8	Information Technology	<1%	<1%
6	7	Health care	<1%	<1%
7	5	Accounting	<1%	<1%
8	10	Telecommunications	<1%	<1%
9	4	Small Business	<1%	<1%
10	14	Utilities / Energy	<1%	<1%

Table 1. Top targeted sectors
Source: Symantec Corporation

Between January 1 and June 30, 2006, the home user sector was the most highly targeted sector, accounting for 86% of all targeted attacks (table 1). As computers in the home users sector are less likely to have well established security measures and practices in place than other sectors, they are much more vulnerable to targeted attacks. Furthermore, as home users represent a fertile resource for identity theft, it is likely that many of the targeted attacks against them are used for fraud or other financially motivated crime.

The number of targeted attacks detected against home users might be inflated due to the way in which they connect to the Internet. Because home users generally connect to the Internet through Internet service providers (ISPs), it is likely that the majority of them share networks that span a single block of IP addresses. As a result, opportunistic attacks targeting a broadband ISP may be noted as targeted attacks, thereby artificially inflating the percentage of targeted attacks against this sector.

Financial services was the second most frequently targeted sector in the first half of 2006. Symantec believes that attackers are increasingly motivated by financial gain; as such, the financial services industry is a logical target for attackers hoping to profit from attack activity. Symantec expects that attacks targeted against the financial services industry will continue to rise as attackers become more profit driven.

The sector most frequently targeted by DoS attacks in the first half of 2006 was the Internet service provider (ISP) sector,[6] which was targeted by 38% of all DoS attacks (table 2). ISPs are popular targets for several reasons. Firstly, they are responsible for providing Internet service to a high number of users. By

[6] The Internet service provider sector is made up of organizations whose primary function is providing Internet as a service.

successfully attacking an ISP, an attacker can effectively create denial of service conditions for a high number of users at one time. Secondly, ISPs also host Web sites and provide Internet access to many potential target organizations. Attackers wanting to target an organization's Web site or networks could do so by targeting the organization's ISP.

Rank	Sector	Proportion of attacks
1	Internet Service Provider	38%
2	Government	32%
3	Telecommunications	8%
4	Transportation	4%
5	Education	3%
6	Accounting	3%
7	Utilities / Energy	3%
8	Insurance	3%
9	Financial Services	2%
10	Information Technology	2%

Table 2. Top sectors targeted by denial of service attacks
Source: Symantec Corporation

Organizations should ensure that a documented procedure exists for responding to DoS events. One of the best ways to mitigate a DoS attack is to filter upstream of the target. For most organizations, this filtering will involve working in conjunction with their Internet service provider (ISP). Symantec also recommends that organizations perform egress filtering on all outbound traffic.[7]

Total number of vulnerabilities disclosed

Symantec documented 2,249 new vulnerabilities in the first half of 2006. This is an increase of 18% over the 1,912 vulnerabilities that were documented in the second half of 2005. It is also a 20% increase over the 1,874 vulnerabilities that were reported in the first half of 2005. Symantec documented a higher volume of vulnerabilities in this reporting period than in any other previous six-month period.[8]

The marked increase in the number of vulnerabilities can be attributed to the continued growth in those that affect Web applications. Web applications are technologies that rely on a Web browser for their user interface, rely on HTTP as the transport protocol, and reside on Web servers. Vulnerabilities affecting Web applications accounted for 69% of all vulnerabilities that were documented by Symantec in the first half of 2006. This is a slight increase over the 68% disclosed in the second half of 2005. It is also a nine percentage point increase over the 60% documented in the first half of 2005.

The high number of these vulnerabilities is due in part to the popularity of Web applications and to the relative ease of discovering vulnerabilities in Web applications compared to other platforms. Web applications are required to accept and interpret input from many different sources, and there are very few restrictions to distinguish valid input from invalid. This is further complicated because Web browsers, the application through which most Web applications operate, are very liberal in what they will accept and interpret as valid input.

[7] Egress traffic refers to traffic that is leaving a network, bound for the Internet or another network.
[8] The *Internet Security Threat Report* has been tracking vulnerabilities in six-month periods since January 2002.

Additionally, Web applications generally have quicker release cycles than traditional desktop and server applications. This provides security researchers with a continually growing source of new applications to audit, particularly as, in many cases, Web applications do not undergo the same degree of quality assurance and testing as other applications.

Web applications also present relatively easy targets. This is because the source code is often readily available to be audited (although in many cases security researchers can also quickly discover vulnerabilities on live Web sites). Compared to other types of applications, researchers can often find many more vulnerabilities in Web applications in a shorter period of time. For instance, Web applications are often susceptible to common types of input validation vulnerabilities, such as cross-site scripting and SQL injection, that are typically easy to discover with a minimal amount of effort and skill.[9]

Symantec recommends that administrators employ a good asset management system or vulnerability alerting service and management system, both of which can help to quickly assess whether a new vulnerability is a viable threat or not. Enterprises should devote sufficient resources to alerting and patch deployment solutions. They should also consider engaging a managed security service provider to assist them in monitoring their networks. Administrators should also monitor vulnerability mailing lists and security Web sites for new developments in vulnerability research.

In order to protect against the exploitation of Web application vulnerabilities, organizations should 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 Secure Development Lifecycle and threat modeling.[10] Symantec recommends the use of secure shared components that have been audited for common Web application vulnerabilities to limit the risk of introducing new vulnerabilities when implementing features from scratch. 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.

Window of exposure, enterprise vendors

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 likely have no official recourse against a vulnerability and instead will have to resort to best practices and workarounds to reduce the risk of successful compromise.

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 of this metric 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 (these form the majority), only exploits for technologies from enterprise vendors (that is, those that generally have widespread deployment) are included.[11]

[9] Cross-site scripting is a vulnerability that allows attackers to inject hostile HTML and script code into the browser session of a Web application user. SQL injection is a vulnerability that can affect Web applications, allowing an attacker to inject their own SQL code into a database query that is made by the vulnerable application.
[10] 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.
[11] Vendors included in this metric are: Microsoft, Sun™, HP®, Symantec/VERITAS, EMC, IBM®, Cisco®, Oracle®, CA™ (Computer Associates), and McAfee®.

In the first six months of 2006, the average patch development time for software developed by enterprise vendors was 31 days. The average exploit code development time during the same period was three days. As a result, the window of exposure for this reporting period was 28 days (figure 2). In the second half of 2005, the window of exposure was 50 days. In the first half of 2005, it was 60 days.

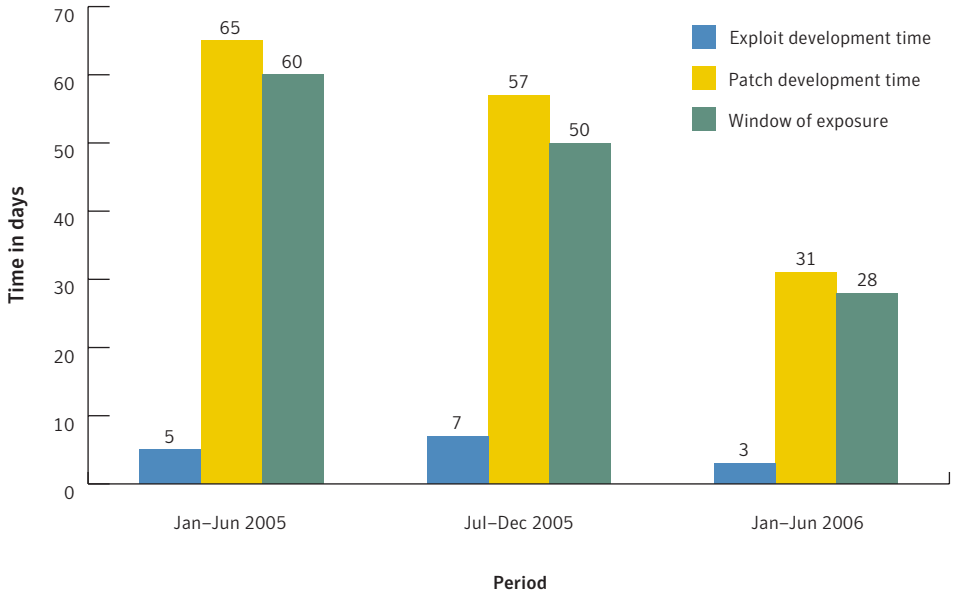

Figure 2. Window of exposure, enterprise vendors
Source: Symantec Corporation

The window of exposure for vulnerabilities in applications developed by enterprise vendors is thus narrowing. While there has been a slight reduction in exploit code development time, the main reason for this narrowing is that patch development time has dropped significantly.

Exploit code for enterprise-vendor vulnerabilities is still being released quickly, forcing administrators to respond rapidly despite a lack of vendor-supplied remediation. However, the decreasing patch development time indicates that enterprise vendors are responding more quickly to vulnerabilities. Despite this, it is critical that organizations follow up with installation of patches.

To minimize the possibility of successful exploitation, administrators need to understand the vulnerabilities and be active in working around them. This may involve making changes to firewall configurations, creating or obtaining IDS/IPS signatures and rules, and locking down services. Administrators should monitor vulnerability mailing lists and security Web sites for new developments in vulnerability research. They should also monitor mailing lists devoted to discussion of security incidents or specific technologies, on which prevention and mitigation strategies may be discussed.

Patch development time, 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.

During the first six months of 2006, Microsoft had an average patch development time of 13 days (figure 3), a significant decrease from 34 days in the last half of 2005. Red Hat also had an average patch development time of 13 days for the first six months of 2006, a drop from 28 days in the last half of 2005. Apple had the third shortest time to patch at 37 days. This is a significant reduction from the 73-day average for 27 vulnerabilities in the second half of 2005.

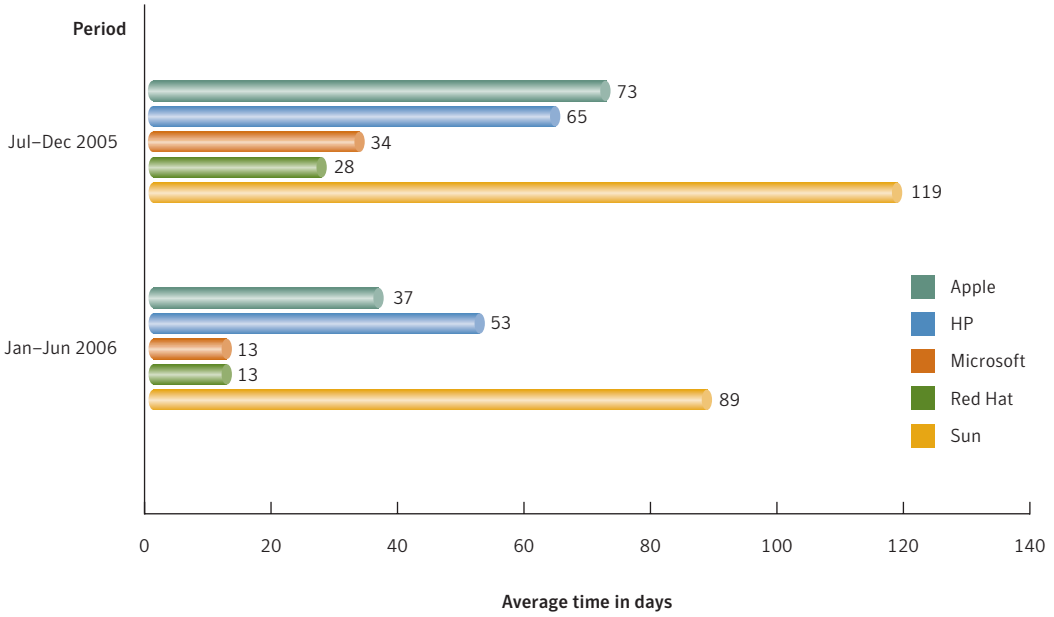

Figure 3. Operating system patch development time
Source: Symantec Corporation

Over the past three reporting periods, Microsoft has had the shortest patch development time of all the operating system vendors. There are many reasons that consumer-oriented vendors such as Microsoft and Apple have lower patch development times than some of the other vendors. Threats to desktop users and consumers generally carry a higher public profile and so there is likely more public pressure for vendors to be responsive and accountable.

Along with Microsoft, Red Hat also had the lowest patch development time during this reporting period. This is likely related to open-source collaboration. If a vendor or a member of the open-source community provides a patch, other vendors can share that patch and incorporate it into their distribution. Linux patches are not released on a fixed schedule; instead, they are often released on a daily basis. This approach differs from Microsoft and Apple, both of whom release their patches less frequently and in large batches to address as many vulnerabilities as possible at a time.

Window of exposure, 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 likely have no official recourse against a vulnerability and instead will have to resort to best practices and workarounds to reduce the risk of successful compromise.

For the first time, in this volume of the *Internet Security Threat Report*, Symantec is assessing the window of exposure for Web browsers. In the first half of 2006, Internet Explorer had a window of exposure of nine days, down considerably from 25 days in the second half of 2005 (figure 4). Apple Safari had a window of exposure of five days, up from zero days in the second half of 2005.[12]

In the first half of 2006, Opera had a window of exposure of two days, down considerably from 18 days during the second half of 2005. In the first half of this year, Mozilla had a window of exposure of one day. In the second half of 2005, Mozilla had a window of exposure of negative two days, meaning that exploit code in that period was generally released after patches were available.

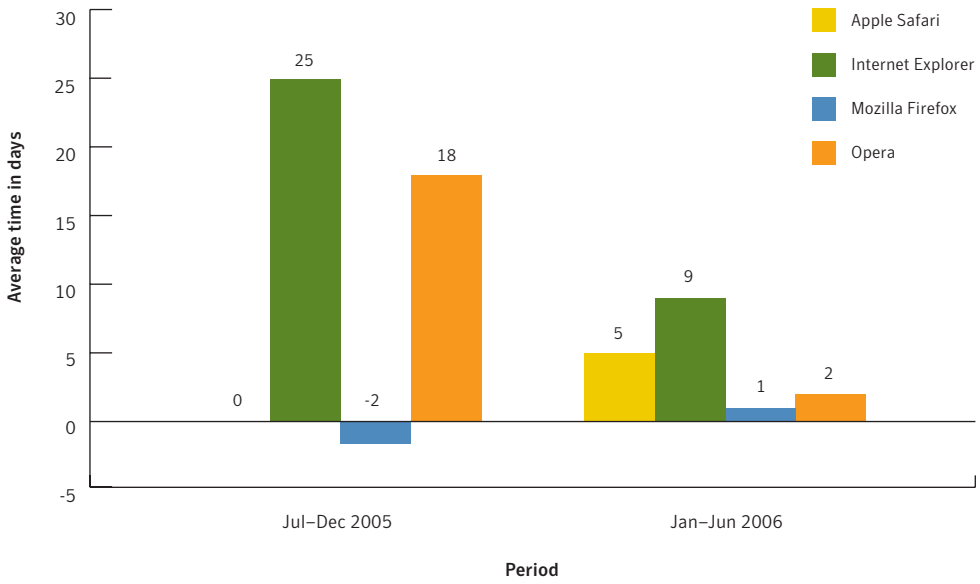

Figure 4. Web browsers window of exposure
Source: Symantec Corporation

In the first half of 2006, the window of exposure for most vendors was smaller than for the second half of 2005. Vendor responsiveness is the key factor in this change, particularly as exploit code development time averages are still very short. It should also be noted that these averages may be influenced by the number of vulnerabilities that are disclosed for each browser.

[12] All patched vulnerabilities affecting Safari in the second half of 2005 were addressed by the vendor at the time of their announcement.

Average patch development times are lower for Web browsers than in other contexts, such as enterprise vendor applications and operating systems. This is noteworthy because some vendors, such as Apple and Microsoft, are included in all of these metrics. This discrepancy may indicate that the patching of browser vulnerabilities is given a higher priority than the patching of other types of vulnerabilities that affect those vendors. This could be attributed to the ubiquity of the Web browser and its high profile as a target for exploitation, which has effectively forced vendors such as Apple and Microsoft to respond more quickly to browser vulnerabilities.

To protect against the exploitation of unpatched vulnerabilities affecting Web browsers, Symantec recommends the deployment of intrusion prevention systems and antivirus at gateways and workstations. Organizations should also closely monitor vulnerability mailing lists and apply necessary patches as required, in a timely manner.

Web browser vulnerabilities

The Web browser is a critical and ubiquitous application that has become a significant target for vulnerability researchers. Traditionally, the focus of security researchers has been on the perimeter: servers, firewalls, and other assets with external exposure. However, a notable shift has occurred, as researchers are increasingly targeting client-side systems, primarily end-user desktop hosts. As part of this shift toward client-side issues, vulnerabilities in Web browsers have become increasingly prominent.

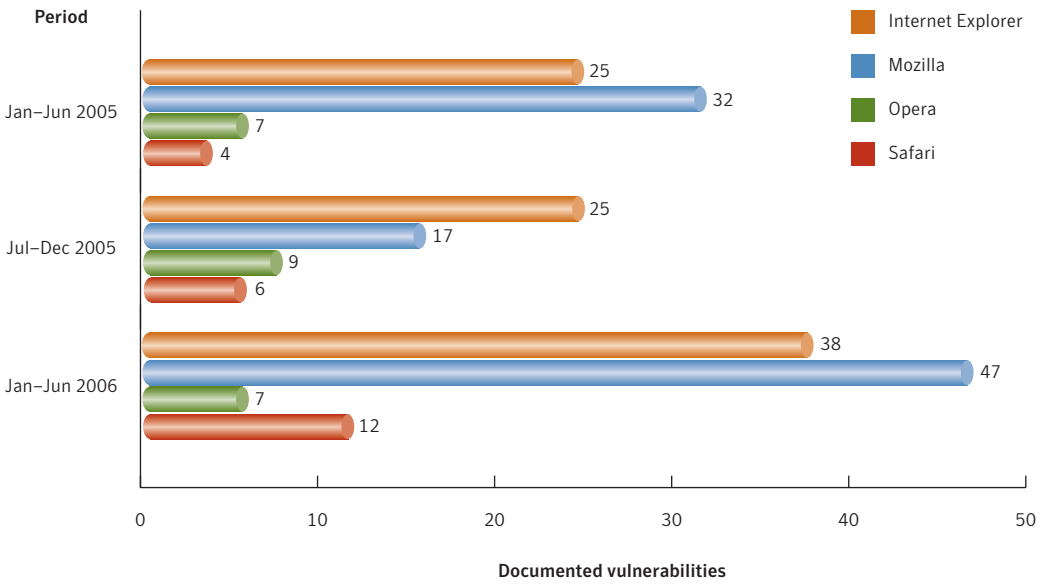

Figure 5. Web browser vulnerabilities
Source: Symantec Corporation

In the first six months of 2006, Symantec documented 47 vulnerabilities that affected Mozilla browsers, including Mozilla Firefox and the Mozilla Browser (figure 5). This is a significant increase over the 17 vulnerabilities that were disclosed in the second half of 2005. The Mozilla Foundation released multiple revisions of Firefox and Mozilla during this period to address the majority of these vulnerabilities.

In the first half of 2006, Symantec documented 38 new vulnerabilities in Microsoft Internet Explorer.[13] This is a 52% increase over the 25 vulnerabilities published in the preceding six-month period. Many of the Internet Explorer vulnerabilities were reported privately to Microsoft and addressed in cumulative security updates over the course of the reporting period. The continued prevalence of Internet Explorer vulnerabilities is likely due to the widespread deployment of the browser.

During this reporting period, 12 vulnerabilities were disclosed that affected Apple Safari. This is double the six reported in the second half of 2005 and triple the four that were disclosed in the first half of 2005. The sharp increase in the number of Apple Safari vulnerabilities over the past 12 months offers evidence that security researchers are increasingly turning their attention to Mac OS X.

Browsers are becoming more complex and feature-rich, which 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 remain a viable attack vector for many client-side vulnerabilities.[14] This is particularly true of Microsoft Windows and other operating systems in which the browser is not disassociated from many other operating system processes and features. This was illustrated by the Excel "zero-day" vulnerability,[15] which Symantec observed in the wild being employed in targeted attacks. The low-key nature of client-side attacks makes them ideal for targeted "zero-day" attacks.[16]

Browser vulnerabilities are a serious security concern, particularly due to their role in online fraud and the propagation of spyware and adware. Organizations should closely monitor vulnerability mailing lists and apply necessary patches as required, in a timely manner. They should also scan their hosts for vulnerable systems to identify hosts that are missing the required patches.

Denial of service attacks

Denial of service (DoS) attacks are a major threat to organizations. A successful DoS attack can render Web sites or other network services inaccessible to customers and employees. This could result in the disruption of organizational communications, a significant loss of revenue, and/or damage to the organization's reputation. Furthermore, as Symantec discussed in a previous *Internet Security Threat Report* (September 2005), criminal extortion schemes based on DoS attacks are becoming more common.[17] During the first six months of 2006, Symantec observed an average of 6,110 DoS attacks per day.

[13] It should be noted that this metric does not include third-party components such as ActiveX components or browser plug-ins; however, if the vendor ships their own ActiveX components or browser plug-ins with the browser, vulnerabilities affecting those components are considered.
[14] Client-side vulnerabilities are those that affect network client applications or that require some degree of user-interaction with data that originates from an external source to be successfully exploited.
[15] http://www.securityfocus.com/bid/18422
[16] 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 public knowledge or known of by the vendor of the affected technology.
[17] Symantec *Internet Security Threat Report*, Volume VIII (September): https://enterprise.symantec.com/enterprise/whitepaper.cfm?id=2238, pp. 11 and 30

The United States was the location of the most DoS targets, accounting for 54% of the worldwide total. The prominence of the United States as a target is not surprising: the country's extensive broadband-Internet infrastructure and the high proportion of Internet-connected organizations situated there make it a very attractive target. China was targeted by the second highest number of DoS attacks, accounting for 12% of the total.

The sector most frequently targeted by DoS attacks in the first half of 2006 was the Internet service provider (ISP) sector, which was targeted by 38% of all DoS attacks. ISPs are popular targets for several reasons. Firstly, they are responsible for providing Internet service to a high number of users. By successfully attacking an ISP, an attacker may be able to effectively create denial of service conditions for a high number of individuals and organizations at one time. Secondly, ISPs also host Web sites and provide Internet access to many potential target organizations. Attackers wanting to target an organization's Web site or networks could do so by targeting the organization's ISP.

Organizations should ensure that a documented procedure exists for responding to DoS events. One of the best ways to mitigate a DoS attack is to filter upstream of the target. For most organizations this filtering will involve working in conjunction with their Internet service provider (ISP). Symantec also recommends that organizations perform egress filtering on all outbound traffic. DoS victims frequently need to engage their upstream ISP to help filter the traffic to mitigate the effects of attacks.

Bot networks

Bot networks are groups of compromised computers on which attackers have installed software that listens for and responds to commands, typically using Internet relay chat (IRC), thereby giving the attacker remote control over the computers. Bots can be used by external attackers to perform DoS attacks against an organization's Web site, which could result in the disruption of organizational communications, a significant loss of revenue, and/or damage to the organization's reputation. Furthermore, bots within an organization's network can be used to attack other organizations' Web sites as well as spread malicious code, both of which can have serious legal and business consequences. Finally, bots can be used by attackers to harvest confidential information from compromised computers.

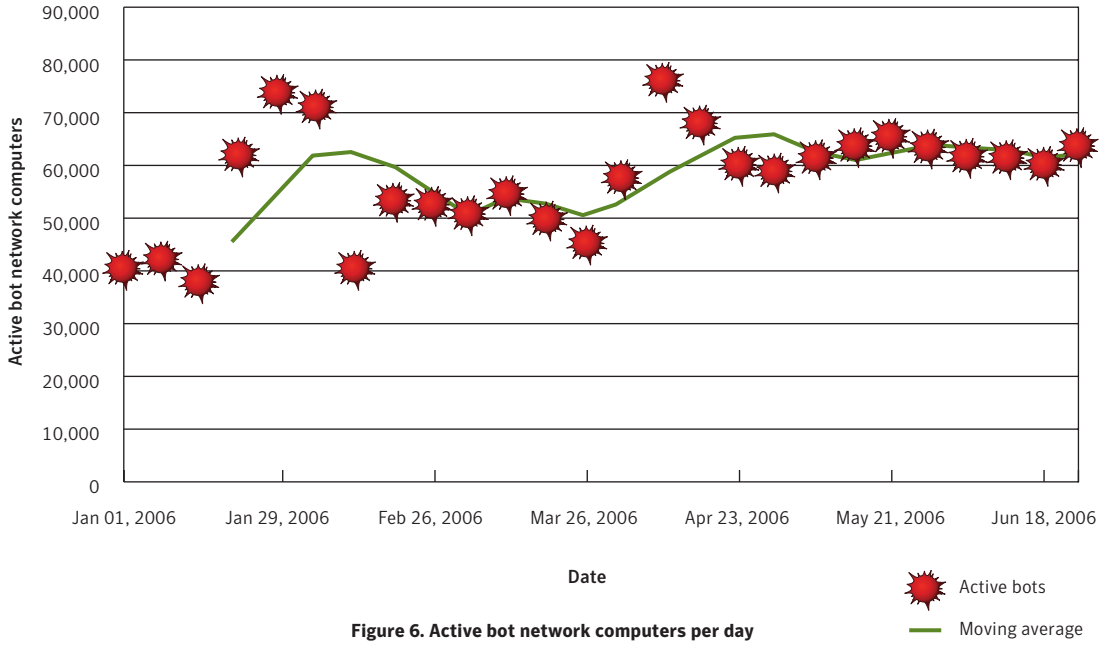

Figure 6. Active bot network computers per day
Source: Symantec Corporation

In the first six months of 2006, Symantec observed an average of 57,717 active bot network computers per day (figure 6). During this period, Symantec observed 4,696,903 distinct bot network computers that were identified as being active at any point in time during the six-month period.

If bots begin to exploit an attack vector that bypasses firewalls and perimeter defenses, the population of bot-infected computers could increase rapidly. This could be particularly dangerous because bot network owners have become more organized and experienced. Furthermore, bot technology is much more established and is more readily available to the public due to the disclosure of bot source code. Finally, some bots and bot networks are reportedly using encrypted channels to communicate, which could make them much more difficult to detect.

Symantec also tracks the number of bot command-and-control servers worldwide. Bot command-and-control servers are computers that bot network owners use to relay commands to bot-infected computers on their bot networks. Symantec identified 6,337 bot command-and-control servers during the first six months of 2006.

The United States had the highest percentage of bot command-and-control computers, with 42% of the worldwide total. 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. The high proportion of bot-infected computers and bot command-and-control servers in the United States is driven by its extensive Internet and technology infrastructure as well as the fact that more than 49 million broadband Internet users are located there.[18]

[18] http://www.oecd.org/document/39/0,2340,en_2649_34449_36459431_1_1_1_1,00.html

China had the highest number of bot-infected computers during the first half of 2006, accounting for 20% of the worldwide total. This ranking represents a rise from third place in the second half of 2005. This coincides with and illustrates a trend that began in 2005, which saw an increase in bot activity in China during that period. Symantec has observed that bots usually infect computers connected to high-speed broadband Internet through large ISPs and that the expansion of broadband connectivity often facilitates the spread of bots.

During the first half of 2006, Beijing was the city with the most bot-infected computers in the world, accounting for almost three percent of the worldwide total. Guangzhou, China ranked second, with just under two percent of the world's bot-infected computers. Seoul, South Korea had the third highest number of bot-infected computers worldwide, accounting for slightly less than two percent of the total. All of the top three cities are large population centers that are cultural and economic centers in their respective countries. Furthermore, all have a large broadband Internet infrastructure.

To prevent bot infections, Symantec recommends that ISPs perform both ingress and egress filtering to block known bot network traffic.[19] They should also filter out potentially malicious email attachments. 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.

End users should employ defense in-depth strategies, including the deployment of antivirus software and a firewall.[20] 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.

Previously unseen malicious code threats

For the first time, in this volume of the *Internet Security Threat Report*, Symantec is tracking the proportion of previously unseen malicious code threats. These are defined as distinct malicious code threats that are detected on Symantec's honeypot computers for the first time before they are detected by other means.[21] This information offers insight into emerging attacker activity, particularly the speed with which attackers are adopting new malicious code tools for use against target computers.

Between January 1 and June 30, 2006, 18% of all distinct malicious code samples detected by the Symantec honeypot had not previously been seen. A high proportion of previously unseen malicious code likely indicates that attackers are more actively attempting to evade detection by signature-based antivirus and intrusion detection systems.

One of the major factors contributing to the increase in previously unseen threats is the number of variants within malicious code families. This indicates that attackers are commonly updating current malicious code to create new variants instead of creating new malicious code "from scratch." This is particularly evident in

[19] 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.
[20] 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.
[21] 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.

the extremely high number of variants in malicious code families such as the Mytob or Beagle families. Attackers and malicious code writers can create new variants in a number of ways, including metamorphic code evolution,[22] changes to the functionality, and run-time packing utilities.[23] The increase in new threats detected during the first six months of 2006 indicates that attackers may be employing these tactics more actively in order to avoid being detected by antivirus software.

Previously unseen threats are particularly dangerous because traditional defenses, such as some signature-based antivirus products, are typically unable to detect them. Administrators should ensure that their networks are protected by perimeter security tools such as intrusion prevention systems (IPS), which will ultimately provide better protection than intrusion detection systems (IDS) or firewalls, neither of which will have rules to protect from previously unseen threats. Organizations should also consider network compliance solutions that will help keep infected mobile users out of the network (and disinfect them before entering). Administrators should also be sure to maintain up-to-date antivirus definitions to ensure that their computers are protected from new threats at the earliest possible time.

Modular malicious code

In the "Future Watch" section of the September 2005 volume of the *Internet Security Threat Report*, Symantec predicted that modular malicious code would become a more prominent security issue in the near future.[24] Modular malicious code works by compromising a computer and then downloading other pieces of code with added functionalities. It initially possesses limited functionality, such as disabling antivirus software and firewalls, but can update itself with code that has new, potentially more damaging capabilities. These may allow it to further compromise the target computer or to perform other malicious tasks.

Modularity in malicious code can serve different purposes. The malicious code may simply attempt to update itself to a more recent version, as is often the case for bots and back door servers. Frequently, modular malicious code is used to download an application that can gather confidential information, which may then be used by attackers for financial gain. By using modular malicious code, attackers may download and simultaneously install a confidential information threat on a large number of compromised computers.

Between January 1 and June 30, 2006, modular malicious code accounted for 79% of the volume of the top 50 malicious code reported to Symantec. This represents a significant decrease from the 88% reported from July to December 2005. The decline in the volume of modular malicious code this period can mainly be attributed to the prevalence of the Blackmal.E worm (also known as the Kama Sutra worm). This worm was the second most widely reported malicious code sample in the current period; however, it did not attempt to download additional components or threats and so is not considered modular. The large volume of Blackmal.E reports thereby caused the overall volume of modular malicious code in the top 50 to decline.

While the volume of modular malicious code has declined since the previous period, the number of modular malicious code samples in the top 50 has remained constant. In both the first half of 2006 and the second half of 2005, 36 unique samples were reported.

[22] Metamorphic code evolution describes a method used by malicious code writers that allows a piece of malicious code to change itself autonomously.
[23] Run-time packing utilities, also known as run-time packers, are traditionally used to make files smaller. Malicious code writers use them to make antivirus detection more difficult. See the "Win32 Viruses and Bots" discussion of the "Malicious Code Trends" section in this document for a more detailed discussion on run-time packers.
[24] Symantec *Internet Security Threat Report*, Volume VIII (September 2005): http://enterprisesecurity.symantec.com/content.cfm?articleid=1539, p. 83

In order to protect against modular malicious code, administrators should implement proxy-based Internet access from inside the organization. This will allow administrators to control which sites can be visited and, in cases of infection, identify hosts that have visited URLs in order to download malicious code updates.

Malicious code threats that expose confidential information

Some malicious code programs are designed specifically to expose confidential information from an infected computer. Threats that expose confidential information 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 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.

In the current period, 30 of the top 50 malicious code samples exposed a user's confidential information in some way. This is the same number as was reported in the previous reporting period but ten more than the 20 reported in the first half of 2005.

Symantec believes that the number of threats to confidential information will likely hold steady or increase in the next six months. In the current period, variants of Mytob accounted for 16 of the 30 information-exposure threats in the top 50 malicious code reports. Bots such as Mytob will likely continue to be common amongst the top 50 reported malicious code samples, as their versatility and modularity make them very popular with attackers.

Trojans

While Trojans dominated the malicious code landscape a year ago, making up 21 of the top 50 malicious code samples, they currently account for only ten of the top 50. They account for less than a quarter of the volume of the top 50 malicious code reported to Symantec during this period.

While some industry observers have claimed that the number of Trojans outnumbers worms and viruses overall, this has not been borne out in the data that Symantec has received from enterprise and consumer customers worldwide. Due to a lack of propagation mechanisms, Trojans are not likely to be seen by as many users or in such high volume as mass-mailing worms.

That said, Trojans are still an important security threat. Five of the top ten new malicious code families reported during the first six months of 2006 were Trojans. Attackers appear to be making a shift toward targeted attacks using these threats. Mass-mailing worms tend to use a "shotgun" approach, sending large quantities of themselves to as many users as possible. However, Trojans are now frequently produced to target specific users and groups. For example, the Mdropper.H Trojan exploited a zero-day vulnerability in Microsoft Word in order to install a variant of the Ginwui back door program.[25] The Word document containing the Mdropper Trojan was spammed to a selected user base using a message with social engineering that was tailored to entice the recipients into opening it. Because of the targeted nature of these attacks, the Trojan was sent to a smaller group of users, making it less conspicuous and less likely to be submitted to antivirus vendors for analysis.

[25] http://www.symantec.com/outbreak/word_exploit.html

To protect against Trojans and to mitigate their effectiveness in the case of infection, users should deploy regularly updated antivirus software and a personal firewall. They should also ensure that all desktop, laptop, and server computers are updated with all necessary security patches from their operating system vendor. They should never view, open, or execute any email attachment unless it is expected and comes from trusted source, and unless the purpose of the attachment is known.

Number of unique phishing messages

Over the first six months of 2006, the Symantec Probe Network detected 157,477 unique phishing messages (figure 7). This is an increase of 81% over the 86,906 unique phishing messages that were detected in the last half of 2005. It is also an increase of 61% over the 97,592 messages detected in the first half of 2005.

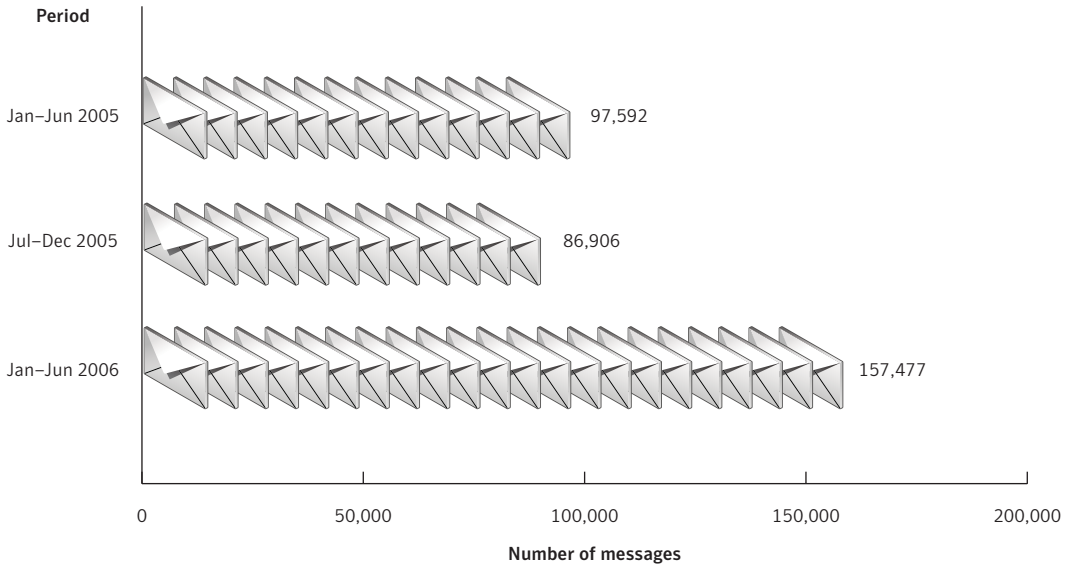

Figure 7. Number of unique phishing messages
Source: Symantec Corporation

This sharp increase over the previous six-month period may be a result of attempts by attackers to bypass filtering technologies by creating multiple randomized messages. These messages attempt to phish the same brands, but include slight variances—such as variations in the URLs included in the phishing message—in order to bypass the use of MD5 checksums or other basic email scanning techniques.[26] Attackers tend to rotate their usage of a particular domain frequently. By using a large number of domains in a short period, they make it more difficult for authorities to shut them down due to the amount of effort involved in tracking and taking down each domain used.

[26] An MD5 checksum is obtained when a message is hashed through an algorithm to obtain a unique value. This technique can be used to identify known spam, phishing, and malicious code email messages.

Phishing activity by sector

For the first time, in this edition of the *Internet Security Threat Report*, Symantec is tracking the sectors of brands being targeting by phishing attacks. Not surprisingly, the financial sector is the most heavily phished, accounting for 84% of phishing sites tracked by the Symantec Phish Report Network and Symantec Brightmail AntiSpam™ this period (figure 8).

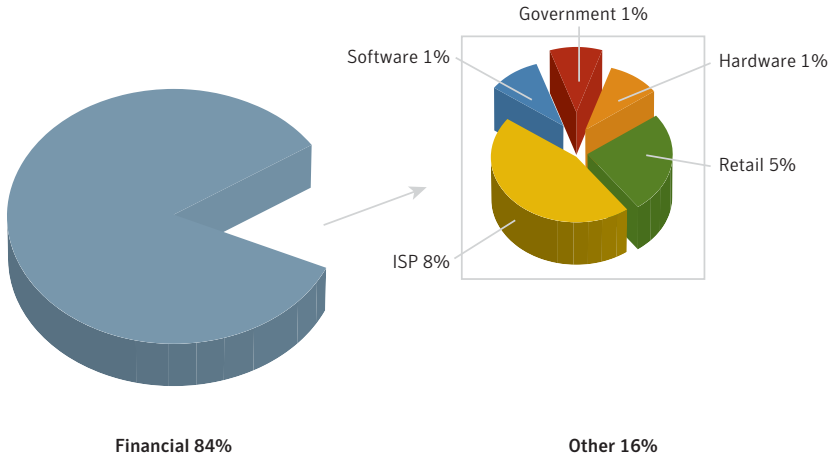

Figure 8. Phishing activity by sector
Source: Symantec Corporation

Phishing attacks against the financial services sector are most likely to produce the greatest monetary gain for attackers. Once an attacker gains access to a target's account through one of these attacks, he or she can initiate wire transfers to remove funds, apply for loans, credit lines, or credit cards. Further evidence of the high concentration of phishing activity targeting the financial sector is the fact that nine of the top ten brands phished this period were from that sector.

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.[27] They should also keep their employees notified of the latest phishing attacks and how to avoid falling victim to them.[28]

Symantec recommends that organizations sign up for a fraud alerting service or 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 site that could be used for phishing. Organizations should also monitor the purchasing of cousin domain names by other entities to identify purchases that could be used to spoof their corporate domains.[29] This can be done with the help of companies that specialize in domain monitoring; some registrars even provide this service.[30]

[27] 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/conline/pubs/alerts/phishingalrt.htm.
[28] A good resource for information on the latest phishing threats can be found at http://www.antiphishing.org.
[29] "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.
[30] See http://markmonitor.com/brandmanagement/index.html for instance.

End users should follow best security practices as outlined in Appendix A of this report. 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.

Six-month volume of spam

Spam is usually defined as junk or unsolicited email sent by a third party. While it is certainly an annoyance to end users and administrators, spam is also a serious security concern as it can be used to deliver Trojans, viruses, and phishing attempts. It can also be sent in sufficient quantities to cause denial of service conditions.

Between January 1 and June 30, 2006, spam made up 54% of all monitored email traffic. This is an increase from the last six months of 2005 when 50% of email was classified as spam. However, it is lower than the first half of 2005, when 61% of email was classified as spam.

In the previous *Internet Security Threat Report* (March 2006), Symantec speculated that the decline of spam detected during that reporting period was due to the implementation of IP filtering, traffic shaping, and ISP policy changes to control spam.[31] The current reversal of this trend indicates that spammers may have found means to circumvent these measures, such as utilizing image-based spam. Since the financial gains from spam are directly related to the number of messages that reach end users, it is in the spammers' best interests to find ways to bypass any defenses that administrators put in place.

Percentage of spam containing malicious code

For the first time, in this volume of the *Internet Security Threat Report*, Symantec will assess the percentage of spam messages containing malicious code. In the first six months of 2006, 0.81% of spam email contained malicious code. This means that one out of every 122 spam messages blocked by Symantec Brightmail AntiSpam during this period contained malicious code. Between January and May, spam containing malicious code dropped steadily before rising again slightly in June. At the beginning of the year, 1.27% of spam email contained malicious code compared to 0.56% at the end of June.

This decline is likely due to two factors. The first is that attaching malicious code to a message increases its chances of being blocked by various means. In some cases, administrators may block all incoming messages with attachments or executable type attachments. Additionally, spam messages with malicious code attachments may be detected by both spam-filtering software and antivirus scanners, decreasing their chances of reaching end users.

The second factor, which is likely a response to the first, is the inclusion of links to Web sites hosting malicious code in spam messages. Rather than attach a malicious code executable to a message, spammers will instead include a link to a Web site that is hosting malicious code. In many cases, the Web site may exploit a client-side vulnerability in the user's browser to install the malicious code

[31] Symantec *Internet Security Threat Report*, Volume VIII (September 2005): http://enterprisesecurity.symantec.com/content.cfm?articleid=1539, pp. 15 and 97.

without the user's knowledge or consent. This technique helps reduce the number of messages that are blocked before reaching the end user and still allows spammers to have their malicious code installed on a recipient's computer.

To protect against malicious code received through spam, users should follow the same precautions used to protect against any malicious code infections, such as employing defense in-depth strategies, including the deployment of antivirus software and a personal 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. 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.

Misleading applications

Misleading applications give false or exaggerated reports of security threats on a user's system in order to persuade users to pay money to purchase software or upgrade to a version of security software that will purportedly remove the "threats" that were found. They are an example of a security risk that uses social engineering to achieve its end.

Misleading applications pose a risk to the user's security in a number of ways. Firstly, the consumer may be unprotected against a wide variety of threats that the fraudulent security product claims to protect them against. A bad or misleading security product can be worse than no security at all because of the illusory sense of protection provided. Secondly, in paying for the fraudulent security product, the end user may disclose personal information, such as banking or credit card information, that could then be used for criminal purposes, such as identity theft or credit card fraud.

During the first six months of 2006, three of the top ten new security risks were misleading applications. They accounted for 50% of the volume of reports for the top ten new security risks in the first half of 2006. Two of the top three reported security risks during this period were misleading applications, including ErrorSafe,[32] which accounted for 30% of submissions of the top ten new security risks.

In order to mitigate the threat posed by misleading applications, Symantec recommends that administrators and users follow the recommended best practices outlined in Appendix A of this report, and exercise caution when installing applications that purport to solve security issues. Enterprises should only install applications that have been reviewed and certified as legitimate applications. Any application should only be deployed as part of an approved security policy.

[32] http://www.symantec.com/security_response/writeup.jsp?docid=2006-012017-0346-99&tabid=1

Future Watch

The previous sections of the *Internet Security Threat Report* have discussed some of the Internet security developments that Symantec observed between January 1 and June 30, 2006. This section 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 with an opportunity to prepare themselves for rapidly evolving and complex security issues.

Increased polymorphism in Win32 malicious code

In the first half of 2006, Symantec Security Response noticed a renewed interest in polymorphic viruses. A polymorphic virus is one that can change its byte pattern when it replicates, thereby avoiding detection by simple string-scanning antivirus techniques. In essence, polymorphic viruses make changes to their code to avoid detection.

Polymorphic and self-mutating viruses appear to be enjoying renewed popularity. Over the past several years, malicious code authors have been developing increasingly sophisticated malicious code that employs these techniques. However, due to the difficulty in implementing these techniques, malicious code authors focused their efforts on developing run-time packers as a means of mass propagation,[33] as opposed to sophisticated malicious code that avoids detection.

With the success of large-scale worms such as Nimda and Code Red and viruses like I Love You, it became apparent that malicious code did not need to be sophisticated in order to infect large numbers of machines. Today, there is an increased focus on targeted attacks and more subtle infection methods. As a result, attackers are increasingly using polymorphic techniques to avoid detection and aid in propagation.

As noted in the "Malicious Code Trends" section of this report, improved unpacking support has reduced the effectiveness of using run-time packers to obfuscate malicious code. As a result, malicious code authors have been forced to employ different means to avoid detection while infecting host systems. This may be the main factor in the emergence of polymorphic malicious code activity observed by Symantec.

During March and April of 2006, a worldwide outbreak of two viruses, Polip[34] and Detnat,[35] signified that polymorphic viruses may be regaining prominence. Detection of complex polymorphic viruses is much more technologically dependent than other types of malicious code. It involves a complex process of cryptological and statistical analysis along with code emulation and data-driven engine designs. It therefore requires experienced analysts to develop detection and removal techniques.

As Polip and Detnat showed, security and antivirus vendors may have difficulty in detecting and protecting against these threats. Symantec has increasingly observed the use of polymorphic techniques in packers, which could lead to increasingly sophisticated and potentially more damaging malicious code being circulated worldwide. This is of particular concern to organizations, as Symantec has observed an increase in the number of attacks and malicious code specifically designed to target specific, individual organizations.

[33] Run-time packers are compression routines that allow an executable file to run even though they are compressed.
[34] http://securityresponse.symantec.com/security_response/writeup.jsp?docid=2006-042309-1842-99
[35] http://securityresponse.symantec.com/security_response/writeup.jsp?docid=2006-032912-3047-99

Due to the difficulty in detecting and removing polymorphic viruses, Symantec speculates that more malicious code authors may begin to use polymorphic techniques at all levels of malicious code development. For enterprises, this could result in increased volumes of targeted malicious code from which they have limited protection. Should more malicious code use these techniques, targeted organizations may be increasingly at risk, as obtaining samples to develop detection signatures will likely be difficult.

To protect against these anticipated threats, Symantec recommends that organizations deploy intrusion prevention solutions on host systems and ensure that their antivirus solutions are able to detect and protect against these types of threats. They should also always ensure that their antivirus definitions are up-to-date.

Web 2.0 security threats and AJAX attacks expected to increase

Web 2.0 is term used to describe Web application technologies and Web sites that use the Internet in a collaborative way to provide services to its users. Web 2.0 technologies rely in large part on the user-as-publisher model of interaction and allow for user-created content to be developed and implemented by large groups of individuals.

Web 2.0 technologies present a number of areas for security concern. Because individuals are able to create and host content on various collaboration platforms such as Weblogs, the possibility exists for those platforms to host exploits and become distribution points for links to fraudulent Web sites, malicious code, and other security threats, such as spyware. Attackers will often take advantage of the implied trust between the community of individual developers and the sites hosting content to compromise individual users and/or Web sites.

Additionally, Web 2.0 technologies rely heavily upon Web services. Web services are services that are designed to support interoperability between hosts over a network. Symantec has already observed one worm that used the Google Search Web service.[36] This attack provided evidence that well-known services are not immune to these sorts of attacks and that the number of users these services have present an attractive opportunity to maximize attackers' efforts. Symantec is concerned that in the rush to develop Web services, the underlying Web applications that use them are not receiving the same level of security auditing as traditional client-based applications and services.

The last several *Internet Security Threat Reports* have highlighted the high percentage of Web application vulnerabilities reported to Symantec. AJAX, which is short for asynchronous JavaScript and XML, is an interactive Web development technique used in Web applications to create a more seamless user experience by exchanging small amounts of data between various Web services used in Web applications. It does this without the knowledge of the user initiating the request in order to present a quicker and smoother user experience, much like a desktop application.

As Web applications continue to gain in popularity, Symantec expects to see an increase in the number of attacks taking advantage of the interconnected, interactive nature of AJAX to increase the number of potential targets. Whereas traditional client-server models process the majority of requests on the server side, AJAX allows for a larger portion of those requests to be processed on the client side. The net result is

[36] http://www.symantec.com/security_response/writeup.jsp?docid=2004-122109-4444-99

that attackers have a greater ability to examine Web application code and, in turn, to develop attacks by exploiting newly exposed Web applications through malicious AJAX applications.

Because AJAX can be used in conjunction with a large number of Web services and enables connectivity between them, this could present additional attack vectors into which attackers could inject hostile content. The potential also exists in AJAX for attackers to exploit the trust relationship inherent in the client-server model utilized in Web applications by creating exploits hosted by malicious Web services that steal poorly stored state or login information on the client-side.

One example of this is cross-site scripting. Cross-site scripting attacks take place when Web applications gather data from a user or other source and then create an output of that data on a user's Web browser. Not only could this allow an attacker to steal confidential information, it could also allow an attacker to insert malicious code onto the host through malicious scripts.

The combination of increased attack surfaces with greater cross-site scripting and content injection attacks unique to AJAX has the potential to expose users to a larger number of attacks that would escape the notice of traditional security solutions. As most servers interpret AJAX-submitted data as valid requests, it is nearly impossible to determine what constitutes malicious activity.

In addition to following the best practices outlined in Appendix A of this report, Symantec recommends that enterprises ensure that access to Web applications, whether in-house or external, are limited to approved applications only. Symantec also recommends that before any Web service or application is implemented, it undergo a secure code audit to ensure that it is not vulnerable to possible attack.

Microsoft Windows Vista™

Microsoft's latest desktop operating system, Windows Vista™, is scheduled for release in 2007. With the various Windows operating systems deployed on an estimated 90% of desktop systems around the world, the introduction of Windows Vista is expected to present new and unique security concerns. In anticipation of its release, security researchers have begun releasing preliminary security analyses of beta versions of the operating system that have been made available to the development community.[37]

The preponderance of malicious code, vulnerabilities, and attacks targeting Microsoft Windows today is due in large part to the widespread deployment of the Microsoft operating systems. In response to this, Microsoft has devoted substantial resources to securing Windows Vista. With the release of Vista, Microsoft has rewritten significant portions of its code base and has performed ongoing exhaustive source code audits in hopes of addressing many of the security issues that are present in its current products. Microsoft believes that the combination of new security technologies that have been (and are still being) integrated into Windows Vista will dramatically decrease its potential susceptibility to attack.

Symantec expects to see a concerted effort by the research community to discover and document shortcomings in Windows Vista as attackers attempt to circumvent these new technologies. If successfully implemented, these new technologies may play a role in decreasing the overall volume of malicious code threats affecting the Windows platform. However, it is not yet clear to what degree they will succeed.

[37] For example, please see "Windows Vista Network Attack Surface Analysis" available at http://www.symantec.com/enterprise/security_response/whitepapers.jsp

Research to date has shown increased risk in some of the new technologies, such as the new Vista network stack,[38] while others can still be disabled or bypassed by attackers, such as driver signing and PatchGuard.[39]

Symantec speculates that the new features and changes to Windows Vista's code base, in conjunction with increased scrutiny from security researchers and malicious code authors, will result in previously unseen attacks. Organizations considering a move to Windows Vista will need to plan their migration carefully. Symantec recommends that they do so first in small, non-critical environments, and that they conduct thorough security audits to reduce possible exposure to attack. Based on currently available research, Symantec suggests that until its public release, Windows Vista should be deployed only in an isolated lab environment.

Increase in number of vulnerabilities reported due to the use of fuzzers

In Volume VIII of the *Internet Security Threat Report* (September 2005), Symantec speculated that the number of newly discovered vulnerabilities would increase due to the use of sophisticated tools for decompiling and analyzing software.[40] As discussed in the "Vulnerability Trends" section of this report, the number of vulnerabilities reported has continued to increase over the past several reporting periods.

The introduction of fuzzers—programs or scripts that are designed to find vulnerabilities in software code or scripts—has automated many of the code auditing tasks that security researchers had previously done manually. When fuzzers are combined with modern debugging and disassembly tools, more vulnerabilities can be discovered and analyzed in less time, resulting in more products receiving increased scrutiny.

These advances in code auditing reduce both the time and the effort involved in finding new vulnerabilities. The use of fuzzers does not require advanced security skills; as a result, the number of amateur researchers discovering security flaws has increased. This may result in a larger number of vulnerabilities being reported that may not be exploitable while lowering the overall quality of vulnerability research. Symantec speculates that if this trend continues, security administrators may pay less attention to vulnerability research and, by doing so, leave their systems susceptible to exploitation of high-severity vulnerabilities.

The advent of fuzzing as a mainstream security research technique will increase the number of vulnerabilities reported. As a result, organizations may well be forced to deal with an increased number of vulnerabilities in the software and technologies deployed in their environments. They may therefore need to devote more resources to vulnerability management.

Furthermore, fuzzing tools have the potential to discover new vulnerabilities in heavily audited programs in which fewer vulnerabilities have been found using traditional source code and binary analysis techniques. Symantec speculates that this may result in more vulnerabilities being discovered in technologies and software previously thought to be patched or secure. In addition to following the best practices outlined in Appendix A of this report, organizations should develop a vulnerability management process that alerts security administrators to potential issues.

[38] http://www.symantec.com/avcenter/reference/ATR-VistaAttackSurface.pdf
[39] http://www.symantec.com/avcenter/reference/Windows_Vista_Kernel_Mode_Security.pdf
[40] Symantec *Internet Security Threat Report*, Volume VIII (September 2005): http://enterprisesecurity.symantec.com/content.cfm?articleid=1539, p. 88

Attack Trends

This section of the *Internet Security Threat Report* will provide an analysis of attack activity that Symantec observed during the period between January 1 and June 30, 2006. An attack is defined as any malicious activity carried out over a network that has been detected by an IDS or firewall. An attack is typically an attempt to exploit a vulnerability in software or hardware. When applicable, attack activity for this period will be compared to that presented in the previous *Internet Security Threat Report*.[41] However, in some cases methodological changes that have been implemented for this reporting period may preclude valid comparison. Wherever appropriate, suggestions for attack remediation will be made along with references to Symantec's best practices, which are outlined in Appendix A of this report.

The Symantec™ Global Intelligence Network monitors attack activity 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. In addition to these sources, Symantec has developed and deployed a honeypot network that is used to identify, observe, and study complete instances of attacker and malicious code activity. This data can help to provide details about how some of the attack activity identified in this section is carried out. 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.

Security devices can monitor for attacks and suspicious behavior at many different levels on the network. Devices such as IDS/IPS, firewalls, proxy filters, and antivirus installations all contribute to the security of an organization. Symantec gathers data from many of these devices. One consequence of this data-gathering scheme is that malicious code data and attack trend data often address the same activity in different ways. For instance, attack trends data is based on the number of infected sources that are attempting to spread through network-based attacks. On the other hand, malicious code data is based primarily on reports of attempted propagation, which includes network-based attacks as well as other methods such as mass mailing. This can lead to different rankings for the same threats presented in the "Attack Trends" and "Malicious Code Trends" sections of this report.

This section of the *Internet Security Threat Report* will discuss:

- Top Web browser attacks
- Distribution of Web browser attacks
- Top wireless threats
- Denial of service attacks
- Top countries targeted by denial of service attacks
- Top sectors targeted by denial of service attacks
- Bot networks
- Bot-infected computers by country
- Bot-infected computers by city
- Top originating countries
- Top targeted sectors

[41] Symantec *Internet Security Threat Report*, Volume VIII (September 2005) and Volume IX (March 2006). Both are available at: http://enterprisesecurity.symantec.com/content.cfm?articleid=1539

Top Web browser attacks

For the first time, in this volume of the *Internet Security Threat Report*, Symantec is tracking the top attacks carried out against Web browsers. Attacks targeting Web browsers can be launched from malicious Web sites, bulletin board sites, legitimate Web sites that have been compromised and, in some cases, through malicious emails. Since Web browser attacks are carried out over network ports that are traditionally not filtered, they represent a serious threat because they cannot be blocked using some traditional perimeter security tools such as firewalls. It should be noted that attacks discussed in this metric include only those that specifically target Web browsers and not libraries or other applications that are accessible through Web browsers.

Successful attacks targeting Web browsers typically allow an attacker to compromise a vulnerable computer and gain the privileges of the user currently logged into the computer. For example, if a user with administrator privileges was running the compromised browser, the attacker could gain administrator privileges.

Once a successful Web browser attack is carried out, the attacker could use the compromised computer as a platform from which to launch further attacks against the network from behind perimeter security defenses. This could allow an attacker to launch attacks against other computers on the network or eavesdrop on internal network traffic. It could also allow an attacker to carry out further Web browser attacks by using DNS manipulation tools and Web servers on the compromised computer to redirect users on the network to a malicious Web site.

The attacks discussed in this section are the most common Web browser attacks detected by the Symantec Global Intelligence Network during the first six months of 2006. They are determined by the percentage of all attacking IP addresses that perform a given Web browser attack. These attacks reflect those carried out on the Internet as a whole and are thus indicative of activity that security administrators are likely to observe on their own networks.

The most common attack carried out against Web browsers between January 1 and June 30, 2006 was the Multiple Browser Zero Width GIF Image Memory Corruption Attack, which accounted for 31% of all detected Web browser attacks (table 3). This attack exploits the vulnerability of the same name,[42] which was first disclosed in September 2002 and affects older Netscape, Mozilla, Galleon, and Opera Web browsers.[43] This attack is carried out when a user loads a Web site containing a graphics interchange format (GIF) image file with a width field that is set to zero.

[42] http://www.securityfocus.com/bid/5665
[43] This issue affects Netscape versions 6.2 through 6.2.3, Mozilla Browser versions 0.9.5 through 1.0, and Galleon Browser versions 1.2.4 through 1.2.6 and 5.1.2 through 6.0.1.

Jan–June 2006 rank	Attack	Jan–June 2006 percent of attackers
1	Multiple Browser Zero Width GIF Image Memory Corruption Attack	31%
2	Microsoft Internet Explorer DHTML Object Race Condition Memory Corruption Attack	19%
3	Microsoft Internet Explorer Remote URLMON.DLL Buffer Overflow Attack	17%
4	Mozilla JavaScript URL Host Spoofing Arbitrary Cookie Access Attack	8%
5	Mozilla Browser BMP Image Decoding Multiple Integer Overflow Attack	7%
6	Microsoft Internet Explorer Bitmap Processing Integer Overflow Attack	3%
7	Mozilla Browser Non-ASCII Hostname Heap Overflow Attack	3%
8	Microsoft Internet Explorer Drag and Drop Attack	3%
9	Mozilla Multiple URI Processing Heap Based Buffer Overflow Attack	2%
10	Microsoft Internet Explorer HTML Document Directive Buffer Overflow Attack	2%

Table 3. Top Web browser attacks
Source: Symantec Corporation

One reason that this attack is so common is that the affected vulnerability is relatively easy to exploit. A publicly available proof-of-concept exploit has been available for some time, allowing attackers to use and distribute malicious image files. Furthermore, using bulletin board and Webmail software, it is very easy for attackers to distribute malicious files to targeted users. Unlike browser attacks that require the browser to load malicious HTML or script code (which is typically filtered and restricted by bulletin board and Webmail software), this attack requires only that a user attempt to view a malicious GIF image.

The second most common attack targeting Web browsers during the first half of 2006 was the Microsoft Internet Explorer DHTML Object Race Condition Memory Corruption Attack, which was used by 19% of all attacking IP addresses. This attack is carried out against the vulnerability of the same name,[44] which was first disclosed in April 2005. The third most popular Web browser attack in the first half of 2006 was the Microsoft Internet Explorer Remote URLMON.DLL Buffer Overflow Attack. It was used by 17% of all detected attacking IP addresses. This attack targets the vulnerability of the same name,[45] which was first disclosed in April 2003. Both of these attacks are likely prominent because of the availability of long-standing publicly available exploit code and the widespread deployment of the Microsoft Internet Explorer Web browser.

In order to protect against Web browser attacks, Symantec advises users and administrators to upgrade all browsers to the latest, patched versions. To reduce exposure to attacks, Symantec recommends that organizations educate users to be extremely cautious about visiting unknown or untrusted Web sites or

[44] http://www.securityfocus.com/bid/13120
[45] http://www.securityfocus.com/bid/7419

viewing or following links in unsolicited emails. Administrators should also deploy Web proxies in order to block potentially malicious script code and implement ActiveX controls in order to stop attacks before they can be carried out.

Distribution of Web browser attacks

This metric will assess the number of attacks that targeted each Web browser during the first six months of 2006. The purpose of this metric is to assess which browsers attackers are targeting the most frequently so that administrators whose networks deploy them can take the appropriate protective measures. The browsers included in this analysis are Microsoft Internet Explorer, Apple Safari, the Mozilla family (including Firefox and the Mozilla browser), Opera, Netscape, and KDE Konqueror.

During the first six months of 2006, Microsoft Internet Explorer was the most frequently targeted Web browser. It was targeted by 47% of all known attacking IP addresses (figure 9). The prominence of Microsoft Internet Explorer is not surprising considering the number of vulnerabilities that affect it. Furthermore, on a worldwide basis, it is the most widely deployed browser.

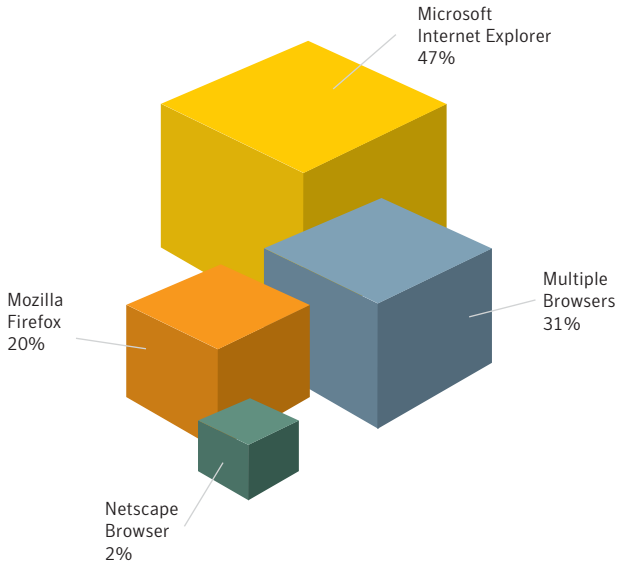

Figure 9. Distribution of attacks targeting Web browsers
Source: Symantec Corporation

Some attacks target vulnerabilities that are present in multiple Web browsers. These vulnerabilities are typically present in numerous browsers because of shared source code, although this is not always the case. Browsers that fall within the "multiple browsers" category include Apple Safari, KDE Konqueror, the Mozilla Browser family, Netscape, Opera, Microsoft Internet Explorer and others.[46] Attacks targeting multiple browsers were the second most common during the first half of 2006, accounting for 31% of all attacking IP addresses.

[46] Safari, Konqueror, and Opera are included in the multiple browsers category but are not listed individually because there were no detected attacks that targeted them as individual browsers.

Mozilla Firefox was the Web browser targeted by the third highest number of detected Web browser attacks during the first half of 2006. Twenty percent of all attacking IP addresses targeted Firefox during this period.

In order to protect against Web browser attacks, Symantec advises users and administrators to upgrade all browsers to the latest, patched versions. To reduce exposure to attacks, Symantec recommends that organizations educate users to be extremely cautious about visiting unknown or untrusted Web sites or viewing or following links in unsolicited emails. Administrators should also deploy Web proxies in order to block potentially malicious script code and implement ActiveX controls in order to stop attacks before they can be carried out.

Top wireless threats

For the first time, in this volume of the *Internet Security Threat Report*, Symantec is tracking threats against wireless network infrastructures. Although many of the security issues facing wireless networks are similar to traditional wired networks, there are some that specifically affect the former. Many new wireless network implementations are insecure, leaving them susceptible to attacks that could allow attackers to steal bandwidth, view or manipulate confidential information, or use compromised networks to carry out further attacks.

The wireless threats discussed in this section are those that Symantec detected being launched against a sample of wireless local area networks between January 1 and June 30, 2006. This discussion will only deal with threats against the wireless network infrastructure itself, it will not deal with attacks against computers deployed on wireless networks.

Current rank	Threat	Proportion of total threats
1	Device Probing for an Access Point	30%
2	Spoofed MAC Address	17%
3	Unauthorized NetStumbler Client	16%
4	Rogue Wireless Access Point	8%
5	Unauthentication Association Denial of Service Attack	6%
6	Radio Frequency Jamming Denial of Service Attack	4%
7	CTS Flood Denial of Service Attack	3%
8	Illegal 802.11 Packet	2%
9	Potential Honeypot Access Point	2%
10	Authentication Flood Denial of Service Attack	2%

Table 4. Top attacks against wireless networks
Source: Symantec Corporation

The most common wireless threat detected between January 1 and June 30, 2006 was a device probing for an access point, which accounted for 30% of all threatening activity (table 4). A device probing for a wireless network access point is one that is noisily trying to connect with an access point using any service set identifier (SSID).[47]

[47] A Service Set Identifier (SSID) is the public name of a wireless network. It is used to allow computer nodes to differentiate one wireless network from another.

An organization's wireless security can be threatened by devices probing for an access point in two ways. The first is by attackers roaming urban areas attempting to locate and connect to wireless networks, a practice that is known as war driving.[48] Attackers carrying out this type of attack often do so using wireless network tools. This type of activity may lead to the compromise of an organization's internal network, potentially giving attackers access to sensitive information through open network shares or by eavesdropping on network communications. It may also give attackers a platform from which to carry out further attacks against other targets anonymously.

The second way in which an organization can be threatened by devices probing for wireless access points is through authorized, albeit poorly configured, computers trying to connect to an access point using any SSID. Although apparently innocuous, this could be more damaging to an organization than war driving.

An attacker may take advantage of authorized computers that are probing for access points in two different ways. The first is through the use of a rogue access point.[49] In this case, an authorized computer probing for an access point might connect to the rogue access point, allowing the attacker to eavesdrop on wireless network communications and potentially gain access to sensitive information.

The second is by using the authorized computer to gain access to the organization's wired local network. It is not uncommon for computers with wireless network hardware to also be connected to a wired network within an office environment. An attacker may take advantage of such computers that are probing for an access point by creating a rogue access point and allowing the authorized computer to connect to it. Once the connection is in place, the attacker may use the authorized computer as a bridge to gain access to the local wired network.

To protect against threats from devices probing for access points, administrators should make certain that their wireless access points are not publicly broadcasting their SSIDs. This will ensure that wireless access points cannot be detected by traditional war-driving tools.[50] Furthermore, users should be taught safe wireless practices, including disabling wireless network software unless it is in use and allowing it to connect to only known and trusted access points.

The second most common wireless threat during the first six months of 2006 was the use of a spoofed MAC address, which accounted for 17% of all wireless threats observed during this period. Access control lists are commonly used to ensure that only legitimate computers are allowed to access wireless networks. These access control lists are often implemented by validating computers based on the MAC address of their network interface card.[51] Since MAC addresses were not developed for purposes of security, an attacker can bypass these access control lists by changing their computer's MAC address to correspond with one that is authorized to access the target network, a practice known as spoofing.

An attacker using a computer with a spoofed MAC address can masquerade as an authorized computer, allowing the attacker to gain access to an organization's internal network. This may allow him or her to gain access to potentially sensitive information through open network shares or by eavesdropping on network communications.

[48] War driving is the commonly used term inspired by war dialing; however, similar attacks have been called war walking, war cycling, war flying, and war busing depending on the attacker's mode of transportation. Furthermore, war chalking includes marking (using chalk) urban areas that can be used to access various wireless network resources.
[49] A rogue access point is a wireless network access point placed by an attacker to attempt to intercept communications between authorized host computers on a local network. If an authorized computer associates with a rogue access point, the owner of the rogue access point will be able to eavesdrop on communications before forwarding them to a legitimate access point.
[50] Wireless access points that do not broadcast their SSID will go unnoticed by tools that strictly attempt to connect to an access point with any SSID. There are some tools, however, that allow attackers to passively scan for access points by eavesdropping on wireless network traffic.
[51] Media Access Control (MAC) addresses are used to uniquely identify a network interface card on a local area network.

To protect against malicious computers with spoofed MAC addresses, Symantec advises administrators to deploy tools that allow for the identification and tracking of such computers. Furthermore, by ensuring that all communications carrying sensitive information on a local wireless network are encrypted, it may be possible to reduce the impact of a successful attack.

The third most common threat detected targeting wireless networks during the first six months of 2006 was unauthorized NetStumbler clients,[52] which accounted for 16% of all detected wireless threats. NetStumbler is a freely available wireless network utility that allows users to identify wireless network access points and attain general information about them. NetStumbler may be used by a network administrator to monitor local network settings and detect rogue access points; however, detection of unauthorized NetStumbler clients typically indicates that a network is being targeted by a war driver.

Although NetStumbler attempts to identify wireless networks by probing for an access point (that is, requesting connection to an access point using any SSID, as discussed above), it is possible to specifically identify NetStumbler versions prior to version 0.4.0. This is because the tool transmits unique data when requesting additional information from a discovered access point.

It is important to be able to distinguish between unauthorized NetStumbler clients and devices that are simply probing for access points. The unauthorized use of a NetStumbler likely indicates malicious activity (such as war driving), rather than a poorly configured authorized computer. Administrators can protect their networks against war drivers who are using NetStumbler by ensuring that their wireless access points are not publicly broadcasting their SSIDs.

Denial of service attacks

Denial of service (DoS) attacks are a major threat to organizations. A successful DoS attack can render Web sites or other network services inaccessible to customers and employees. This could result in the disruption of organizational communications, a significant loss of revenue, and/or damage to the organization's reputation. Furthermore, as Symantec discussed in a previous *Internet Security Threat Report* (September 2005), criminal extortion schemes based on DoS attacks are becoming more common.[53]

For this version of the Internet Security Threat Report, Symantec has changed the methodology used to obtain and record attack data. As a consequence of this methodological change, any comparison with attack data gathered in previous periods would be invalid; therefore, this discussion will focus only on the period between January 1 and June 30, 2006.

Although there are numerous methods for carrying out DoS attacks, Symantec derives the data for this metric by measuring attacks carried out by flooding a target with SYN requests.[54] This type of attack works by overwhelming a target with SYN requests and not completing the initial request, which prevents other valid requests from being processed. In many cases, SYN requests with forged IP addresses are used to carry out an attack, allowing a single attacking computer to initiate multiple connections. This results in unsolicited traffic, known as backscatter, being sent to other computers on the Internet (whose IP addresses were spoofed). This backscatter is used to derive the number of DoS targets observed throughout the reporting period. Backscatter is only one method of obtaining DoS statistics and for the purposes of this report is only intended to provide a high-level overview of overall DoS activity.

[52] http://www.stumbler.net
[53] Symantec *Internet Security Threat Report*, Volume VIII (September 2005): https://enterprise.symantec.com/enterprise/whitepaper.cfm?id=2238, p. 11 and 30
[54] The TCP protocol requires a three-way exchange to be carried out before any data is sent. The SYN request is the first phase of the three-way exchange. Once a SYN request is received by a server, a SYN-ACK is sent in response. The final step is an ACK response, completing the connection negotiation process.

During the first six months of 2006, Symantec observed an average of 6,110 DoS attacks per day (figure 10). DoS attacks are generally carried out by a wide variety of attackers, from amateurs who simply download a freely available tool, to owners of highly organized bot networks whose primary purpose is to carry out coordinated attacks.[55]

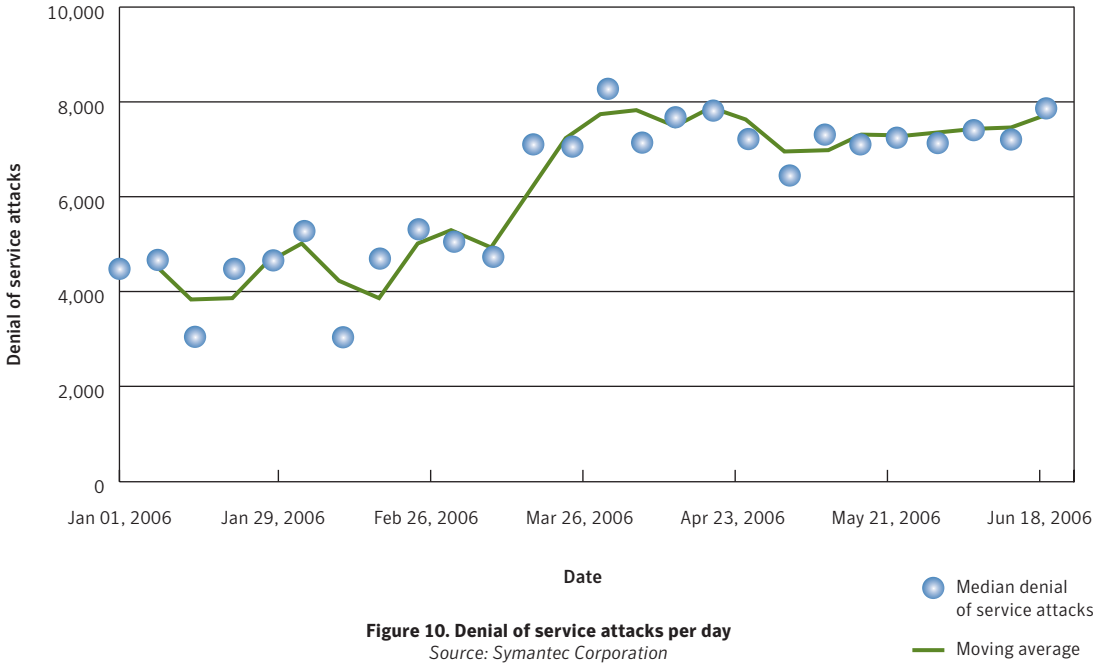

Figure 10. Denial of service attacks per day
Source: Symantec Corporation

Defending against DoS attacks that use forged source addresses is difficult, as spoofed addresses make filtering based on the IP address very complicated. Some operating systems have configuration options that may be used to make the computers less prone to resource exhaustion, thereby making them more resilient against DoS attacks. Administrators should optimize this to minimize the effects of DoS attacks.

Organizations should ensure that a documented procedure exists for responding to DoS events. One of the best ways to mitigate a DoS attack is to filter upstream of the target. For most organizations, this filtering will involve working in conjunction with their Internet service provider (ISP). Symantec also recommends that organizations perform egress filtering on all outbound traffic.

[55] An example of a denial of service attack tool is the Smurf tool, which is designed to carry out ICMP flood attacks against target computers.

Top countries targeted by denial of service attacks

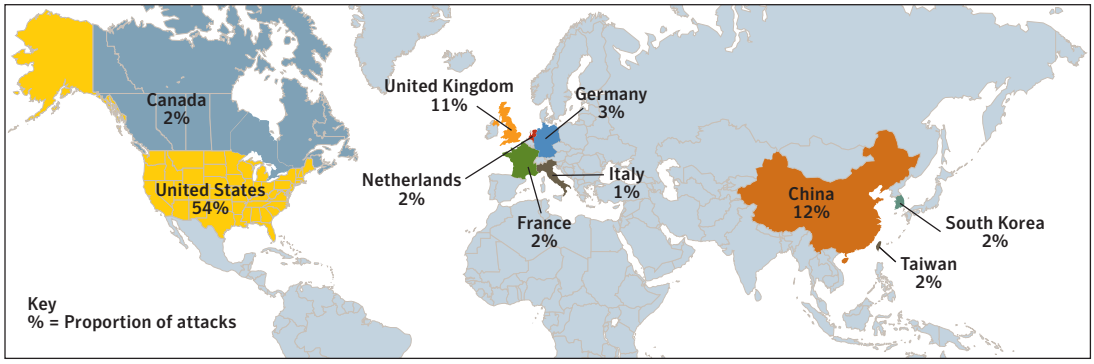

Figure 11. Top countries targeted by denial of service attacks
Source: Symantec Corporation

For the first time, in this volume of the Symantec *Internet Security Threat Report*, Symantec is tracking the geographic location of targets of denial of service (DoS) attacks. Insight into the locations targeted by these attacks is valuable in determining global trends in DoS attack patterns. It may also help administrators and organizations in affected countries to take the necessary steps to protect against or minimize the effects of DoS attacks.

Between January 1 and June 30, 2006 the United States was the location of the most DoS targets, accounting for 54% of the worldwide total (figure 11). The prominence of the United States as a target is not surprising. The country's extensive broadband Internet infrastructure and its high proportion of Internet-connected organizations make it a very attractive target.

China was targeted by the second highest number of DoS attacks, accounting for 12% of the total. The United Kingdom was the third most common target, accounting for 11% of all detected attacks. Like the United States, both China and the United Kingdom have an extensive broadband Internet infrastructure. Both countries are also regional and global political and economic centers. As a result, attackers who are acting on financial or political motives may choose to target these countries.

Organizations should ensure that a documented procedure exists for responding to DoS events. One of the best ways to mitigate a DoS attack is to filter upstream of the target. For most organizations, this filtering will involve working in conjunction with their Internet service provider (ISP). Symantec also recommends that organizations perform egress filtering on all outbound traffic. DoS victims frequently need to engage their upstream ISP to help filter the traffic to mitigate the effects of attacks.

Top sectors targeted by denial of service attacks

Rank	Sector	Proportion of attacks
1	Internet Service Provider	38%
2	Government	32%
3	Telecommunications	8%
4	Transportation	4%
5	Education	3%
6	Accounting	3%
7	Utilities / Energy	3%
8	Insurance	3%
9	Financial Services	2%
10	Information Technology	2%

Table 5. Top sectors targeted by denial of service attacks
Source: Symantec Corporation

This metric will assess the sectors that are most commonly targeted by DoS attacks. This data could help administrators and organizations in affected sectors take the necessary steps to prepare defenses against DoS attacks.

The sector most frequently targeted by DoS attacks in the first half of 2006 was the Internet service provider (ISP) sector, which was targeted by 38% of all DoS attacks (table 5). ISPs are popular targets for several reasons. First, they are responsible for providing Internet service to a high number of users. By successfully attacking an ISP, an attacker can effectively create denial of service conditions for a high number of users at one time. Second, ISPs also host Web sites and provide Internet access to many potential target organizations. Attackers wanting to target an organization's Web sites or networks could do so by launching a DoS attack against the organization's ISP.

The second most popular target of DoS attacks during the first half of 2006 was the government sector, which was targeted by 32% of all detected attacks. Government Web sites typically provide essential services and information. They are also high-profile sites, so it is logical that the government sector is a popular target for DoS attacks.

The telecommunications sector was the third most popular target of DoS attacks,[56] accounting for eight percent of all detected attacks during the period. The telecommunications sector is likely an attractive target to attackers attempting to deny access to telecommunications services, such as voice over IP (VoIP), and the companies that host them. Also, many telecommunications companies also provide Internet services, similar to the ISP sector, and so will be popular targets for the same reasons as organizations in that sector.

[56] The telecommunications sector is made up of organizations that provide various telecommunications services. While this could include the provision of Internet services, it is not usually their primary function.

Organizations should ensure that a documented procedure exists for responding to DoS events. One of the best ways to mitigate a DoS attack is to filter upstream of the target. For most organizations, this filtering will involve working in conjunction with their Internet service provider (ISP). Symantec also recommends that organizations perform egress filtering on all outbound traffic.

Bot networks

Bot networks are groups of compromised computers on which attackers have installed software that listens for and responds to commands, typically using Internet relay chat (IRC), thereby giving the attacker remote control over the computers. The software used to compromise and control these computers, known as bot software, is often modular malicious code. As such, after it has infected a computer, it may be upgraded to include new functionality, including exploit code that can target new vulnerabilities. (For more information on modular malicious code, please see the "Malicious Code Trends" section of this report).

Bots can have numerous effects on Internet users, including home users, small businesses, and large organizations. A single infected host within a network (such as a laptop that was compromised outside the local network and then connected to the network, either directly or by VPN) can allow a bot to propagate to other computers that are normally protected against external attacks by corporate firewalls.

Bots can be used by external attackers to perform DoS attacks against an organization's Web site, which can render Web sites or other network services inaccessible to customers and employees. This could result in the disruption of organizational communications, a significant loss of revenue, and/or damage to the organization's reputation. Furthermore, bots within an organization's network can be used to attack other organizations' Web sites, which can have serious business and legal consequences. Finally, bots can be used by attackers to harvest confidential information from compromised computers.

This metric explores the number of active bot network computers that the Symantec™ Global Intelligence Network has detected and identified during the first 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.[57]

As a consequence of this, Symantec does not identify all bot network 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 activity, and can thus provide an understanding of current bot activity across the Internet as a whole.

As has been mentioned previously, for this version of the *Internet Security Threat Report*, Symantec has implemented new methodologies to obtain and record attack data. This extends to the identification of bot network activity. These methodological changes have expanded Symantec's insight into attack and bot network activity to give a more accurate view of global trends. As a consequence of these changes, any comparison with attack data gathered in previous periods would be invalid. As a result, this discussion will focus only on the current period between January 1 and June 30, 2006.

[57] It should be noted that Symantec has identified a number of bots that propagate through means other than exploiting network-based vulnerabilities. In spite of this, Symantec believes that the population of network propagating bots is a good indicator of overall bot activity trends.

During this period, Symantec observed an average of 57,717 active bot network computers per day (figure 12). Symantec also observed 4,696,903 distinct bot network computers that were identified as being active at any one (or more) point in time during the six-month period.

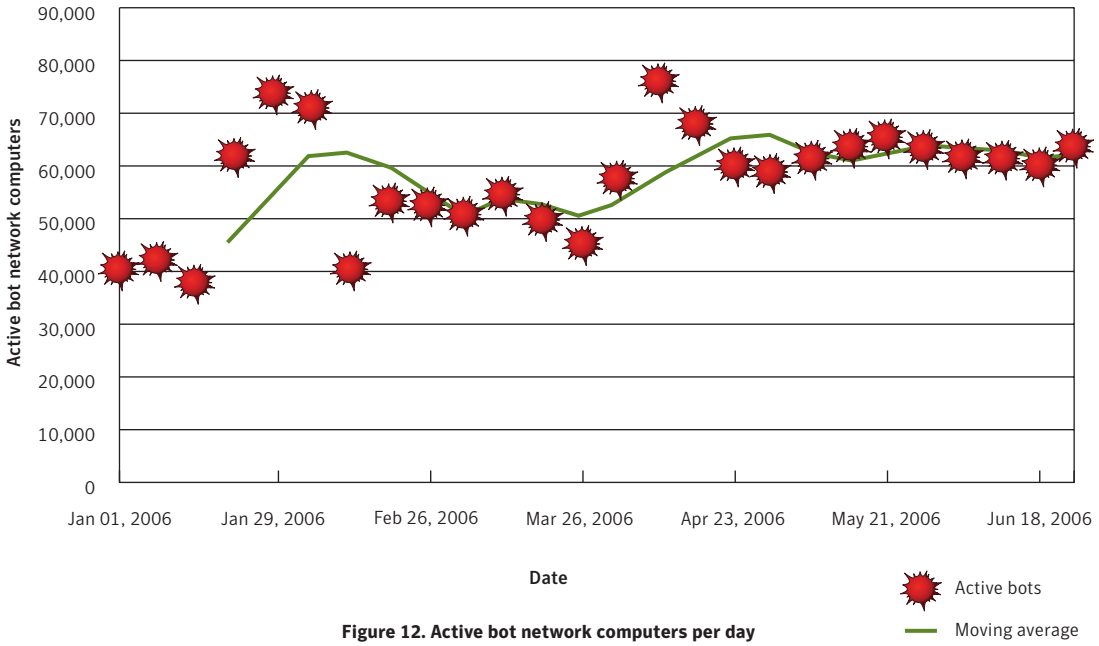

Figure 12. Active bot network computers per day
Source: Symantec Corporation

Symantec also tracks the number of bot command-and-control servers worldwide. Bot command-and-control servers are computers that bot network owners use to relay commands to bot-infected computers on their bot networks. Symantec identified 6,337 bot command-and-control servers during the first six months of 2006.

Throughout this reporting period the number of active bot network computers each day remained relatively constant. As discussed in the previous volume of the *Internet Security Threat Report*, this behavior likely indicates that the population of network-propagating bot-infected computers has reached the saturation point.[58] In the same discussion, Symantec speculated that bot network populations rise and fall in a boom-and-bust cycle in which numbers of bot-infected computers will increase for a period of time, level off, then decrease for another period before beginning the cycle again. The current leveling off is consistent with that cycle.

Symantec believes that bot network owners are increasingly discreet about the number of machines they bring online at any one time. This is due to the increased awareness among end users and organizations of bots and bot networks. Large numbers of bot network machines acting in a coordinated fashion are often easily identifiable, making it easier for ISPs to detect and shutdown bot networks.

[58] Symantec *Internet Security Threat Report*, Volume IX (March 2006): http://enterprisesecurity.symantec.com/content.cfm?articleid=1539, pp. 36

If bots begin to exploit an attack vector that bypasses firewalls and perimeter defenses, the population of bot-infected computers could increase rapidly. This possible boom period could have a greater impact on the Internet than earlier ones because bot network owners have become more organized and experienced. Furthermore, bot technology is much more entrenched due to the public disclosure of bot source code. Finally, some bots and bot networks are reportedly using encrypted channels to communicate, which could make them much more difficult to detect.[59]

To prevent bot infections, Symantec recommends that ISPs perform both ingress and egress filtering to block known bot network traffic. 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 the enterprises notify their ISPs of any potentially malicious activity.

To reduce exposure to bot-related attacks, end users should employ defense in-depth strategies, including the deployment of antivirus software and a firewall. Creating and enforcing policies that identify and limit applications that can access the network may also be helpful in limiting the spread of bot networks. 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. In order to do this, Symantec calculates the number of computers worldwide that are known to be infected with bots and assesses what percentage are situated in each country.

This metric can help analysts understand how bot-infected computers are distributed globally. The identification of bot-infected computers is important, as a high percentage 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.

Furthermore, Symantec tracks the global distribution of bot command-and-control servers. Bot command-and-control servers are computers that bot network owners use to relay commands and instructions to other computers on their bot 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.

For this version of the *Internet Security Threat Report*, Symantec has implemented new methodologies to obtain and record attack data as well as to identify bot network activity. These changes have expanded Symantec's view into attack and bot network activity to give a more accurate view of global trends. As a consequence of these methodological changes, any direct comparison between bot network data from

[59] For instance, W32.Nugache.A@mm was used to encrypt peer-to-peer communication to talk to other bots in a bot network. For more details, please see http://www.symantec.com/security_response/writeup.jsp?docid=2006-043016-0900-99

the current period and the previous periods would be invalid. Despite this, the relative ranks of bot network countries will be included in this discussion, as Symantec believes that they can still offer some insight into global bot network trends.

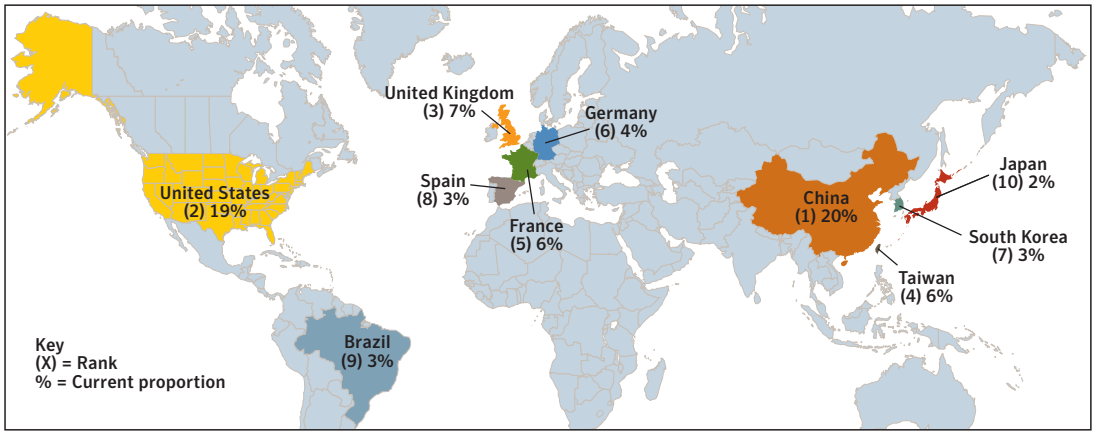

Figure 13. Top countries by bot-infected computers
Source: Symantec Corporation

China had the highest number of bot-infected computers during the first half of 2006, accounting for 20% of the worldwide total (figure 13). This ranking represents a rise from third place in the second half of 2005. 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.

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. Frequently, ISPs will focus their resources on meeting growing broadband demand at the expense of implementing security measures, such as port blocking and ingress and egress filtering. As a result, ISPs that are expanding their services rapidly may have security infrastructures that are underdeveloped relative to their needs. Symantec believes that bot activity in China will continue to rise as broadband Internet continues to be adopted at a rapid rate.

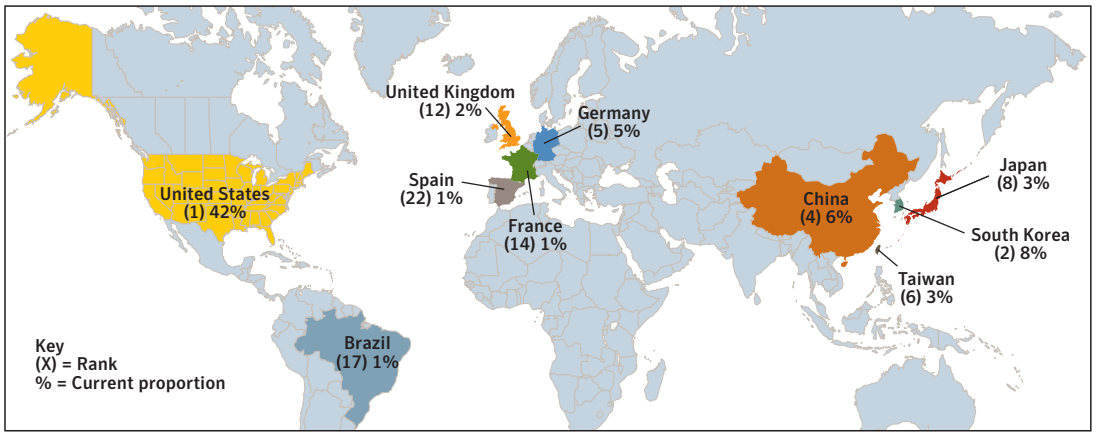

Figure 14. Distribution of command-and-control servers in top ten bot-infected countries
Source: Symantec Corporation

Although China had the most bot-infected computers worldwide, it had only the fourth highest number of known command-and-control servers worldwide (figure 14). This discrepancy likely indicates that the majority of bot-infected computers within China are being controlled from servers in other countries. However, this does not mean the person controlling the server is necessarily situated elsewhere. 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. Attackers frequently use previously compromised computers to host command-and-control servers, allowing them to obscure their actual location. For example, an attacker in China could control a command-and-control server in South Korea to administer bot-infected computers all over the world.

In the first six months of 2006, the United States had the second highest number of bot-infected computers. Nineteen percent of bot-infected computers worldwide were situated there. The United States was also the site of 42% 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. The high proportion of bot-infected computers and bot command-and-control servers in the United States is driven by its extensive Internet and technology infrastructure and the fact that more than 49 million broadband Internet users are located there.[60]

The United Kingdom had the third highest number of bot-infected computers worldwide, accounting for seven percent of the worldwide total. It ranked second in the world in the second half of 2005. This drop is likely an indication that the security infrastructure in the United Kingdom is beginning to catch up to the growth of Internet connectivity. The United Kingdom accounted for only two percent of all known command-and-control servers worldwide. This indicates that the majority of bot network computers inside the country are likely controlled by servers in other countries.

[60] http://www.oecd.org/document/39/0,2340,en_2649_34449_36459431_1_1_1_1,00.html

Bot-infected computers by city

Rank	City	Country	Proportion
1	Beijing	China	2.91%
2	Guangzhou	China	1.67%
3	Seoul	South Korea	1.61%
4	Madrid	Spain	1.50%
5	Hangzhou	China	1.41%
6	Los Angeles	United States	1.20%
7	Taipei	China	1.10%
8	Shanghai	China	1.02%
9	Paris	France	0.99%
10	Chicago	United States	0.99%

Table 6. Top cities by bot-infected computers
Source: Symantec Corporation

In addition to identifying top bot-infected countries, Symantec also tracks the distribution of bot-infected computers by city. As with the previous metric, the identification of bot-infected computers is important, as a high percentage of infected machines 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 city.

During the first half of 2006, Beijing was the city with the most bot-infected computers in the world, accounting for almost three percent of the worldwide total (table 6). Guangzhou, China ranked second, with just under two percent of the world's bot-infected computers. Seoul, South Korea had the third highest number of bot-infected computers worldwide, accounting for slightly less than two percent of the total. All of the top three cities in this category are large population centers that are cultural and economic centers in their respective countries. Furthermore, all have a large broadband Internet infrastructure.

In previous volumes of the *Internet Security Threat Report*, Symantec speculated that the number of computers with high-speed Internet in a region is a significant factor in determining the number of bot-infected computers. Considerations such as the type of industries situated in a city may also have a strong influence on the percentage of bot-infected computers.

Another factor may be the rate of growth in broadband connectivity. Symantec believes that new broadband customers may not be aware of the additional security precautions that need to be taken when exposing a computer to an always-on high-speed Internet connection. Furthermore, the addition of many new customers, with the corresponding increase in infrastructure and support costs, may inhibit the ability of ISPs to respond to reports of network abuse and infection.

Symantec recommends that organizations employ defense in-depth strategies, including firewalls and adequate perimeter filtering. Furthermore, administrators are advised to subscribe to a vulnerability alerting service, and to apply necessary patches across the enterprise in a timely manner. End users should always deploy antivirus software and a firewall and should ensure that antivirus definitions are updated regularly.

Top originating countries

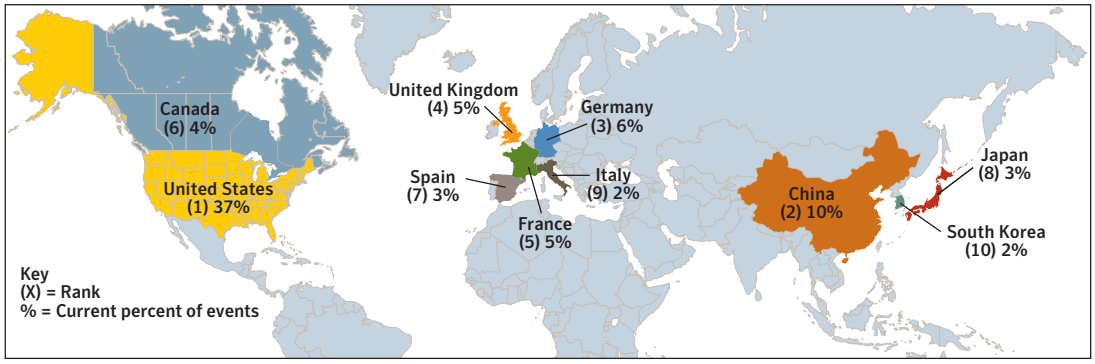

Figure 15. Top originating countries
Source: Symantec Corporation

This section will discuss the top countries of attack origin. This metric only discusses the location of the computer from which the attack originates and not the actual location of the attacker. While it is simple to trace an attack back to the computer from which it was launched, that computer may not be the attacker's own system. Attackers frequently hop through numerous systems or use previously compromised systems to obscure their location prior to launching the actual attack. For example, an attacker in China could launch an attack from a compromised system located in South Korea against a Web server in New York. Further complicating the matter is that international jurisdictional issues often prevent proper investigation of an attacker's real location.

During the first six months of 2006, the United States ranked as the top country of attack origin, accounting for 37% of the worldwide total (figure 15). Attack activity originating in the United States increased by 29% in this period, which is 13 percentage points above the average increase of 16%. This is likely driven by recent growth in broadband infrastructure there. As has been stated in previous volumes of the *Internet Security Threat Report*, an increase in broadband connectivity in a country often leads to an increase in attacks and bot infections originating there. During the last half of 2005, the number of broadband Internet users in the United States increased by nearly seven million,[61] which is the largest increase in volume in the country's history.

[61] http://www.oecd.org/document/39/0,2340,en_2649_34449_36459431_1_1_1_1,00.html

China remained in second position for the first half of 2006, accounting for ten percent of all attacking IP addresses. Attack activity originating in China increased by 37% over the previous reporting period, well above the 16% average increase. Symantec expects attack activity originating in China to continue to rise as broadband Internet continues to be adopted there.

Germany rose from fourth position during the last half of 2005 to third position in this reporting period. Six percent of all attacking IP addresses were situated there. The United Kingdom dropped from third position to fourth with five percent of observed attacking IP addresses. The United Kingdom's drop is likely an indicator that computer security infrastructure there is catching up with its broadband growth, which may have resulted in sufficient security measures being put in place.

Top targeted sectors

Current Rank	Previous Rank	Sector	Current Proportion of Targeted attacks	Previous Proportion of Targeted attacks
1	1	Home user	86%	93%
2	2	Financial Services	14%	4%
3	6	Government	<1%	<1%
4	3	Education	<1%	2%
5	8	Information Technology	<1%	<1%
6	7	Health care	<1%	<1%
7	5	Accounting	<1%	<1%
8	10	Telecommunications	<1%	<1%
9	4	Small Business	<1%	<1%
10	14	Utilities / Energy	<1%	<1%

Table 7. Top targeted sectors by proportion of targeted attacks
Source: Symantec Corporation

Although many attackers select targets randomly, some attack computers within a specific sector, industry, or organization. For the purposes of this metric, these attacks are referred to as "targeted attacks." For this discussion, a targeted attack is identified as an IP address that has attacked at least three sensors in a given industry to the exclusion of all other industries during the reporting period. This metric has been redesigned from previous volumes of the *Internet Security Threat Report* to include home users as well those sectors that were previously assessed.

Between January 1 and June 30, 2006, the home user sector was the most highly targeted sector, accounting for 86% of all targeted attacks (table 7). As computers in the home user sector are less likely to have well established security measures and practices in place, they may be more vulnerable to targeted attacks. Furthermore, as home users represent a fertile resource for identity theft, it is likely that many of the targeted attacks are used for fraud or other financially motivated crime.

It should be noted that the number of targeted attacks detected against home users might be inflated due to the way in which they access the Internet. It is likely that the majority of home users share networks that span a single block of IP addresses. As a result, opportunistic attacks targeting a broadband ISP may be noted as targeted attacks, thereby artificially inflating the percentage of targeted attacks against this sector. This assertion is supported by findings outlined in the "Top sectors targeted by denial of service attacks" section above that show ISPs as the primary target of DoS attacks.

Financial services was the second most frequently targeted sector in the first half of 2006. As was discussed in the previous *Internet Security Threat Report*, Symantec believes that attackers are increasingly motivated to conduct on-line criminal activities by financial gain.[62] The financial services industry is typically considered a popular target for attackers hoping to profit from attack activity. Symantec expects that attacks targeted against the financial services industry will continue to rise as attackers become more profit-driven.

Government was the third most frequently targeted sector during the first half of 2006, although it accounted for less than one percent of all targeted attacks. Attackers may target government organizations for many reasons, including politically motivated attacks and attempts to gain access to government records for the purposes of identity theft.

[62] Symantec *Internet Security Threat Report*, Volume VIII (September 2005) p. 4, and Volume IX (March 2006) p. 19. Both are available at: http://enterprisesecurity.symantec.com/content.cfm?articleid=1539

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.[63] Symantec also maintains one of the world's most comprehensive vulnerability databases, currently consisting of over 18,000 vulnerabilities (spanning more than a decade) affecting more than 30,000 technologies from over 4,000 vendors. The following discussion of vulnerability trends is based on a thorough analysis of that data. (Please note that all numbers presented in this discussion have been rounded off to the nearest whole number. As a result, some cumulative percentages may exceed 100%.)

This section of the Symantec *Internet Security Threat Report* will discuss vulnerabilities that have been disclosed between January 1 and June 30, 2006. It will compare them with those disclosed in the two previous six-month periods and discuss how current vulnerability trends may affect potential future Internet security activity. Where relevant, it will also offer protection and mitigation strategies. Symantec's recommendations for best security practices can be found in Appendix A at the end of this report.

This section of the Symantec *Internet Security Threat Report* will discuss:

- Total number of vulnerabilities disclosed
- Web application vulnerabilities
- Easily exploitable vulnerabilities
- Easily exploitable vulnerabilities by type
- Patch development time for operating systems
- Window of exposure for enterprise vendors (consisting of patch development time and exploit code development time)
- Window of exposure for Web browsers (consisting of patch development time and exploit code development time)
- Web browser vulnerabilities
- Exploit code release period

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

It should be noted that, unlike other reports in the *Internet Security Threat Report*, the "Vulnerability Trends Report" is based on data that often changes over time. This is because entries in the vulnerability database are frequently revised as new information emerges. For instance, vulnerabilities may be attributed to a particular reporting period after that period has ended because additional information has become available after that time. Conversely, entries may be removed after a reporting period has ended because they are subsequently deemed not to have been vulnerabilities. Because of this, statistics and percentages reported in one volume of the *Internet Security Threat Report* may not agree with information presented in subsequent volumes. As a result, some of the comparative data for previous reporting periods that is presented within this report may differ from the data presented in previous volumes of the *Internet Security Threat Report*.

Total number of vulnerabilities disclosed

Symantec documented 2,249 new vulnerabilities in the first half of 2006 (figure 16). This is an increase of 18% over the 1,912 vulnerabilities that were documented in the second half of 2005. It is also a 20% increase over the 1,874 vulnerabilities that were reported in the first half of 2005. Symantec documented a higher volume of vulnerabilities in this reporting period than in any other previous six-month period.[64]

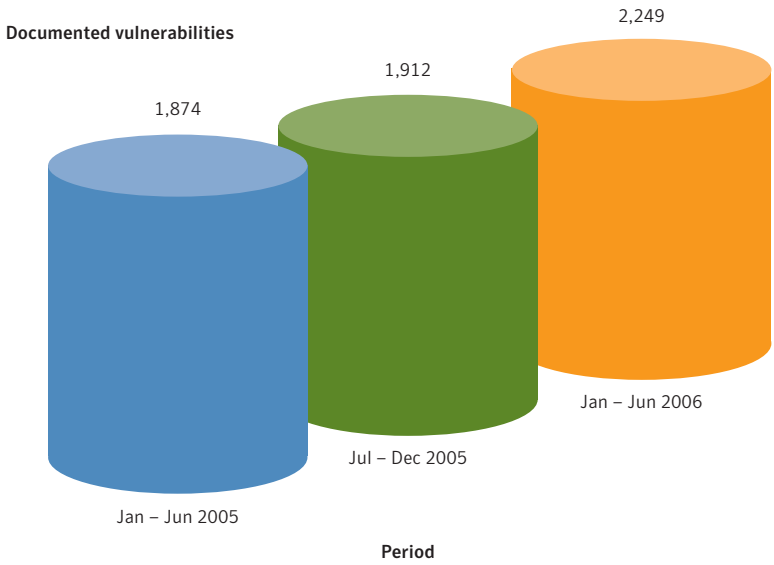

Figure 16. Total volume of vulnerabilities
Source: Symantec Corporation

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

The marked increase in the number of vulnerabilities can be attributed to the continued growth in those that affect Web applications. This is due to the relative ease of discovering vulnerabilities in Web applications compared to other applications. Additionally, Web applications generally have quicker release cycles than traditional desktop and server applications. This provides security researchers with a continually growing source of new applications to audit, particularly as, in many cases, Web applications do not undergo the same degree of quality assurance and testing as other applications. This will be discussed in greater detail in the "Web application vulnerabilities" section below.

Another factor in the general growth of vulnerability volume is that security researchers have better tools at their disposal than in previous periods. In a previous *Internet Security Threat Report*, Symantec speculated that advanced research tools would make the discovery of vulnerabilities easier than ever before.[65] That appears to be the case. For example, recent advances in fuzzing tools and techniques have made it easier for security researchers to automate vulnerability discovery.[66] As well, there are numerous disassembly and debugger tools that are specifically customized for security research.[67] Furthermore, virtualization appears to be becoming more accessible to security researchers due to the availability of new virtualization software.[68]

Symantec recommends that administrators employ a good asset management system, patch management system or service, and a vulnerability alerting service, all of which can help to quickly assess whether a new vulnerability is a viable threat or not. They should also monitor vulnerability mailing lists and security Web sites for new developments in vulnerability research.

Web application vulnerabilities

Web applications are technologies that use a browser for their user interface, rely on HTTP as the transport protocol, and reside on Web servers. Examples of Web-based applications include content management systems, e-commerce suites (such as "shopping cart" implementations), Weblogs, and Web-based email. An increasing number of traditional software vendors are re-implementing their existing applications with Web-based user interfaces.

Vulnerabilities in these technologies are particularly threatening because they are typically exposed to the Internet through a Web server and because they are often required to be publicly available. Web-based attacks may be challenging to detect and prevent because they are often easy to obfuscate. While many IDS vendors provide generic signatures for these attacks, there may not be signatures that are application-specific or that account for all variants of an attack. In the worst case scenario, exploitation of Web application vulnerabilities could enable a successful attacker to compromise an entire network by gaining access through a single vulnerable system. Vulnerabilities in these technologies can also give an attacker access to confidential information from databases without having to compromise any servers.

[65] Symantec *Internet Security Threat Report*, Volume VIII (September 2005): http://enterprisesecurity.symantec.com/content.cfm?articleid=1539, p. 88
[66] Fuzzing is a testing, quality assurance, and security research technique that typically involves randomly generating data to use as input to an application.
[67] Security researchers and reverse engineers have developed a number of plug-ins for Data Rescue IDA Pro disassembler and Ollydbg debugger, such as those found here: http://www.openrce.org/downloads/
[68] http://www.securityfocus.com/columnists/397

Vulnerabilities affecting Web applications accounted for 69% of all vulnerabilities that were documented by Symantec in the first half of 2006 (figure 17). This is a slight increase over the 68% seen in the second half of 2005. It is also higher than the 60% proportion in the first half of 2005.

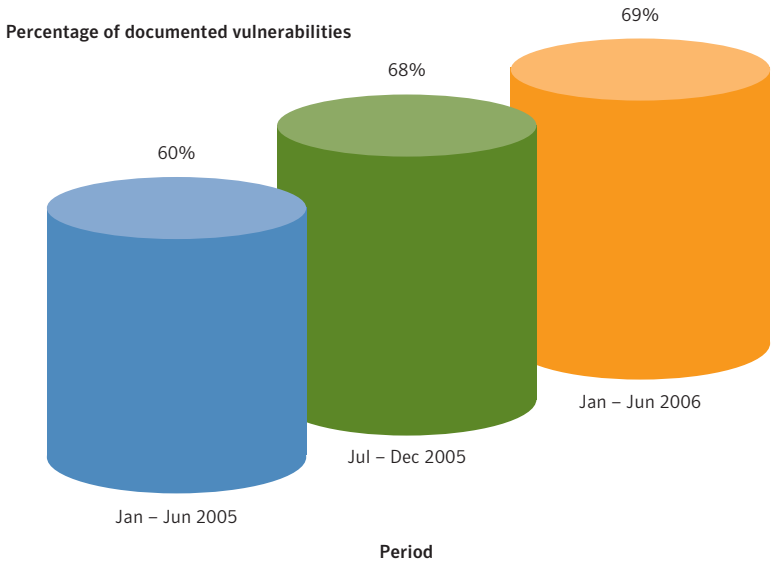

Figure 17. Web application vulnerabilities
Source: Symantec Corporation

As was discussed in the "Total vulnerabilities disclosed" section, Web applications generally have quicker release cycles than traditional desktop and server applications. This provides security researchers with a continually growing source of new applications to audit, particularly as, in many cases, Web applications do not undergo the same degree of quality assurance and testing as other applications.

Web applications are required to accept and interpret input from many different sources, and there are often very few restrictions to distinguish valid input from invalid input. Web applications can host malicious content that may affect clients but be otherwise innocuous to the server. This trait increases the susceptibility of Web applications to attack, as the application must also be aware of malicious input that is hostile to its clients and not just itself.

This is further complicated because Web browsers, the application through which most Web applications operate, are very liberal in what they will accept and interpret as valid input. Because of different browser implementations, some malicious content may be harmful to one browser but not to another. This creates further confusion for a Web application that is trying to determine which input is invalid and which is potentially malicious.

Many security researchers opt for a "quantity" over "quality" approach, meaning that the vulnerabilities that can be discovered most easily will take precedence over those that take longer to research. Researchers who favor this approach often choose Web applications because they present easy targets.

This is because the source code is often readily available to be audited (although in many cases security researchers can also quickly discover vulnerabilities on live Web sites). As a result, researchers can often find many more vulnerabilities in Web applications in a shorter period of time than in other applications. For instance, Web applications are often susceptible to common types of input validation vulnerabilities such as cross-site scripting and SQL injection that are typically easy to discover with a minimal amount of effort and skill.[69]

In order to protect against the exploitation of Web application vulnerabilities, Symantec recommends that administrators employ a good asset management system to track what assets are deployed and which may be affected by the discovery of new vulnerabilities. Vulnerability assessment technologies may 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.

Organizations should 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 Secure Development Lifecycle and threat modeling.[70] Symantec recommends the use of secure shared components that have been audited for common Web application vulnerabilities to limit the risk of introducing new vulnerabilities when implementing features from scratch. 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.

Easily exploitable vulnerabilities

Previous versions of the *Internet Security Threat Report* assessed vulnerabilities according to their ease of exploitation. However, over the past six months, the Symantec Vulnerability Database has adopted the Common Vulnerability Scoring System (CVSS).[71] Various criteria for the previous "ease of exploit" metric are incompatible with the CVSS standard and thus that method of categorizing and analyzing vulnerabilities is no longer supported by the Symantec Vulnerability Database. This version of the *Internet Security Threat Report* will instead discuss the volume of easily exploitable vulnerabilities.

Easily exploitable vulnerabilities present a serious threat to organizations because they can be exploited with a minimal amount of skill and effort. Easily exploitable vulnerabilities fall into one of two classes:

- Vulnerabilities that have exploit code associated with them or for which exploit code is known to be available. Previous versions of the *Internet Security Threat Report* referred to this class as "exploit code available."
- Vulnerabilities that do not require exploit code for successful exploitation. Previous versions of the *Internet Security Threat Report* referred to this class as "no exploit code required."

[69] Cross-site scripting is a vulnerability that allows attackers to inject hostile HTML and script code into the browser session of a Web application user. SQL injection is a vulnerability that can affect Web applications, allowing an attacker to inject their own SQL code into a database query that is made by the vulnerable application.
[70] 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.
[71] http://www.first.org/cvss/

In the first six months of 2006, 80% of newly disclosed vulnerabilities were considered easily exploitable (figure 18). This is a slight increase over the 79% of the easily exploitable vulnerabilities that were disclosed in the second half of 2005 and a larger increase over the 73% of vulnerabilities that were considered easily exploitable in the first half of 2005.

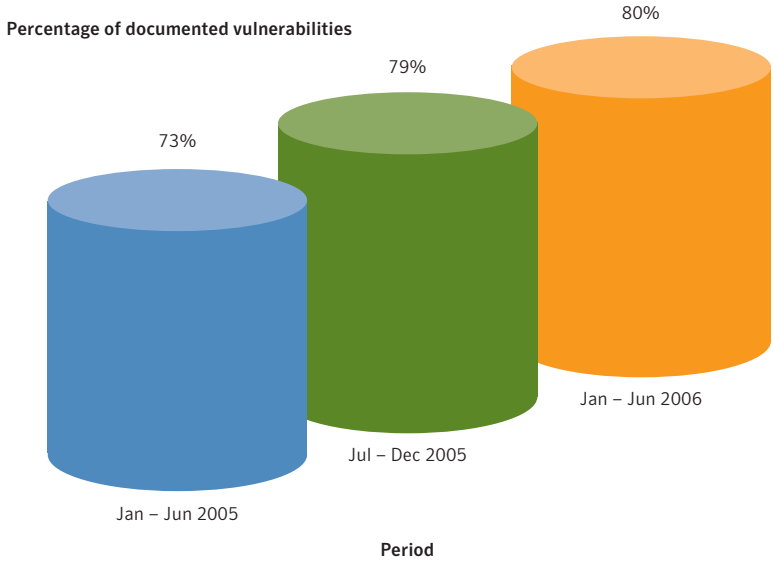

Figure 18. Easily exploitable vulnerabilities
Source: Symantec Corporation

The rise in the percentage of easily exploitable vulnerabilities is due mainly to the continued increase in Web application vulnerabilities. They make up the majority of easily exploitability vulnerabilities, although exploit code is also being actively developed for other threats as well. At any given time, four out of five vulnerabilities either have exploit code available or are easily exploitable without exploit code. This gives attackers a large pool of vulnerabilities to exploit. Given this high percentage, there is a reasonable probability that an organization will be affected by one or more of these vulnerabilities.

Easily exploitable vulnerabilities by type

To give enterprises an idea of the distribution of easily exploitable vulnerabilities, this version of the *Internet Security Threat Report* will provide a breakdown of easily exploitable vulnerabilities according to the type of affected application. The purpose of this metric is to assess which types of applications are currently being affected by easily exploitable vulnerabilities so that organizations that deploy these applications can take any steps necessary to protect their assets. For the purposes of this discussion, Symantec analyzes easily exploitable vulnerabilities according to the following six categories:[72]

- Browser vulnerabilities
- Client-side vulnerabilities
- Local vulnerabilities (those that do not require remote access but instead require only local access to exploit)
- Server vulnerabilities
- Web application vulnerabilities
- Other vulnerabilities (those that do not discretely fall into the previous categories)

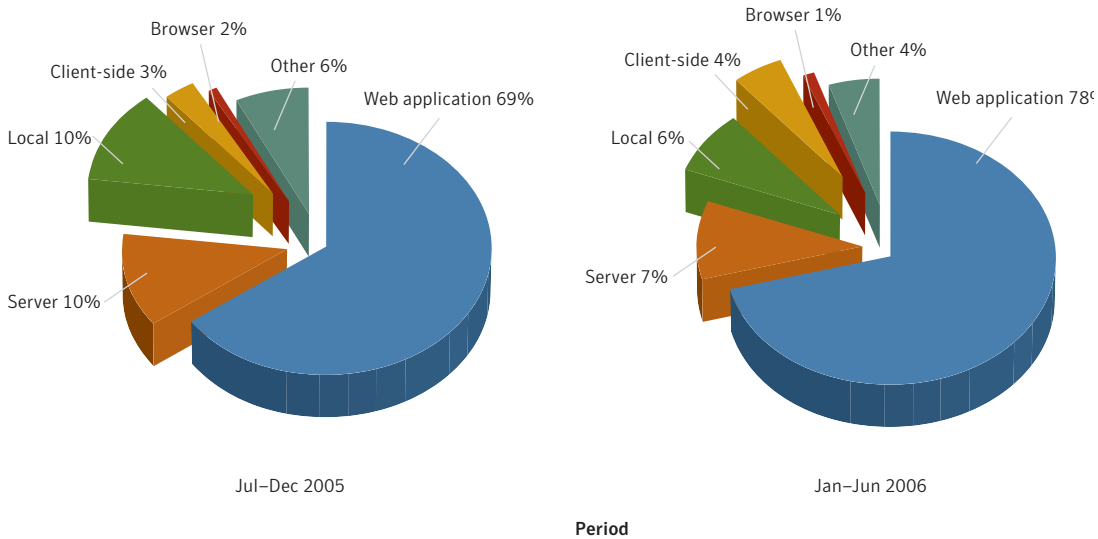

Figure 19. Easily exploitable vulnerabilities by type
Source: Symantec Corporation

Over the first six months of 2006, 78% of easily exploitable vulnerabilities affected Web applications (figure 19). This continued the increase that was evident in the two previous six-month periods, during which Web applications accounted for 69% and 61% of easily exploitable vulnerabilities respectively. In part, Web applications dominate this metric because they make up the majority of vulnerabilities that were documented over the last three periods. Furthermore, because many common Web application vulnerability types, such as cross-site scripting or SQL injection, do not require exploit code for successful exploitation, they are considered easily exploitable.

[72] These categories are explained in-depth in Appendix C of this report.

Server vulnerabilities made up seven percent of easily exploitable vulnerabilities in the first half of 2006. This is down from ten percent in the second half of 2005 and 14% in the first half of that year. The drop in the proportion of easily exploitable server vulnerabilities is a reflection of an overall drop in vulnerabilities affecting servers. However, servers, which have traditionally been the target of network worms, are still a higher risk than the other categories because attackers still see them as attractive targets.

Network perimeter defenses such as firewalls are effective measures against server attacks. Enterprises should restrict access to all ports and services that are not required to be publicly accessible. While many server attacks can be prevented with perimeter security measures, there are some public services that must accept traffic from arbitrary hosts on the Internet. Symantec recommends NIDS/NIPS to detect and protect against these attacks.

Over the past three reporting periods, local vulnerabilities accounted for the third highest percentage of easily exploitable vulnerabilities, with six percent in the first six months of 2006, ten percent in the second half of 2005, and 11% in the first half of 2005. In general, it is relatively easy to develop exploit code for local vulnerabilities because of the amount of control that the attacker has over the local environment, particularly as local attacks are generally executed by insiders who already have access to affected hosts. A local attacker can gather more information from the host operating system and has fewer of the variables to deal with that can often complicate the exploitation of remote vulnerabilities. Local attacks may also be used when remote attackers compromise a low-privileged service and need a means of gaining administrative access.

Host-based IDS/IPS systems can help to prevent local attacks. Features such as file integrity checking, behavioral intrusion prevention, and memory protection in the form of address space layout randomization (ASLR) can help to prevent or complicate attacks.[73] Organizations should also limit local access to critical hosts.

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 widely deployed 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.

During this period, Microsoft had a patch development time of 13 days, based on a sample set of 22 vulnerabilities (figure 20), a significant decrease from the 34 days in the last half of 2005, with 27 vulnerabilities patched. Red Hat, with 42 vulnerabilities to patch during this period, also had an average patch development time of 13 days for the first six months of 2006, a drop from the 28 days in the last half of 2005, when there were 98 Red Hat vulnerabilities.

[73] Address space layout randomization is a security feature that can prevent exploitation of buffer overflows and other memory corruption vulnerabilities by randomizing certain sections within the address space of a process.

For the current reporting period, Apple had the third shortest time to patch at 37 days for 21 vulnerabilities. This is a significant reduction from the 73-day average for 27 vulnerabilities in the second half of 2005. During this period, HP had an average patch development time of 53 days for the seven vulnerabilities it had to patch. This is down from 65-day average for the 15 patched vulnerabilities over the previous six months. Finally, in the first six months of 2006, Sun had an average patch development time of 89 days for sixteen patched vulnerabilities, down from 119 days in the second half of 2005 for 18 patched vulnerabilities.

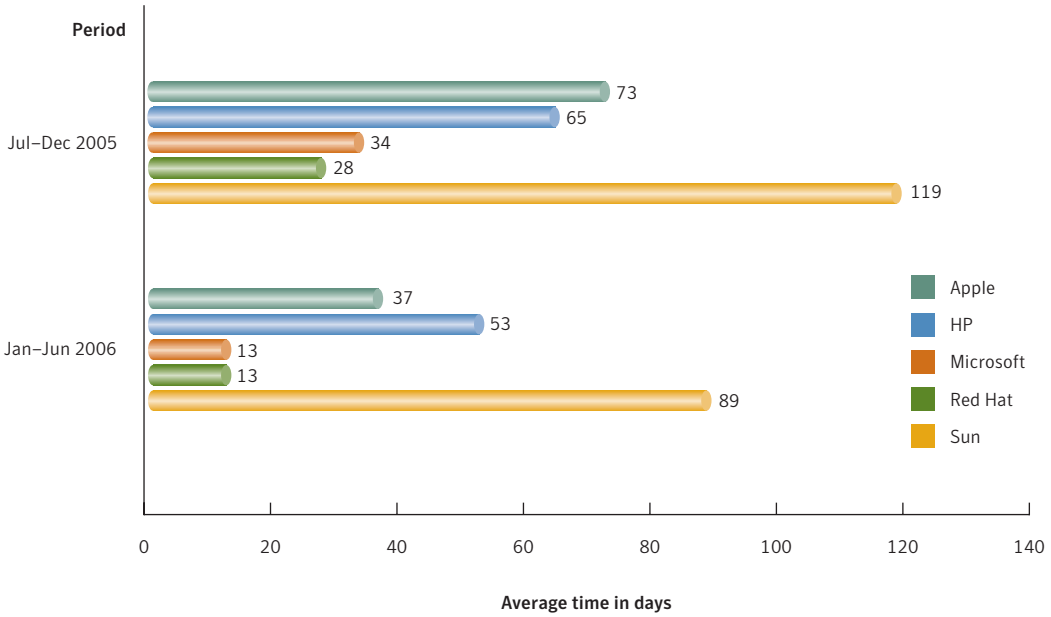

Figure 20. Operating system patch development time
Source: Symantec Corporation

Over the past six months, each of the five vendors had shorter average patch development times than in the previous two six-month periods. Linux vendor patch development times were generally shorter than those of the commercial UNIX vendors, HP and Sun. Over the past three reporting periods, Microsoft has had the shortest patch development time of all the operating system vendors.

Along with Microsoft, Red Hat had the lowest patch-development time during this reporting period. This is likely related to open-source collaboration. If a vendor or a member of the open-source community provides a patch, other vendors can share that patch and incorporate it into their distribution. Linux patches are not released on a fixed schedule; instead, they are often released on a daily basis. This approach differs from Microsoft and Apple, both of whom release their patches less frequently and in large batches to address as many vulnerabilities as possible at a time.

There are many reasons that the consumer-oriented vendors such as Microsoft and Apple have lower patch development times than some of the other vendors. Threats to desktop users and consumers generally carry a higher public profile and so there is likely more public pressure for vendors to be responsive and accountable.

The commercial UNIX vendors, HP and Sun, have the longest average patch development times. While both vendors release patches frequently, in many cases patches for optional third-party components are released later than patches for core operating system components. This likely drove up the average patch development time for these vendors during the first six months of 2006.

Window of exposure

Attackers use custom-developed code known as exploit code or exploits to take advantage of vulnerabilities to compromise a computer. The time lapse between the publication of an initial vulnerability report and the appearance of third-party exploit code is known as the "exploit code development time."[74] Exploit code development time is a concern to enterprises because it is a measurement of how long it takes for the average exploit to become public. If an exploit is published before a patch is available, administrators must implement other protective measures to reduce the risk of attack.

When a vulnerability is announced, the vendor in whose product it was found must develop and release a set of code known as a patch that will secure the vulnerability. 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." Until a patch is developed, released, and applied, computers on which the vulnerability resides may be susceptible to successful attack, particularly if exploit code developed for that vulnerability is available.

The difference between the exploit code development time and the patch development time is known as the "window of exposure." During this period of time, and until a patch is released, computers on which the vulnerable applications reside may be susceptible to successful compromise. This metric will assess the window of exposure in two contexts: applications developed by enterprise vendors and Web browsers.

The intent of this metric is to determine for how long after a vulnerability is announced a computer on which a vulnerable application resides is likely to be susceptible to a successful attack. The window of exposure is calculated by subtracting the exploit code development time from the patch availability time in each for the three contexts.

Window of exposure for enterprise vendors

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 that 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 (which form the majority), only exploits for technologies from enterprise vendors (that is, those that are generally widely deployed) are included.[75]

In the first six months of 2006, the average patch development time for enterprise vendors was 31 days (figure 21). During the same period, the average exploit code development time for vulnerabilities affecting enterprise vendors was three days. As a result, the window of exposure for this reporting period was 28 days.

[74] It should be noted that the data included in this discussion is limited to public examples of exploit code that Symantec has associated with specific vulnerabilities. There are many instances in which a private or commercial exploit may be available, but this data cannot be consistently tracked since exploit publication dates are not available.

[75] Vendors included in this metric are: Microsoft, Sun™, HP®, Symantec, EMC, IBM®, Cisco®, Oracle®, CA™ (Computer Associates), and McAfee®.

In the second half of 2005, the window of exposure for vulnerabilities affecting enterprise vendors was 50 days, based on a patch development time of 57 days and an exploit code development time of seven days. The window of exposure for the first half of 2005 was 60 days, based on a patch development time of 65 days and an exploit code development time of five days.

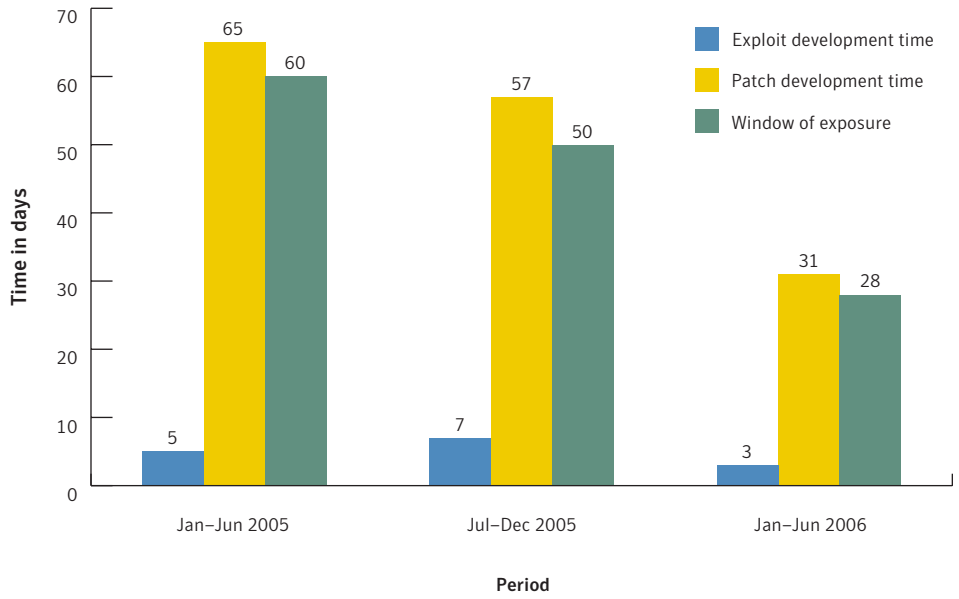

Figure 21. Window of exposure, enterprise vendors
Source: Symantec Corporation

The window of exposure for vulnerabilities in applications developed by enterprise vendors is thus narrowing. While there has been a slight reduction in the average exploit code development time, the main reason for this narrowing is that the average patch development time has dropped significantly.

Exploits for enterprise-vendor vulnerabilities are still being released quickly, forcing administrators to respond rapidly despite a lack of vendor-supplied remediation. However, the decreasing average patch development time indicates that enterprise vendors are responding more quickly to vulnerabilities. Despite this, it is critical that organizations follow up with the timely installation of patches, as attackers are still actively exploiting old vulnerabilities.

To minimize the possibility of successful exploitation, administrators need to understand newly disclosed vulnerabilities and be active in working around them. This may involve making changes to firewall configurations, creating or obtaining IDS/IPS signatures and rules, and locking down services. Symantec recommends that administrators employ a good asset management system or vulnerability alerting service. Both of these services can provide an understanding of the potential risk of new vulnerabilities, help to quickly assess whether they are a viable threat or not, and provide relevant protection/mitigation information. Administrators should monitor vulnerability mailing lists and security Web sites for new developments. They should also monitor mailing lists devoted to the discussion of security incidents or specific technologies, on which prevention and mitigation strategies may be discussed.

Window of exposure for Web browsers

This metric will assess the windows of exposure for four widely deployed Web browsers: Microsoft Internet Explorer, Mozilla (including Firefox and the Mozilla browser), Opera, and Apple Safari. The window of exposure will be calculated by computing the difference in days between the average patch development time and the average exploit code development time for vulnerabilities in these operating systems. Due to the number of browsers assessed in this metric, it will be necessary to chart exploit development time and patch development time separately.

Exploit code development time, Web browsers

The time lapse between the publication of an initial vulnerability report and the appearance of third-party exploit code is known as the "exploit code development time." In the first half of 2006, the average exploit code time for vulnerabilities in Apple Safari was zero days, the same average it had in the second half of 2005. In the first half of 2006, vulnerabilities in Internet Explorer had an average exploit code development time of one day, an increase over the second half of 2005, when it was zero days.

Between January 1 and June 30, 2006, the average exploit code development time for vulnerabilities in the Mozilla family of browsers was two days, down from seven days in the second half of 2005. In the first half of 2006, the average exploit code development time for Opera vulnerabilities was zero days, the same as in the second half of 2005.

The limited number of exploits that have been developed for vulnerabilities in Apple Safari and Opera make it difficult to perceive a long-term trend. In each case, the sample set includes only one or two public exploits, which were released within the first day of the initial disclosure of the affected vulnerability.

On the other hand, there are more exploits available for Internet Explorer and Mozilla vulnerabilities. As such, a more accurate understanding can be gained for Internet Explorer and Mozilla. The average exploit development time for Internet Explorer is still very short, as it is a high priority for attackers who are actively researching vulnerabilities and developing exploit code. This is likely because of the widespread deployment of the Microsoft browser.

Patch development time, Web browsers

The time period between the disclosure date of a vulnerability and the release date of an associated patch is known as the "time to patch." If exploit code is created and made public during this time, computers may be immediately vulnerable to widespread attack. It should be noted that this metric includes all patched vulnerabilities affecting the browser, regardless of their severity.

During the current reporting period, Apple Safari had an average patch development time average of five days, up from zero days in the second half of 2005 (figure 22).[76] In the first six months of 2006, Microsoft had an average patch development time of ten days for Internet Explorer vulnerabilities, down from the 25 days in the second half of 2005.

Between January and June 2006, Mozilla had an average patch development time of three days, slightly lower than the five-day average during the second half of 2005. During this reporting period, Opera had an average patch development time of two days. In the second half of 2005, it was 18 days.

[76] All patched vulnerabilities affecting Safari in the second half of 2005 were addressed by the vendor at the time of their announcement.

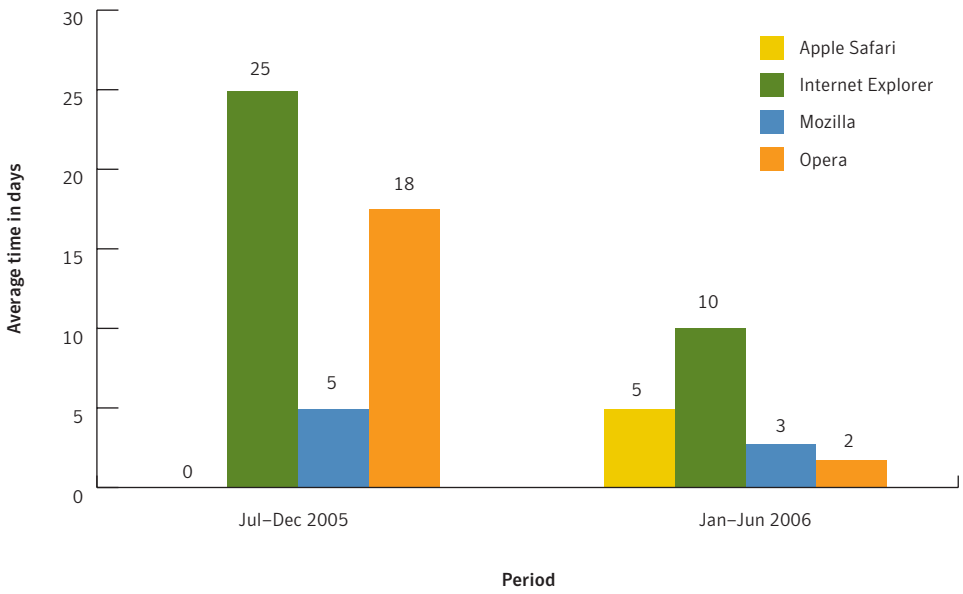

Figure 22. Patch development time, Web browsers
Source: Symantec Corporation

There does not appear to be any discernible trend in patch development times for Web browsers. This may be because these times are influenced by the number of vulnerabilities that are disclosed for each browser. Mozilla is the only vendor whose patch development time has decreased consistently over the past three six-month periods. Generally speaking, Internet Explorer has the longest patch development times of any browser. This may be due to the vendor's practice of issuing patches on a regular monthly schedule.

Window of exposure, Web browsers

The window of exposure is the difference between the average patch development time and the average exploit code development time for vulnerabilities in the selected Web browsers. In the first half of 2006, Internet Explorer had a window of exposure of nine days, down considerably from 25 days in the second half of 2005 (figure 23). During this reporting period, Apple Safari had a window of exposure of five days, up from zero days in the second half of 2005.

In the first half of 2006, Opera had a window of exposure of two days, down considerably from 18 days during the second half of 2005. In the first six months of 2006, Mozilla had a window of exposure of one day. In the second half of 2005, Mozilla had a window of exposure of negative two days, meaning that exploits were generally released after patches were available.

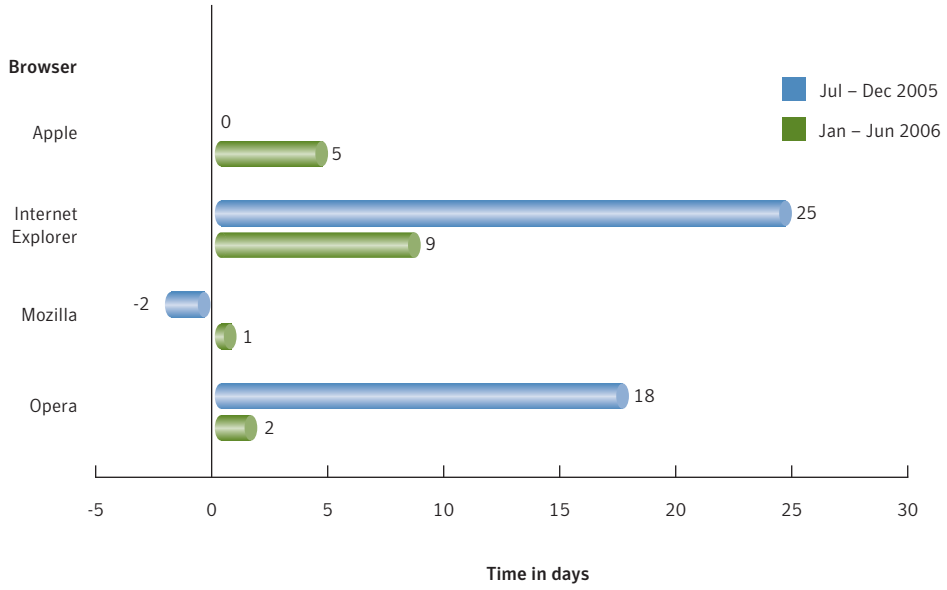

Figure 23. Window of exposure, Web browsers
Source: Symantec Corporation

In the first half of 2006, the window of exposure for most vendors was smaller than for the second half of 2005. Vendor responsiveness appears to be the key factor in this change, particularly as exploit development time averages are still very short.

Average patch development times for browsers are generally shorter than the patch development times in other contexts, such as enterprise and operating system. This is noteworthy because some vendors, such as Apple and Microsoft, are included in all of these metrics. This may indicate that these vendors are giving a higher priority to vulnerabilities in browsers than in other contexts. This may be because of the ubiquity of the Web browser and its high profile as a target for exploitation, effectively forcing vendors such Apple and Microsoft to respond more quickly to browser vulnerabilities.

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. In order to provide protection against the exploitation of unpatched vulnerabilities affecting Web browsers, Symantec recommends that organizations deploy intrusion prevention systems and regularly updated antivirus software at gateways and workstations. Organizations should also closely monitor vulnerability mailing lists and apply necessary patches as required, in a timely manner.

In order to protect against Web browser attacks, Symantec advises users and administrators to upgrade all browsers to the latest, patched versions. To reduce exposure to attacks, 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 and implement ActiveX controls to stop attacks before they can be carried out.

Web browser vulnerabilities

The Web browser is a critical and ubiquitous application that has, in the past few years, been a growing target for vulnerability researchers. Traditionally, the focus of security researchers has been on the perimeter: servers, firewalls, and other assets with external exposure. However, a notable shift has occurred, as researchers are more frequently targeting client-side systems, primarily end-user desktop hosts. As part of this shift toward client-side issues, vulnerabilities in Web browsers have become increasingly prominent.

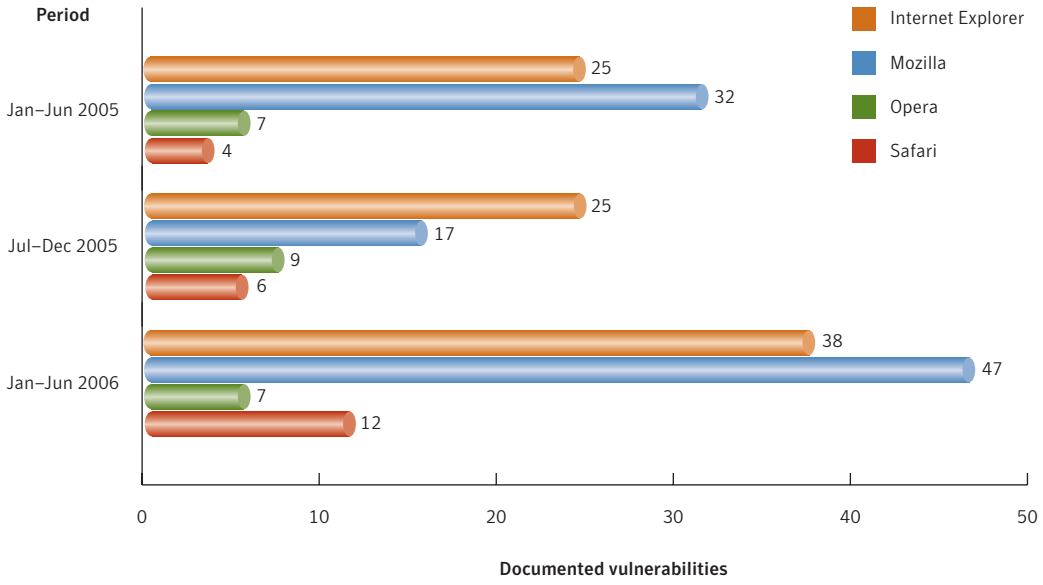

Figure 24. Web browser vulnerabilities
Source: Symantec Corporation

In the first six months of 2006, Symantec documented 47 vulnerabilities that affected Mozilla browsers, including Mozilla Firefox and the Mozilla Browser (figure 24). This is a significant increase over the 17 vulnerabilities that were disclosed in the second half of 2005. The Mozilla Foundation released multiple revisions of Firefox and Mozilla during this period to address the majority of these vulnerabilities.

In the first half of 2006, Symantec documented 38 new vulnerabilities in Microsoft Internet Explorer.[77] This is a 52% increase over the 25 vulnerabilities published in the preceding six-month period. Many of the Internet Explorer vulnerabilities were also reported privately to Microsoft and addressed in cumulative security updates over the course the reporting period. The continued prevalence of Internet Explorer vulnerabilities is likely due to its widespread deployment.

During this reporting period, Symantec documented 12 vulnerabilities that affected Apple Safari, double the six reported in the second half of 2006 and triple the four that were disclosed in the first half of 2006. The sharp increase in the number of Apple Safari vulnerabilities over the past twelve months offers evidence that security researchers are increasingly turning their attention Mac OS X.

[77] It should be noted that this metric does not include third-party components such as ActiveX components or browser plug-ins. However, if the vendor ships their own ActiveX components or browser plug-ins with the browser, vulnerabilities affecting those components will be considered.

Browser fuzzing is an automated vulnerability discovery technique that works by testing the browser against randomly generated input. It has been a noticeable factor in the prominence of browser vulnerabilities, particularly for Internet Explorer and Safari. Both of these browsers have a high number of vulnerabilities that are known to have been discovered using fuzzing techniques. While fuzzing is not a new technique, improved browser-fuzzing tools have made it easier to discover vulnerabilities in more obscure and less-audited code paths. (For more on browser fuzzing, please refer to the "Future Watch" section of this report.)

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 Microsoft Windows and other operating systems in which the browser is not disassociated from many other operating system processes and features.

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. In order to provide protection against the exploitation of unpatched vulnerabilities affecting Web browsers, Symantec recommends that organizations deploy intrusion prevention systems and regularly updated antivirus software at gateways and workstations. Organizations should also closely monitor vulnerability mailing lists and apply necessary patches as required, in a timely manner.

In order to protect against Web browser attacks, 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 and implement ActiveX controls to stop attacks before they can be carried out.

Exploit code release period

While exploit code development times provide an estimate of the time it takes for exploit code to be developed, the exploit code release period measures the period of time after vulnerability disclosure during which associated exploit code continues to be developed and released. This is an important security consideration because a significant number of exploits are published months after the initial disclosure of the affected vulnerability and many of those may be improved versions of earlier exploit code. These exploits may be released after a vendor-developed patch has been made available; however, some enterprises may delay patching a vulnerability for which there is no publicly available exploit code or known attacks. This discussion is limited to exploit code for vulnerabilities that affect enterprise vendors.

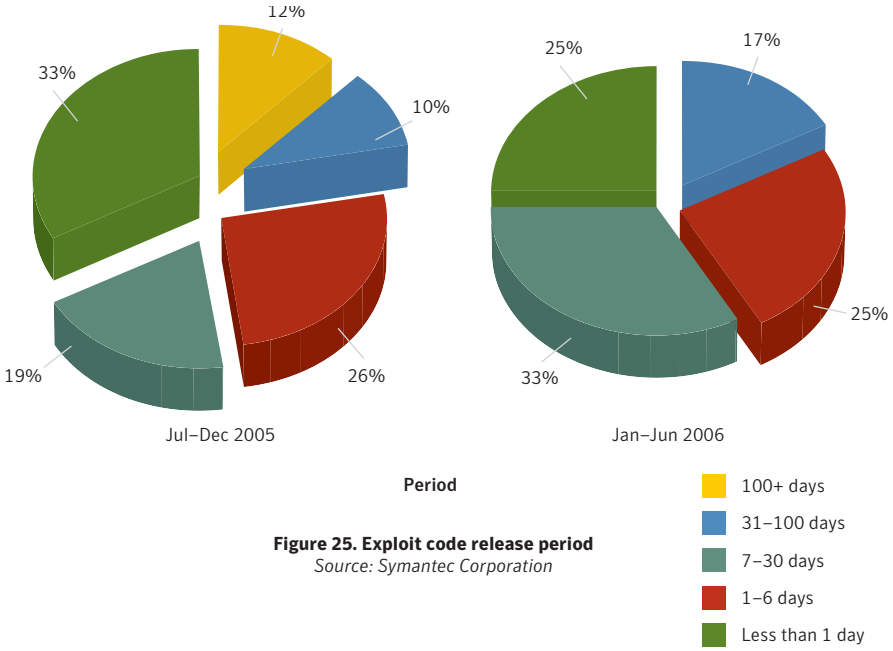

Figure 25. Exploit code release period
Source: Symantec Corporation

Exploit code release periods are broken down into five categories: less than one day, one to six days, seven to thirty days, 31 to 100 days, and more than 100 days. In the first half of 2006, 25% of exploit code was released in less than one day, which is a decrease from 33% in the second half of 2005 (figure 25). During the first six months of 2006, 33% of exploits were released one to six days after initial vulnerability disclosure, an increase over the 19% in the second half of 2005.

The proportion of exploit code released between seven and thirty days was relatively unchanged. During the current reporting period, it was 25%, down slightly from 26% in the second half of 2005. In the first half of 2006, 17% exploits were released between 31 and 100 days after the release of the associated vulnerability. This is an increase over the ten percent proportion in the second half of 2005.

During the first six months of 2006, no exploit code was released in the 100+ day range. In the second half of 2005, 12% of exploit code was released during this period, as was 11% of exploit code in the first half of that year. It should be noted that data for the 100+ day range may change, as exploits may be released after the current reporting period. The data from the current reporting period will likely change after publication of the *Internet Security Threat Report* to reflect the addition of exploit development in the 100+ day range.

The current reporting period was marked by a drop in same-day exploit code and a rise in exploit code published in one to six days. This is likely due to a number of factors. Firstly, some security researchers appear to be withholding proof-of-concept exploit code and technical vulnerability details for a certain amount of time or indefinitely. This is usually due to an agreement that is made with the vendor when the researcher privately reports the vulnerability to the vendor. While this may be effective in limiting the amount of publicly available information about the vulnerability, which could subsequently delay the development of public exploits, it is not an altogether effective protective measure.

Secondly, many exploits that are publicly released come from a source other than the researcher who discovered the vulnerability, so there is frequently a time lapse between the initial vulnerability publication and the release of an exploit. Furthermore, private exploit code may be available for some time before being made public. Since there is no way to measure how long a private exploit has been circulating, the exploit publication date is the date on which the exploit first becomes general public knowledge.

While much of the security emphasis is placed on short exploit code development time, long exploit code release periods are also a cause of concern to organizations. For instance, a vulnerability may initially be considered low risk due to lack of immediate public exploit code and exploitation; because of this, administrators may delay the patching process. If exploit code surfaces a relatively long time after the vulnerability is disclosed, the organization may be caught off guard because the threat was initially perceived as low risk and therefore not addressed in a timely manner. Longer exploit release periods can also result in the development of more reliable exploit code due to the increased time for testing and quality assurance.

Data in other sections of the *Internet Security Threat Report* suggests that older vulnerabilities are still a viable attack vector. When determining the remediation priority, it is important for organizations to evaluate the potential risk of the vulnerability if an exploit is available. Vulnerability managers should incorporate ratings systems that account for the potential risk that vulnerabilities may pose to applications and devices on their network. The Common Vulnerability Scoring System is an example of such a system.[78]

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

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 January 1 and June 30, 2006.

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.

The "Malicious Code Trends" section will discuss:

- Top ten new malicious code families
- Previously unseen malicious code threats
- Malicious code types and worms
- Win32 viruses and bots
- Exposure of confidential information
- Instant messaging threats
- Modular malicious code
- Propagation vectors

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. Additionally, Symantec recommends that ingress and egress filtering be put in place on perimeter devices to detect anomalous 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 operating system vendor. 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.

Top ten new malicious code families

While mass-mailing worms typically dominate the top 50 malicious code samples reported to Symantec, five of the top ten new malicious code families reported during the first six months of 2006 were Trojan horse programs. The rest of the top ten new families consisted of two viruses, two worms and one back door server (table 8).

Rank	Sample	Type	Vectors	Impact
1	Polip	Virus	File sharing, P2P	Lowers security settings
2	Bomka	Trojan, Backdoor	Spam	Drops other malcode
3	Gobrena	Trojan	Spam	Downloads Goldun Trojan
4	Detnat	Virus	Filesharing	Downloads Lineage Trojan
5	Ecup	Worm	P2P	
6	Rajump	Backdoor	N/A	Allows remote access
7	Nebuler	Trojan	N/A	Sends information to remote sites, downloads other threats
8	Awax	Trojan	N/A	Downloads and installs other threats
9	Yamanner	Worm	Yahoo! Web mail	Sends email addresses from contact list to a remote host
10	TopFox	Trojan	N/A	Logs keystrokes

Table 8. Top ten new malicious code families
Source: Symantec Corporation

The most prevalent new malicious code family this period was that of the Polip virus.[79] Polip is a polymorphic virus; that is, it can change its byte pattern when it replicates, thereby avoiding detection by simple string-scanning techniques. (For more discussion on polymorphic viruses, please see the "Future Watch" section in this report.)

Polip attempts to attach its code to all .exe and .scr files on an infected computer when they are opened. The virus also has the ability to update itself through the Gnutella peer-to-peer network to allow remote access to the infected computer. It can also make itself available for download by other users from the Gnutella network even if Gnutella software is not installed on the infected computer. Instead, the virus connects to the network to make itself available to Gnutella clients on the network for download. Finally, Polip tries to lower the overall security of the computer by deleting files related to certain antivirus applications.

The second most common new malicious code family reported between January 1 and June 30, 2006 was Bomka.[80] This Trojan is downloaded from a link that is included in spam email sent by another Trojan program named Spamlia.[81] The email uses social engineering techniques to convince its recipients that the link is the download location for a video clip.

[79] http://www.symantec.com/security_response/writeup.jsp?docid=2006-042309-1842-99
[80] http://www.symantec.com/security_response/writeup.jsp?docid=2006-012514-0250-99
[81] http://www.symantec.com/security_response/writeup.jsp?docid=2005-122917-3955-99

When Bomka is installed on a victim's computer, it uses rootkit techniques in order to obscure its presence.[82] Bomka also allows a remote attacker to gain full access to the compromised computer by including a back door server component. This could result in the exposure of confidential information. This threat attempts to generate revenue for the attacker by installing a Trojan named Adclicker on the infected computer.[83] Adclicker then drives traffic to certain Web sites that simulate clicks on banner advertisements, a practice known as "click fraud."[84]

The third most frequently reported new malicious code family during this reporting period was also a Trojan, Gobrena.[85] Similar to Bomka, this Trojan is most commonly transmitted through spam email. However, Gobrena is sent as an attachment to the spam email instead of being downloaded through an embedded link. When executed, Gobrena simply downloads and executes the Goldun Trojan on the compromised computer.[86] When Goldun is installed, it attempts to steal the user's e-Gold account information.[87] Modular malicious code combinations such as this will be discussed at greater length in the "Modular malicious code" section below.

Previously unseen malicious code threats

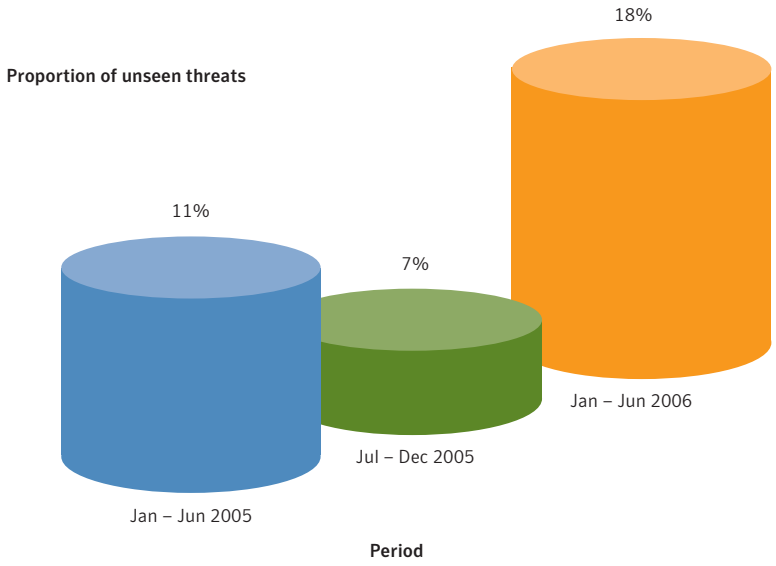

Figure 26. Previously unseen threats as a proportion of all threats
Source: Symantec Corporation

For the first time, in this volume of the *Internet Security Threat Report*, Symantec is tracking the proportion of previously unseen malicious code threats. These are defined as distinct malicious code threats that are detected on Symantec's honeypot computers for the first time before they are detected by other means. This information offers insight into emerging attacker activity, particularly the speed with which attackers adopt new malicious code tools for use against target computers.

[82] A rootkit is a component that uses stealth to maintain a persistent and undetectable presence on the machine. Actions performed by a rootkit, such as installation and any form of code execution, are done without end user consent or knowledge.
[83] http://www.symantec.com/security_response/writeup.jsp?docid=2002-091214-5754-99
[84] Click fraud is the act of using illegitimate means, such as a script or program, to imitate the act of a legitimate user clicking on a pay-per-click banner advertisement on a Web page. This act generates revenue for the owner of the page hosting the advertisement. Click fraud is a felony in some jurisdictions.
[85] http://www.symantec.com/security_response/writeup.jsp?docid=2006-052911-1759-99
[86] http://www.symantec.com/security_response/writeup.jsp?docid=2005-010715-5330-99
[87] e-Gold is an Internet payment system.

Between January 1 and June 30, 2006, 18% of all distinct malicious code samples detected by the Symantec honeypot had not previously been seen (figure 26). A high proportion of previously unseen malicious code likely indicates that attackers are more actively attempting to evade detection by signature-based antivirus and IDS.

One of the major factors contributing to the increase in previously unseen threats is the number of variants within malicious code families. This indicates that attackers are commonly updating current malicious code to create new variants instead of creating new malicious code "from scratch." This is particularly evident in the extremely high number of variants in malicious code families such as the Mytob or Beagle families. Attackers and malicious code writers can create new variants in a number of ways, including metamorphic code evolution,[88] changes to the functionality, and run-time packing utilities. The increase in new threats detected during the first six months of 2006 indicates that attackers may be employing these tactics more actively in order to avoid being detected by antivirus software.

Previously unseen threats are particularly dangerous because traditional defenses, such as some signature-based antivirus products, are typically unable to detect them. Administrators should ensure that their networks are protected by perimeter security tools such as intrusion prevention systems, which will ultimately provide better protection than IDS or firewalls, neither of which will have rules to protect from previously unseen threats. Organizations should also consider network compliance solutions that will help keep infected mobile users out of the network (and disinfect them before entering). Administrators should also be sure to maintain up-to-date antivirus definitions to ensure that their computers are protected from new threats at the earliest possible time.

Malicious code types and worms

In the first six months of 2006, worms continued to dominate the top 50 malicious code reports. They made up 38 of the top 50 unique malicious code samples, accounting for 75% of the volume of top 50 malicious code reports between January 1 and June 30, 2006. This is an increase over 60% in the previous period and 24% in the first half of 2005 (figure 27). (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 numbers of malicious code types in the Top 50 malicious code reports may exceed 50 and cumulative percentages may exceed 100%.)

[88] Metamorphic code evolution describes a method used by malicious code writers that allows a piece of malicious code to change itself autonomously.

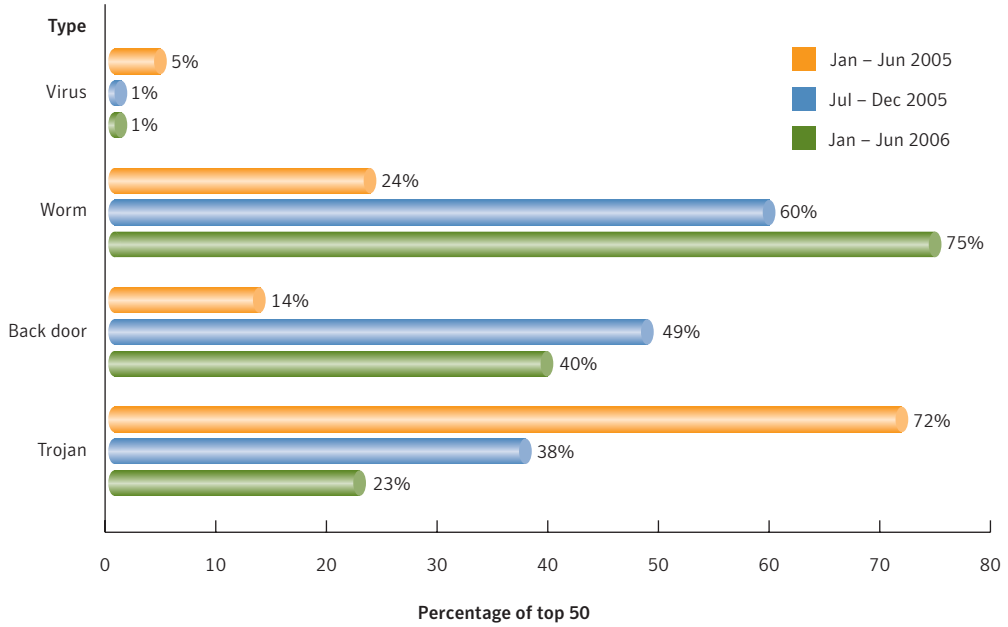

Figure 27. Malicious code types by volume
Source: Symantec Corporation

The increase in worms can largely be attributed to the dominance of mass-mailing worms such as Sober, Netsky, Beagle, and Mytob variants. Additionally, since mass-mailing worms have efficient propagation mechanisms, they are more likely to be reported in high volumes than Trojans, which have no propagation mechanisms.

Back doors were the second most frequently reported malicious code type during the second half of 2006, accounting for 24 of the top 50 malicious code samples. They made up 40% of the volume of the top 50 malicious code reports, a decrease from 49% the second half of 2005, but still significantly higher than the first half of 2005 when they accounted for only 14% of the volume.

The prevalence of back doors in the top 50 samples is due to the number of variants of the Mytob family,[89] which accounted for 16 of the top 50 samples during the first six months of 2006. The slight decline from the previous period is mainly due to the decline in reports of Spybot,[90] Gaobot,[91] and Randex[92] variants, of which only Spybot remains in the top 50 malicious code samples. This will be discussed further in the "Win32 viruses and bots" section below.

While Trojans dominated the malicious code landscape a year ago, making up 21 of the top 50 malicious code samples, they currently account for only ten of the top 50 samples. They also account for less than a quarter of the volume of the top 50 malicious code reported to Symantec during this period.

[89] http://www.symantec.com/security_response/writeup.jsp?docid=2005-022614-4627-99
[90] http://www.symantec.com/security_response/writeup.jsp?docid=2003-053013-5943-99
[91] http://www.symantec.com/security_response/writeup.jsp?docid=2003-112112-1102-99
[92] http://www.symantec.com/security_response/writeup.jsp?docid=2004-072612-2522-99

While some industry observers have claimed that Trojans outnumber worms and viruses overall, this has not been supported by the data that Symantec has received from enterprise and consumer customers worldwide. Due to a lack of propagation mechanisms, Trojans are not likely to be seen by as many users or in such high volume as mass-mailing worms.

Additionally, attackers appear to be making a shift towards targeted attacks using Trojans. Mass-mailing worms tend to use a "shotgun" approach, sending large quantities of themselves to as many users as possible. However, Trojans are now frequently being designed to target specific users and groups. For example, the Mdropper.H[93] Trojan exploited a zero-day vulnerability in Microsoft Word in order to install a variant of the Ginwui back door program.[94] The Word document containing the Mdropper Trojan was spammed to a selected user base using a message with social engineering tailored to entice the users into opening it. Because of the targeted nature of these attacks, the Trojan was sent to a smaller group of users, making it less conspicuous and less likely to be submitted to antivirus vendors for analysis.

Win32 viruses, worms, and bots

Win32 threats are executable programs that operate by using the Win32 API (application program interface), which provides a document interface by which software can interact with different components of the Windows platform. These forms of malicious code work on at least one Win32 platform.[95] For the first time in years, Win32 threats have shown a decline; however, this may be due to changes in Symantec's reporting methods that were made during this period. The change in reporting is due to the fact that Symantec developed new run-time unpacking technology in response to the increase in worm variants, especially those that contain "bot" components, such as the Spybot family,[96] that consists of thousands of variants.[97]

Over the last two years, attackers have intensively utilized run-time packers and wrappers to create new variants in order to "hide" known code from pattern-matching antivirus techniques.[98] Using these tools, attackers could rapidly generate new variants without needing to write new code. As a result, antivirus vendors were required to create new definitions each time the same piece of malicious code was packed or wrapped. Vendors were thus forced to create new antivirus definitions for each variant.

To counter these tactics, over the past six months Symantec has made numerous scanning engine improvements to detect packed threats without needing to create new variant definitions. As a result, Symantec Security Response needs to release far fewer new definitions. Consequently, the number of Win32 variants being identified has decreased.

[93] http://www.symantec.com/security_response/writeup.jsp?docid=2006-051911-0706-99
[94] http://www.symantec.com/outbreak/word_exploit.html
[95] Win32 platforms include Windows 2000 and XP as well as 32-bit versions of Windows 2003 and Vista.
[96] http://www.symantec.com/security_response/writeup.jsp?docid=2003-053013-5943-99
[97] For more on the rise of variants, please see the Symantec *Internet Security Threat Report*, Volume IX (March 2006): http://enterprisesecurity.symantec.com/content.cfm?articleid=1539, pp. 18, 69
[98] A wrapper is similar to a run-time packer but can allow a script file, such as JavaScript, to be presented in executable file format.

New Win32 viruses and worms

As of June 30, 2006, the total number of Win32 variants had surpassed 46,000. In the first half of 2006, Symantec documented 6,784 new Win32 viruses and worms (figure 28), almost 40% less than in the same period last year. This decline is largely attributed to the new technology and reporting features of packed threats in Symantec products that were described above. Due to these changes, Symantec anticipates that there will be a similar decline of Win32 threats for the second half of 2006 as well.

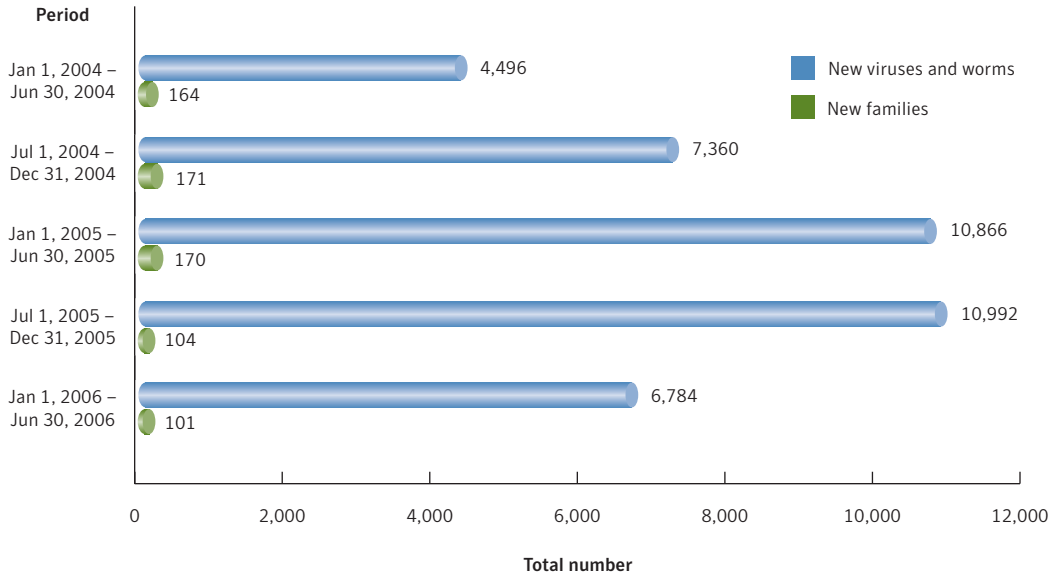

Figure 28. New Win32 virus and worm variants
Source: Symantec Corporation

During the first six months of this year, the number of new Win32 families has also declined. During this period, 101 new Win32 families were detected, down from 104 in the second half of 2005 and 170 in the first half of that year. As discussed in the previous edition of the *Internet Security Threat Report*, the ready availability of source code for various malicious code families makes it easier to modify an existing family to create a new variant than to create an entirely new family.[99] This is likely the reason for the decline in the number of new families over the past two reporting periods.

[99] Symantec *Internet Security Threat Report*, Volume IX (March 2006): http://enterprisesecurity.symantec.com/content.cfm?articleid=1539, pp. 18 and 69.

Win32 bots

In the "Attack Trends" section of the previous *Internet Security Threat Report*, Symantec reported that "the number of bot-infected computers appears to have reached the carrying capacity of its environment."[100] The "leveling off" that was observed at that time continued through this reporting period. In the first six months of 2006, bots accounted for 22% of the top 50 malicious code, up slightly from the 20% reported in the second half of 2005 (figure29).

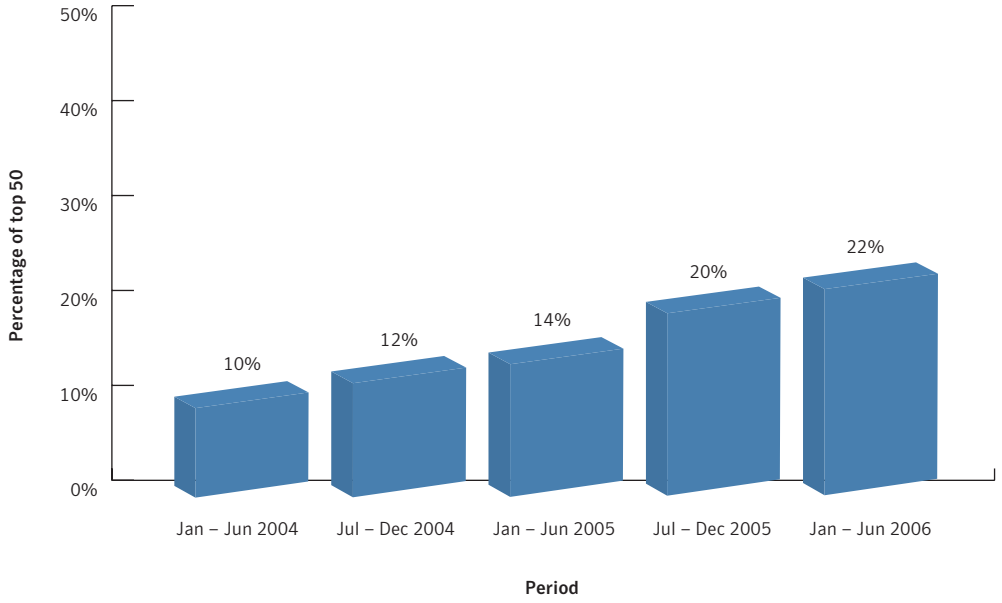

Figure 29. Volume of bots reported
Source: Symantec Corporation

At that time, Symantec speculated that the leveling off of bot infection was due to the widespread and effective implementation of anti-bot security measures, such as firewalls and other perimeter defenses.[101] The apparent success of these measures may have caused malicious code authors to concentrate their efforts on other areas, such as more targeted attacks using Trojans, as was discussed in the "Malicious code types and worms" section above.

It appears that attackers prefer to create variants of existing bots rather than creating entirely new families.[102] There were 23 bot variants in the top 50 malicious code reports belonging to only five different families this period, compared to 21 bot variants from eight different families in the second half of 2006. It is likely that this is due to the desire of attackers to gain the maximum return on their investment of time.

A drop in the number of new remotely exploitable vulnerabilities in Windows services with available exploit code may also be contributing to this leveling off of reported bots. The decreased availability and effectiveness of remotely exploitable vulnerabilities in default services, which is discussed in the "Propagation vectors" section below, has likely necessitated a change in tactics.

[100] Symantec *Internet Security Threat Report*, Volume IX (March 2006): http://enterprisesecurity.symantec.com/content.cfm?articleid=1539, pp. 36
[101] Symantec *Internet Security Threat Report*, Volume IX (March 2006): http://enterprisesecurity.symantec.com/content.cfm?articleid=1539, pp. 36
[102] Symantec *Internet Security Threat Report*, Volume IX (March 2006): http://enterprisesecurity.symantec.com/content.cfm?articleid=1539, pp. 81-82

Attackers appear to have moved away from previous favorites such as Spybot, Gaobot, and Randex to Mytob. This is supported by the fact that, with the exception of Spybot, previously prevalent bots that focus on exploiting service vulnerabilities—such as Gaobot and Randex—are absent from the top 50 malicious code reports in this reporting period. Since Mytob uses SMTP as a propagation vector, as well as the ability to exploit remote vulnerabilities, it is more likely to reach a large number of targets. That said, it should be noted that the discovery of a new remotely exploitable vulnerability with reliable exploit code could easily trigger a resurgence of these bots once the code has been added to their propagation mechanisms.

Exposure of confidential information

Threats that expose confidential information on a compromised computer are a concern to all users, in home, small business, and enterprise environments alike. These threats may expose sensitive data such as system information, confidential files and documents, or cached logon credentials. Some threats, such as back doors, could give a remote attacker complete control over a compromised computer.

Threats to confidential information are a particular concern because of their potential 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.

In the first six months of 2006, 30 of the top 50 malicious code samples exposed a user's confidential information in some way. This is the same number as was reported in the second half of 2005 but ten more than the 20 reported in the first half of 2005 (figure 30).

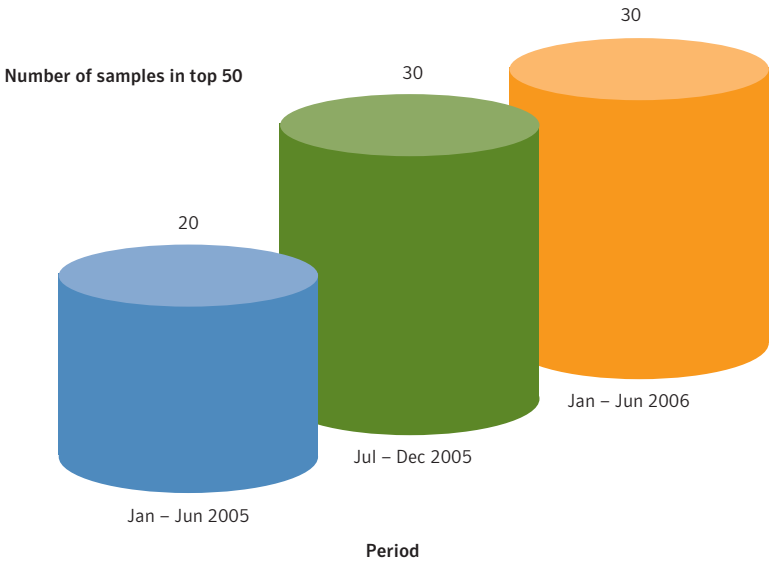

Figure 30. Exposure of confidential information
Source: Symantec Corporation

Symantec believes that the number of threats to confidential information will likely hold steady or increase over the next six months. In the current period, variants of Mytob accounted for 16 of the 30 information-exposure threats in the top 50 malicious code reports. Bots such as Mytob will likely continue to be common amongst the top 50 reported malicious code samples, as their versatility and modularity make them very popular with attackers.

Instant messaging threats

Instant messaging (IM) is widely deployed by users in both home and enterprise environments. However, it is generally unprotected and unmonitored in both contexts, leaving it vulnerable to attacks. This is particularly worrisome for corporate entities, as IM is rapidly becoming a key part of enterprise communications and because confidential information is often exchanged on these networks.

As one of the most successful and widely deployed applications on the Internet, IM has become a potent means for the propagation of viruses, worms, and other threats. The infection of one computer can result in messages being sent to all users in an IM contact list on that machine, creating the potential for rapid proliferation. Furthermore, social engineering tactics are particularly well suited to IM, as the parties communicating over it are inherently trusted.

During the first six months of 2006, AOL Instant Messenger was the IM protocol most commonly used by IM-related malicious code to propagate, accounting for 59%. This is an increase over the 44% of IM-related malicious code that used this protocol in the second half of 2005 (figure 31). It is important to note that just as some malicious code may use multiple propagation vectors, some IM malicious code can also employ multiple IM protocols; as a result, the cumulative percentages presented in this discussion may exceed 100%.

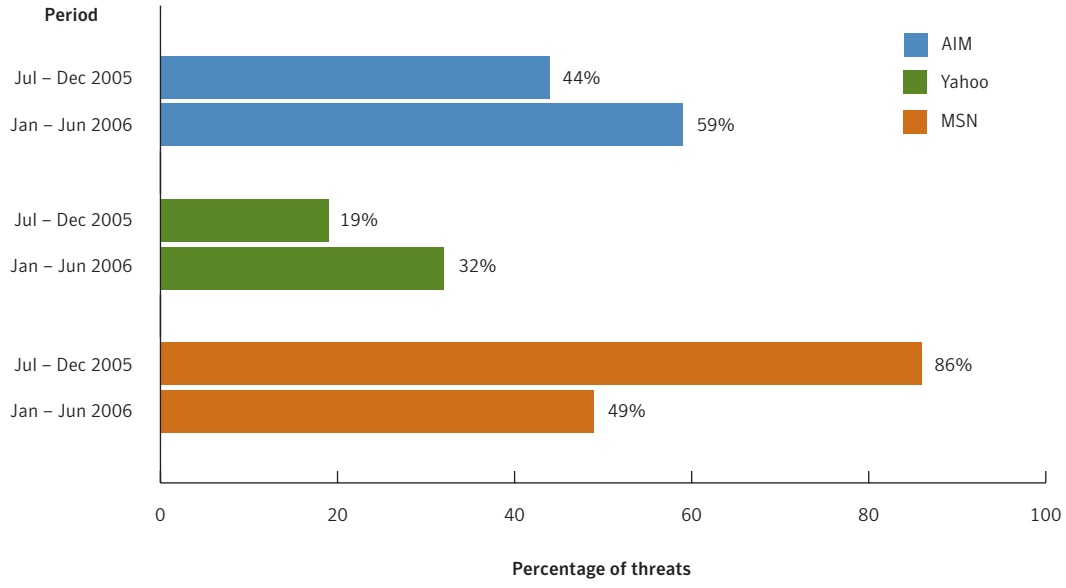

Figure 31. Percentage of instant messaging threat propagation by protocol
Source: Symantec Corporation

Some bots, such as variants of Spybot, Gaobot, and Randex, commonly used AOL Instant Messenger (AIM) to propagate during the first six months of 2006, as did variants of Esbot.[103] This may indicate that it is easier for attackers to incorporate propagation components for this protocol than for others and that a reliable propagation component exists.

The next most frequently targeted IM protocol was MSN Messenger, which was used by 49% of IM-related malicious code this period. This represents a decrease from the previous period when 86% of IM-related malicious code used this protocol. This decline may indicate that changes were made to the MSN protocol that required attackers to write new propagation modules for it. Once the changes have been examined by malicious code authors, they may make necessary adjustments to existing modules, and this protocol may experience renewed malicious code activity.

Yahoo! Instant Messenger was used by 32% of IM-related malicious code to propagate. This makes it the third most frequently targeted protocol for the period. This represents a sharp increase over the 19% of malicious code using this protocol in the previous period.

With the latest releases of Yahoo! Instant Messenger and Windows Live Messenger (formerly MSN Messenger), it was announced that the two protocols would be interoperable.[104] This will allow users from one network to communicate with users of the other without having to install multiple IM clients. This change may also encourage attackers to concentrate their efforts on these protocols, since they will likely enable attackers to reach a larger group of users. It is likely that any malicious code that propagates through one of these protocols in the future will also propagate through the other, allowing attackers to reach a larger user base with minimal effort.

Modular malicious code

In the "Future Watch" section of the September 2005 volume of the *Internet Security Threat Report*, Symantec predicted that malicious code would become a more prominent security issue.[105] Modular malicious code works by compromising a computer and then downloading other pieces of code with added functionalities. It initially possesses limited functionality, such as disabling antivirus software and firewalls, but can update itself with additional code that has new, potentially more damaging capabilities. These may allow it to further compromise the target computer or to perform other malicious tasks.

Modularity in malicious code can serve different purposes. The malicious code may simply attempt to update itself to a more recent version, as is often the case for bots and back door servers. Frequently, modular malicious code is used to download another application to gather confidential information. As previously noted, threats to confidential information may be used by attackers for financial gain. By using modular malicious code, attackers may be able to download and simultaneously install a confidential information threat on a large number of compromised computers.

[103] http://www.symantec.com/security_response/writeup.jsp?docid=2006-070512-0211-99
[104] http://get.live.com/messenger/overview
[105] Symantec *Internet Security Threat Report*, Volume VIII (September 2005): http://enterprisesecurity.symantec.com/content.cfm?articleid=1539, p. 83.

Between January 1 and June 30, 2006, modular malicious code accounted for 79% of the volume of top 50 malicious code reported to Symantec (figure 32). This represents a significant decrease from the 88% reported from July to December 2005. The decline in volume of modular malicious code this period can mainly be attributed to the prevalence of the Blackmal.E worm (also known as the Kama Sutra worm). This worm was the second most widely reported malicious code sample in the current period; however, it did not attempt to download additional components or threats and so is not considered modular. The large volume of Blackmal.E reports thereby caused the overall volume of modular malicious code in the top 50 to decline.

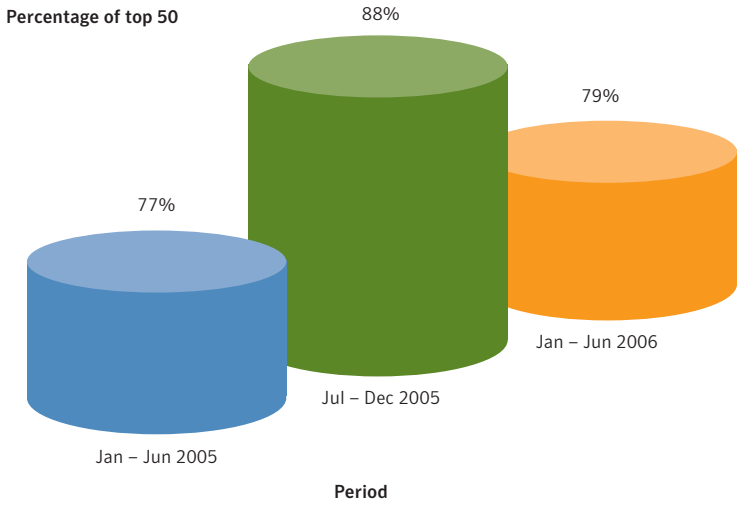

Figure 32. Volume of modular malicious code
Source: Symantec Corporation

While the volume of modular malicious code has declined since the previous period, the number of modular malicious code samples in the top 50 has remained constant. In both the first half of 2006 and the second half of 2005, 36 unique samples were reported to Symantec (figure 33). The most widely reported malicious code sample this period—Sober.X[106]—is a modular malicious code sample that employs a downloader component. The worm contains an algorithm to begin downloading files from a number of Web sites on January 6, 2006 and every week thereafter.

[106] http://www.symantec.com/security_response/writeup.jsp?docid=2005-111915-0848-99

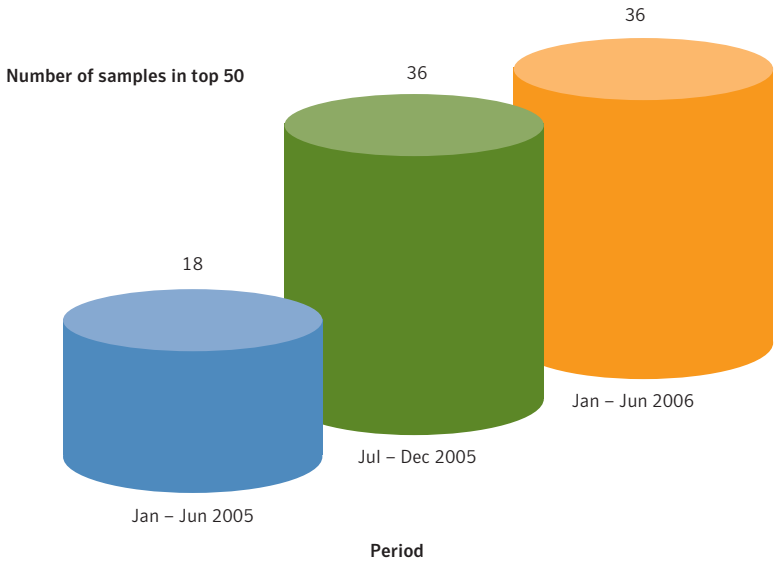

Figure 33. Modular malicious code samples
Source: Symantec Corporation

Propagation vectors

Worms and viruses use various means to transfer themselves, or propagate, from one computer to another. These are collectively referred to as propagation mechanisms. Propagation mechanisms can include a number of different vectors, such as Simple Mail Transfer Protocol (SMTP), Common Internet File System (CIFS), peer-to-peer services (P2P), and remotely exploitable vulnerabilities. Some malicious code may even use other malicious code as a propagation vector by locating a computer that has been compromised by a back door server and using it to upload and install itself. (It is important to note that many malicious code samples employ multiple vectors in an effort to increase the probability of successful propagation, as a result, cumulative percentages included in this discussion may exceed 100%.)

In the first half of 2006, SMTP was the most commonly used propagation vector (figure 34). This is not surprising, as this protocol is heavily involved in the delivery of email, one of the most widely employed applications on the Internet. In addition to being used as a malicious code infection vector, SMTP is also used to send Trojans in spam email.

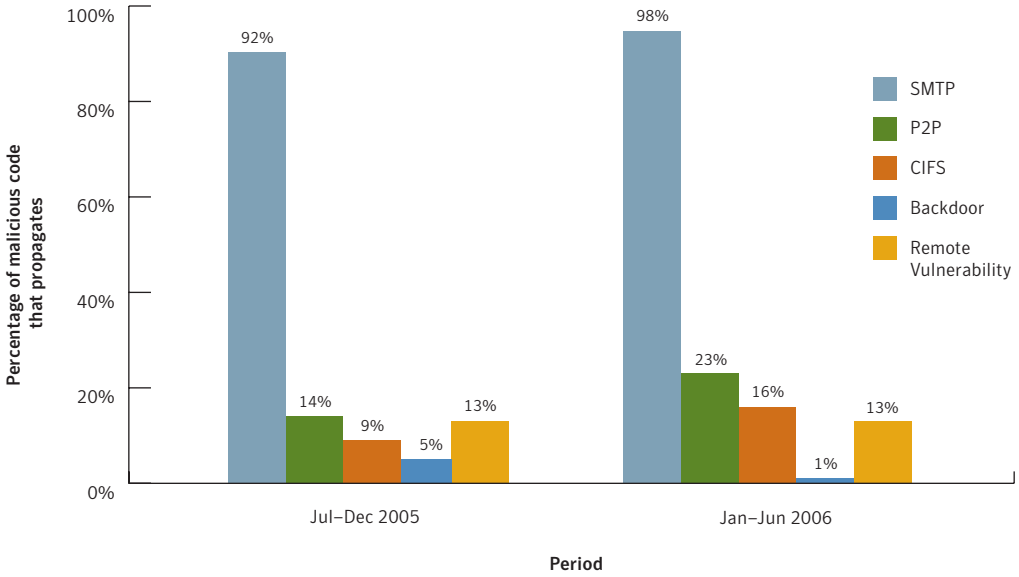

Figure 34. Malicious code propagation vectors
Source: Symantec Corporation

In the first half of 2006, 38 of the top 50 malicious code samples that propagate did so by SMTP, an increase over the 26 in the second half of 2005. Put another way, during this period one out of every 122 email messages scanned by Symantec Brightmail Antispam contained malicious code. Malicious code that propagates by SMTP accounted for 98% of the volume of top 50 malicious code reports with propagation mechanisms this period.

In the first half of 2006, the top 50 malicious code samples was dominated by variants of the Netsky, Beagle, and Mytob worms, all of which are mass-mailing worms. As a result, malicious code that propagated by SMTP accounted for 98% of malicious code in the top 50 samples that propagate (figure 34). This is an increase over the 92% of the volume of the top 50 malicious code reports in the second half of 2005. All of the top ten malicious code samples this period utilized SMTP as a propagation vector, demonstrating the continued effectiveness of this vector. Furthermore, the most prolific mass-mailing worm this period, Sober.X, uses SMTP as its sole propagation vector, as do multiple variants of Mytob.

Organizations can protect against SMTP threats by blocking all email attachments at the mail gateway. If there is a business need for email attachments, only those that are considered safe should be allowed. If other attachment types are accepted, they should always be scanned by antivirus products with up-to-date definitions and should only be accepted from trusted sources.

In the first six months of 2006, six of the top 50 malicious code samples that propagate used CIFS as a vector, accounting for 16% of the total volume. This is a slight increase over the second half of 2005, when seven of the top 50 samples used this vector, accounting for nine percent of the total volume for that period. The rise in the use of CIFS as a propagation mechanism in this period is mainly due to its use by the Blackmal.E worm, which was the second most frequently reported malicious code sample during this period.

As was discussed in the introductory paragraph to this section, some malicious code actually uses other malicious code to propagate. For instance, some variants of Spybot will search for back door servers that are installed on previously compromised computers and use the back door to install themselves. This strategy takes advantage of the fact that if a computer has already been compromised it is likely to have a weak security posture, which could allow subsequent malicious code installations to go undetected.

In the first half of 2006, only one of the top 50 samples that propagate did so by this method, accounting for one percent of the volume of top 50 reports for the period. In the second half of last year, only two of the top 50 malicious code samples that propagate used this vector, accounting for five percent of the volume of top 50 reports. In the first half of 2005, four samples accounting for 35% of the volume of top 50 reports used this vector. As has been noted previously in this discussion, the decline in the current period can likely be attributed to the drop of Gaobot and Randex from the top 50 malicious code reports.

The use of peer-to-peer (P2P) file-sharing networks as a propagation vector for malicious code experienced a slight increase this period. Between January and June 2006, 12 of the top 50 samples, accounting for 23% of the volume of top 50 reports, used P2P networks as a propagation mechanism (figure 34). This is up from eight samples accounting for 14% of the total volume of reports in the previous period.

The slight increase in worms propagating through P2P networks this period can largely be attributed to the recent Feebs worm,[107] as well as the presence of highly reported variants of Beagle[108] and Netsky,[109] both of which utilize this vector. Additionally, two of the top ten new malicious code families—Polip and Ecup[110]—also use P2P as a propagation mechanism. It is likely that P2P will remain a vector that is employed by malicious code authors in the future, but it is unlikely to regain the prominence it achieved in the past.

Malicious code that uses remotely exploitable vulnerabilities to propagate is heavily dependent upon unpatched computers to spread. Use of this vector thus relies upon the discovery of new remote service vulnerabilities that allow code execution. In the current period, eight malicious code samples in the top 50 samples that propagate utilized a remotely exploitable vulnerability to do so (figure 34). This is a slight decrease from the ten samples that used this vector in the previous period but still higher than the five samples using this vector in the same period last year. This is supported by the drop in exploits by type for the server category, from 154 to 129, as was discussed in "easily exploitable vulnerabilities by type" discussion of the "Vulnerability Trends" section in this report.

While fewer unique samples employed this vector in the current period, they appear to have experienced only slightly less success than in the previous period. In the current period, ten percent of the total volume of malicious code samples that propagate were reported to exploit vulnerabilities, compared to 13% in the previous period. This is down from the 38% of reports in the same period last year.

During this period, the majority of the malicious code samples that exploit vulnerabilities to propagate were Mytob variants. Seven of the eight malicious code samples using this vector were variants of this bot. The remaining sample was Spybot.

[107] http://www.symantec.com/security_response/writeup.jsp?docid=2006-013122-5631-99
[108] http://www.symantec.com/security_response/writeup.jsp?docid=2006-020216-2847-99
[109] http://www.symantec.com/security_response/writeup.jsp?docid=2004-032110-4938-99
[110] http://www.symantec.com/security_response/writeup.jsp?docid=2006-053111-0818-99

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 first six months of 2006. In particular, it will consist of three sub-sections, 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, often 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 Symantec *Internet Security Threat Report* will discuss phishing activity that Symantec detected between January 1 and June 30, 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 the Symantec Phish Report Network, an extensive antifraud community in which members 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 message, or lure, can be used many times in different phishing attempts.

A phishing attempt can be defined as an instance of a phishing message being sent to a single user. A single phishing message can be used in numerous distinct phishing attempts, usually targeting different end users. 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.

This section of the Symantec *Internet Security Threat Report* will discuss the following:

- Number of unique phishing messages
- Number of blocked phishing attempts
- Phishing activity by sector
- Number of unique phishing Web sites and/or brands being phished

Number of unique phishing messages

Over the first six months of 2006, the Symantec Probe Network detected 157,477 unique phishing messages (figure 35). This equates to 865 unique phishing messages a day. It represents an 81% increase over the 86,906 unique phishing messages that were detected in the last half of 2005. It is also an increase of 61% over the 97,592 messages detected in the first half of 2005.

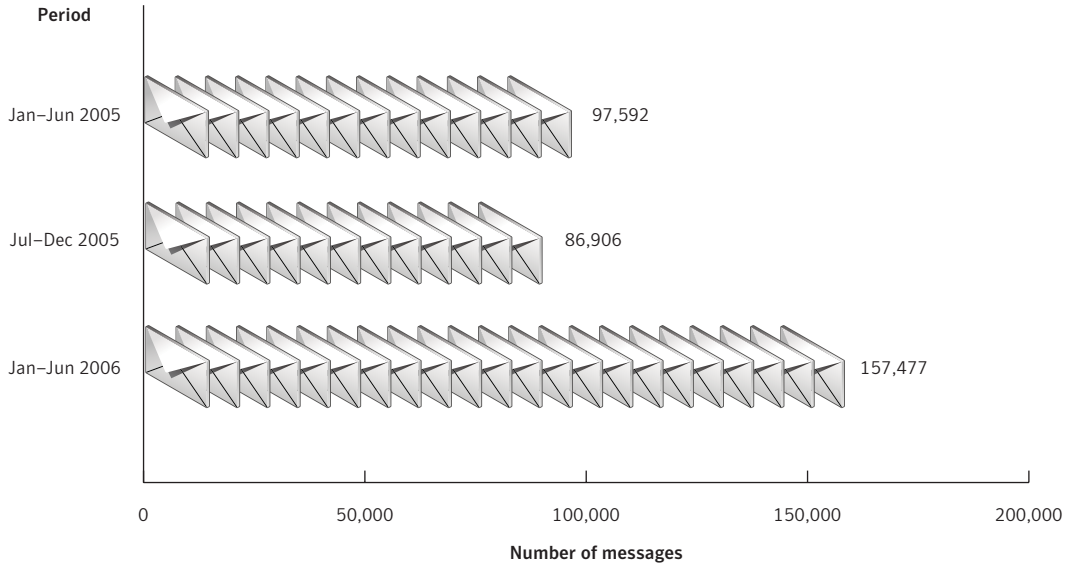

Figure 35. Number of unique phishing messages
Source: Symantec Corporation

The sharp increase over the previous six-month period may be a result of attempts by attackers to bypass filtering technologies by creating multiple randomized messages. These messages may attempt to phish the same brands but include slight variations in order to bypass the use of MD5 checksums or other basic email scanning techniques such as Bayesian filters.[111] These variations often consist of minor changes or differences in the URLs that are included in the email messages. By using a large number of domains in a short period, attackers are able to increase the longevity of each one, making it more difficult for authorities to shut them down because of the amount of effort involved in tracking and taking down each domain used.

[111] An MD5 checksum is obtained when a message is hashed through an algorithm to obtain a unique value. This technique can be used to identify known spam, phishing, and malicious code email messages. Bayesian filters use mathematical probabilities to determine whether a message is spam or not based on the usage of certain words.

Blocked phishing attempts

The number of blocked phishing attempts is derived from the total number of phishing messages that Symantec Brightmail AntiSpam antifraud filters block. Antifraud filters are rules that are created by Symantec Security Response that detect and block known phishing messages. Once the filters have been created they are deployed across the Symantec Brightmail AntiSpam global customer base where they prevent known phishing email messages from reaching end users.

The number of phishing attempts blocked by Symantec Brightmail AntiSpam in the first six months of 2006 indicates a decrease in phishing activity from the previous reporting period. In the first half of 2006, Symantec blocked 1.30 billion phishing attempts, an 11% decrease from the 1.46 billion phishing attempts detected in the last six months of 2005 (figure 36). It is still 25% higher than the 1.04 billion blocked phishing attempts detected in the first six months of 2005.

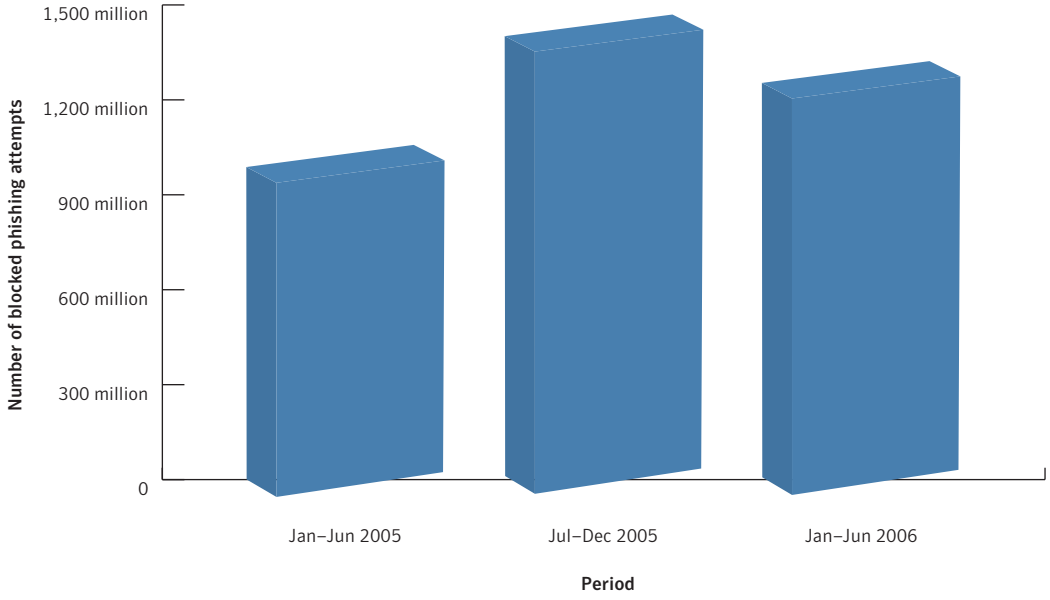

Figure 36. Blocked phishing attempts
Source: Symantec Corporation

Phishing messages that are blocked at the globally distributed mail servers of Symantec Brightmail AntiSpam customers are reflective of phishing activity targeting email users across the Internet. As a result, Symantec believes that the slight decrease in blocked messages may be indicative of more targeted attacks in phishing activity. As noted previously, the number of unique phishing messages is on the rise; this likely reflects an attempt by phishers to bypass current filtering attempts, most of which use previous phishing messages as the basis of detection and subsequent blockage. For this reason, attackers may be sending a higher number of unique messages but in lower volumes and to more focused groups and individuals. For example, if the brand being phished is an Australian bank, the attacker may limit the list of recipients to those with email addresses in the .au domain since those are the users most likely to associate with that brand.

Phishing activity by sector

For the first time, in this edition of the *Internet Security Threat Report*, Symantec is tracking the sectors of companies that are being targeted by phishing attacks. Not surprisingly, the financial sector is the most heavily phished, accounting for 84% of phishing sites tracked by the Symantec Phish Report Network and Symantec Brightmail AntiSpam during this period (figure 37).

As was established in the introduction to this section, phishing is usually conducted for financial gain. Phishing attacks against the financial services sector are most likely to produce the greatest monetary gain for attackers. Once an attacker gains access to a target's account through one of these attacks, he or she may be able to initiate wire transfers to remove funds, or apply for loans, credit lines, or credit cards.

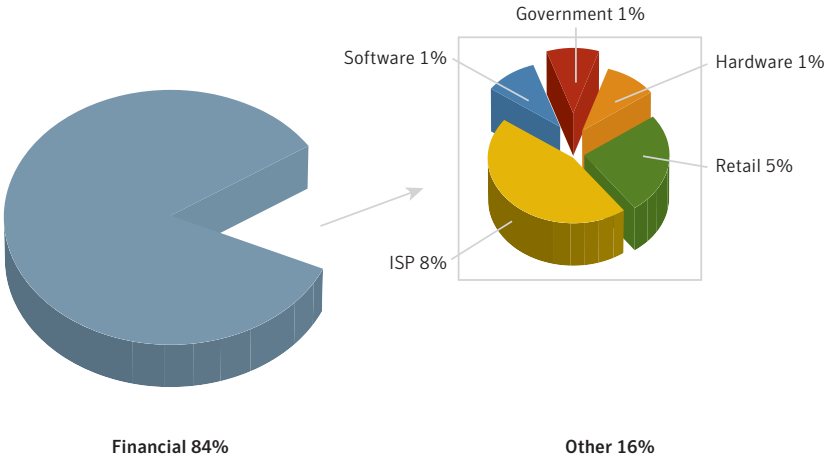

Figure 37. Phishing activity by sector
Source: Symantec Corporation

Phishing activity that targeted Internet service provider (ISP) accounts made up the second largest percentage of attacks this period, accounting for eight percent of the total volume. While access to a user's ISP account may not provide immediate financial gain for the attacker, it could benefit them in other ways. The attacker could use these accounts to access the ISP's outgoing email servers in order to send more spam or phishing messages. Since a major ISP's email servers are less likely to be on DNS blocklists (DNSBL),[112] this tactic increases the probability that the attacker's emails will reach their destination.

The third most widely phished industry in the first half of 2006 was the retail sector. This sector is mostly made up of online retailers or e-commerce sites. Access to a user's e-commerce site account does not provide the immediate financial benefit that an online bank account would, but it could still hold potential rewards for the attacker. The attacker could log on to the user's account and order products by paying with any credit cards that are stored in the system for that user. He or she could then specify a shipping address to which they have access during the checkout process. Once the goods are delivered, the attacker could then resell the merchandise for profit. Since there are more steps involved and the gain is not immediate, it is easy to see why phishing attacks conducted through the retail sector are less desirable to phishers than would be those through financial sector.

[112] A DNS blocklist (DNSBL) is a list of IP addresses that are known to send unwanted email traffic. The DNSBL is used by email software to either allow or reject email coming from IP addresses on the list.`

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. 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.[113]

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.[114] They should also keep their employees notified of the latest phishing attacks and how to avoid falling victim to them.[115]

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.[116] This can be done with the help of companies that specialize in domain monitoring; some registrars even provide this service.[117]

End users should follow best security practices, as outlined in Appendix A of this report. 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.[118] Additionally, network administrators can review Web proxy logs to determine if any users have visited known phishing sites.

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 January 1 and June 30, 2006.

[113] 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.
[114] 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/conline/pubs/alerts/phishingalrt.htm
[115] A good resource for information on the latest phishing threats can be found at http://www.antiphishing.org
[116] "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.
[117] See http://markmonitor.com/brandmanagement/index.html for instance.
[118] http://www.fbi.gov/cyberinvest/inetschemes.htm

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.

This section of the Symantec *Internet Security Threat Report* will explore the following:

- Spam as a percentage of all email
- Top spam categories
- Top ten countries of spam origin
- Percentage of spam containing malicious code

Spam as a percentage of all email

Symantec calculates the percentage of email that is spam by dividing the total number of emails that are identified as spam by Symantec Brightmail AntiSpam filters by the total of the inbound email messages received by the sample customer base. Between January 1 and June 30, 2006, spam made up 54% of all monitored email traffic. This is an increase over the last six months of 2005 when 50% of email was classified as spam. However, it is lower than the first half of 2005, when 61% of email was classified as spam.

While the six-month average was 54%, analysis of the month-to-month spam data reveals a decline and subsequent rise in the percentage of email that was determined to be spam between January 1 and June 30, 2006. In January, 55% of email was categorized as spam. By March this number had declined to 51%, but by the end of June it had climbed back up to 55%.

The temporary mid-term decline likely did not reflect an actual decrease in overall spam activity but instead was likely a statistical anomaly caused by an increase in image spam.[119] Since this type of spam does not contain any text, it is more difficult to block using traditional means. To respond to this, Symantec developed a new class of effective detection technology. After the deployment of this technology, spam activity recorded by Symantec returned to previous levels, indicating that the new measures were effective.

Additional methods can be implemented to block image spam, such as blocking email messages with image file attachments or stripping the attachments at the mail gateway. It should be noted that this would also potentially block any legitimate email messages containing these file types. Administrators should carefully examine the business effects of this type of mitigation before implementing it in the enterprise.

[119] Image spam is a spam email message that does not contain any regular text. Instead, the spam message is implemented as an image, either attached to the email or downloaded from a remote Web site.

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 first six months of 2006 was related to health services and products (figure 38). Health-related spam made up 26% of all spam on the Internet during this time. The next largest spam category was adult spam, which made up 22% of all spam. The next most common type of spam was related to commercial products. It made up 19% of all spam.

It is not surprising that health-related and adult spam make up close to half of all spam. These categories traditionally have the highest "click-through" rates, as they tend to be more difficult to market through more legitimate and traditional means. "Click-through" is a term to describe when a user clicks a link that contains uniquely identifiable information about its originator. Typically, the originator receives financial compensation for each click-through. As spammers have an economic incentive to have a high click-through rate, in order to increase their return on investment, it is reasonable to conclude that they could be changing their content to that which has a higher click-through rates. This in turn makes sending higher volumes of spam in these categories more appealing.

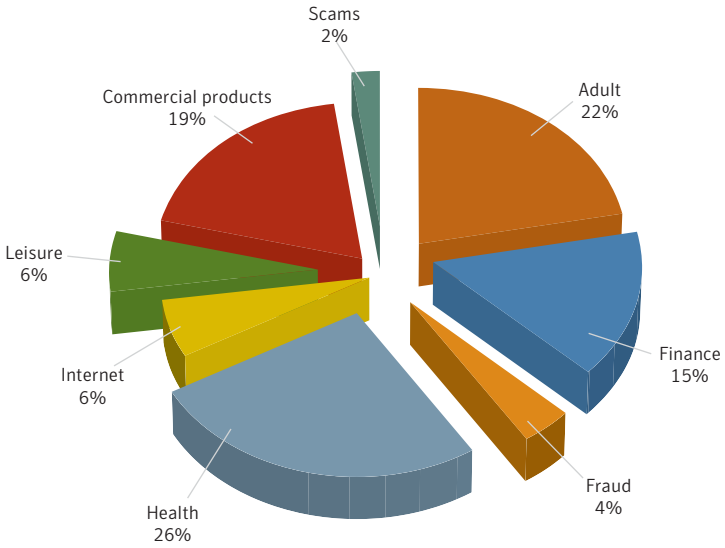

Figure 38. Spam categories
Source: Symantec Corporation

"Adult" spam messages are those that contain pornographic content, sell products of a sexually explicit nature, and/or direct users to a sexually explicit Web site. This category of spam is frequently cited as a concern for organizations because of the need to keep sexually explicit material out of the workplace, primarily for legal issues.

In the previous edition of the Symantec *Internet Security Threat Report*, it was speculated that the amount of adult spam activity would likely increase due to a transition from traditional sexually explicit, HTML-based content to shorter plain-text messages that are more likely to bypass upstream filtering. This prediction appears to have been borne out, as the percentage of adult spam rose during the current reporting period.

Because of the attention it receives, adult spam is often thought to be the most common type of spam. Historically, however, adult spam has only made up around ten percent of all spam. In the current period, however, it accounted for 22%. This rise is likely due to the previously noted transition away from traditional sexually explicit, HTML-based content to shorter plain-text messages.

Top ten 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.

This discussion is based on data gathered by customer installations of Symantec Brightmail AntiSpam. This data includes the originating server's IP address, against which frequency statistics are summarized. Each IP address is mapped to a specific country and charted over time.

During the first six months of 2006, 58% of all spam detected worldwide originated in the United States (figure 39). 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. Since spammers often use bots to send their bulk mailings, this correlation is not surprising. The United States was also the top country of spam origin in the second half of 2005, when 56% of spam originated there (table 9).

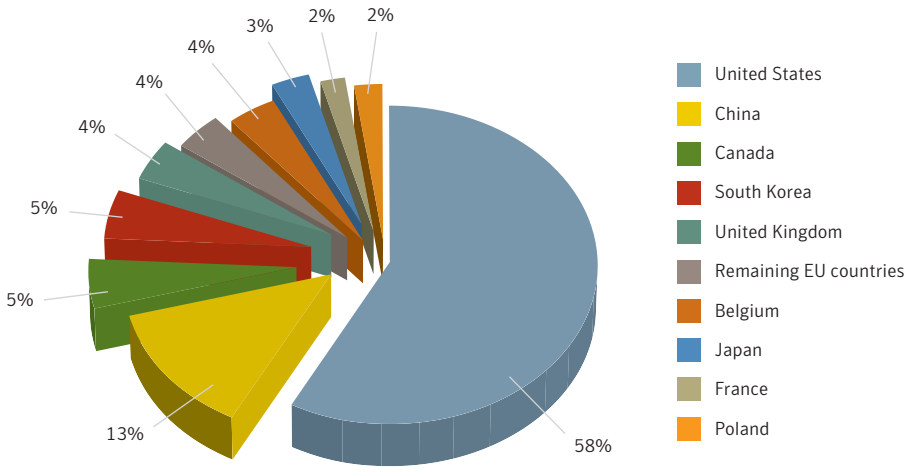

Figure 39. Top ten countries of spam origin
Source: Symantec Corporation

China remained the second highest country of spam origin in the first half of 2006. Thirteen percent of spam during this period originated there, compared to 12% in the second half of last year. Symantec believes that this continuing increase is likely related to technological advancements being made in China, particularly the continued growth in broadband connectivity there. As noted in the "Attack Trends" section of this report, China was also the country with the highest number of bot-infected computers during the first six months of 2006, likely as a result of wider adoption of broadband Internet usage.

Country	Jan–Jun 2006	Jul–Dec 2005
United States	58%	56%
China	13%	12%
Canada	5%	7%
South Korea	5%	9%
United Kingdom	4%	3%
Remaining EU Countries	4%	2%
Belgium	4%	4%
Japan	3%	3%
France	2%	2%
Poland	2%	N/A

Table 9. Top ten countries of spam origin
Source: Symantec Corporation

Percentage of spam containing malicious code

For the first time, in this volume of the *Internet Security Threat Report*, Symantec is assessing the percentage of spam messages that contain malicious code. Malicious code that is delivered by spam is often a bot that can be used in turn to deliver more spam. If an end user has a sufficiently low security posture that they could receive spam messages with malicious attachments in the first place, their computer would also make a good candidate to send spam. Such activity would also be more likely to remain undetected by the user for an extended period of time, although it may be detected by a third party and added to a DNS block list. For organizations, this could prevent users' emails from successfully reaching their intended recipients.

In the first six months of 2006, 0.81% of all spam email contained malicious code. This means that one out of every 122 spam messages blocked by Symantec Brightmail AntiSpam contained malicious code. This is worrisome, as 54% of all email during this period was identified as spam, as was established in the "Spam as a percentage of all email" discussion above.

Since January 2006, spam containing malicious code dropped steadily before rising again slightly in June. At the beginning of the year, 1.27% of spam email contained malicious code compared to 0.56% at the end of June (figure 40).

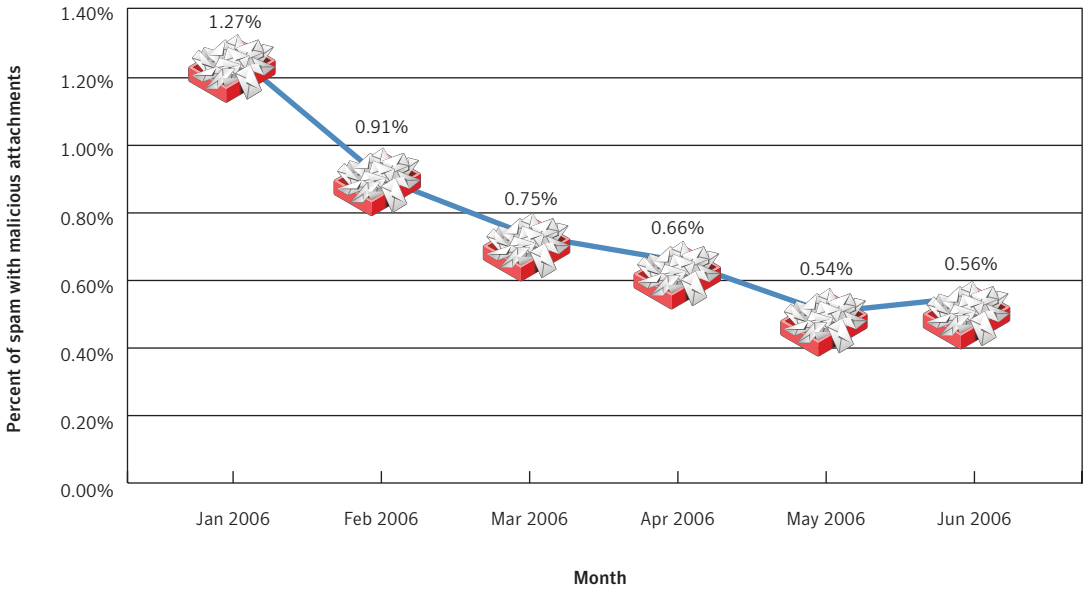

Figure 40. Spam containing malicious code
Source: Symantec Corporation

The five-month decline is likely influenced by two factors. The first is that attaching malicious code to a message increases its chances of being blocked by various means. In some cases, administrators may block all incoming messages with attachments or executable type attachments. Additionally, spam messages with malicious code attachments may be detected by both spam-filtering software and antivirus scanners, decreasing their chances of reaching end users.

The second factor, which is likely a response to the first, is the inclusion of links to Web sites hosting malicious code in spam messages. Rather than attach a malicious code executable to a message, spammers will include a link to a Web site that is hosting malicious code instead. In many cases, the Web site may exploit a client-side vulnerability in the user's browser to install the malicious code without their knowledge or consent. This technique helps reduce the number of messages that are blocked before reaching the end user and still allows the spammer to install malicious code on a recipient's computer.

As discussed in the "Attack Trends" section of this report, bots may be reaching a saturation point. Because it is more difficult for spammers to infect new hosts with their bots, they may also have moved away from attaching them to spam messages because it is simply no longer worth the added effort involved.

To protect against malicious code that is received through spam, users should follow the same precautions used to protect against any malicious code infections. Employing defense in-depth strategies, including the deployment of antivirus software and a personal firewall will help protect against these threats. 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. 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. Finally, users should always perform daily tasks such as browsing Web sites and reading email as an unprivileged user with minimal access rights in order to limit the consequences of a potential malicious code infection.

Security Risks

Symantec uses the term "security risks" to refer to a number of malicious programs, such as adware, spyware, misleading applications, and other programs that users may not want on their system.[120] While security risks are not categorized as malicious code, Symantec monitors them using many of the same methods used for 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.

Top ten security risks

There was little change in the top ten security risks over the first six months of 2006. Most new activity observed during this period consisted of new variants of previously reported security risk programs. The top three security risk programs reported in the first half of 2006 were the same as those reported in the second half of 2005. All three are adware programs, as were eight of the top ten security risks.

Depending upon its functionality and the context in which it is deployed, adware can constitute a security risk. In some cases, these programs may gather information from the user's computer, such as Internet browser usage or other computing habits, and relay this information back to a remote computer. It may also do so by occupying bandwidth, thereby diminishing the functionality and availability of a computing system. Adware can also gather details about the user's computer, which can create a security risk.

[120] Other examples of security risks reported to Symantec in this reporting period are trackware—programs that track system activity, gather system information, or track user habits and relay this information to a third-party organizations—and dialers—programs that use a computer or modem to dial out to a toll number or Internet site, typically to accrue charges.

Between January 1 and June 30, 2006, the most frequently reported security risk program was Hotbar,[121] an adware program that accounted for 24% of the top ten security risks reported to Symantec (table 10). It was the second most frequently reported security risk program in the last six months of 2005.

First detected in 2003, Hotbar adds graphical skins to Internet Explorer, Microsoft Outlook, and Outlook Express toolbars. It also adds its own toolbar and search button to Internet Explorer. These custom toolbars have keyword-targeted advertisements built into them. For example, if a user searches for "mortgages," the toolbar will display mortgage-related advertisements and links from Hotbar's advertising affiliates. Hotbar also monitors Web browsing habits, which may be used for targeted marketing.

Rank	Risk name	Risk type
1	Hotbar	Adware
2	Websearch	Adware
3	BetterInternet	Adware
4	InstantAccess	Dialer
5	NDotNet	Adware
6	Aurora	Adware
7	Lop	Adware
8	Iefeats	Adware
9	Istbar	Adware
10	ISearch	Spyware

Table 10. Top ten security risks
Source: Symantec Corporation

Websearch was the second most frequently reported security risk program over the first six months of 2006.[122] An adware program, it made up 22% of the top ten security risks reported to Symantec during this period. Websearch features a number of noteworthy attributes. It modifies Internet Explorer's default home page and search settings, installs itself as a toolbar to Internet Explorer, and adds a number of icons to the system tray. It also sends user information to a predetermined Web site, including keywords from searches.

One interesting technique that Websearch uses is a "watchdog process," which prevents the manual removal of components of the program.[123] If a user attempts to stop a process associated with the adware program, a second running process restarts it as soon as it has been stopped. This increases the difficulty of removing the program.

BetterInternet was the third most commonly reported security risk in the first half of 2006,[124] making up nine percent of the top ten security risks. It was also the third most common security risk in the second half of 2005. BetterInternet is a browser helper object (BHO),[125] which means that it may display advertisements on the computer on which it is installed. It may also download and install files on the compromised computer, such as other security risks from the same vendor, updates, and/or other applications from the vendor's partners.

[121] http://www.symantec.com/security_response/writeup.jsp?docid=2003-080410-3847-99&tabid=1
[122] http://securityresponse.symantec.com/avcenter/venc/data/adware.websearch.html
[123] http://www.symantec.com/avcenter/reference/techniques.of.adware.and.spyware.pdf
[124] http://securityresponse.symantec.com/avcenter/venc/data/adware.betterinternet.html
[125] 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 that are used to read programs within the browser do so with BHOs. BHOs can also be used to install security risks on a user's Web browser using ActiveX controls.

BetterInternet also gathers system information from the computer on which it is installed. It also sends information that is not personally identifiable—such as running processes, registry entries, hostname, Windows serial number and product ID, network card MAC address, and software version information—to a remote server.

Top ten reported security risks—notable characteristics

Different security risk programs have different characteristics. These may relate to ways in which the program is placed on the user's computer, the ways in which it resists attempts to remove it, and the risks that the program poses to the confidentiality of the user's data. The following sections will discuss some of the characteristics inherent in the top ten security risk programs reported in the first half of 2006.

As the previous section noted, there was little significant change between the top ten security risks detected in each of the last two reporting periods. As a result, there has been little significant change observed in techniques used by the most common security risks. Instead, these applications are still using similar tricks to those observed in the last six months of 2005.

The lack of new techniques used by the top ten most security risks may indicate that the creators of security risks are switching their focus to areas requiring less effort for greater return. As security vendors have improved their products to deal with the various tricks used by security risk vendors, circumventing security products requires increasingly sophisticated ways to install a security risk on a system and prevent easy removal. This would appear to be borne out by the sharp increase over last six months in what Symantec terms "misleading applications," which will be discussed at greater length in "Top ten new security risks" section below.

Anti-removal techniques

Security risks may implement different techniques to resist attempts to remove them from the user's computer. In the first six months of 2006, five of the top ten security risks employed various techniques to avoid removal from systems (table 11). The following paragraphs will describe some of the anti-removal techniques that Symantec has observed over the past six months.

Risk name	Risk type	Anti-removal technique
Hotbar	Adware	N/A
Websearch	Adware	Watchdog processes
BetterInternet	Adware	N/A
InstantAccess	Dialer	N/A
NDotNet	Adware	N/A
Aurora	Adware	Process injection
Lop	Adware	Auto-updates with automatically repacked versions
Iefeats	Adware	Hides part of program in alternate data stream
Istbar	Adware	Exclusive file lock
ISearch	Spyware	Hooks Kernel APIs to prevent its files from being removed

Table 11. Anti-removal techniques in top ten security risks
Source: Symantec Corporation

Some security risks inject their own code into processes running on the system to make themselves more difficult to remove. This can cause system instability, degrade performance, and reduce security. It may also allow the security risk program to run with the same permissions as the program into which it has been injected. This can make it very difficult for administrators or users to remove these programs manually, without specialized tools or in-depth knowledge. Of the top ten security risks this period, only the adware program Aurora deployed process injection.

Run-time packers are programs that are used to reduce the size of programs.[126] As a result, the programs require less time to download. Run-time packers can also obfuscate the content of a file, so that it cannot be easily recognized by antivirus or antispyware programs, unless they understand the packer format. This technique is commonly used by creators of adware and spyware programs, as well as malicious code authors. For instance, the adware program Lop is dynamically repacked each time it is downloaded, thereby making detection and removal more difficult.

As was discussed in the "Top security risks" section above, watchdog processes may be used by a security risk to avoid removal. They do this by allowing security risks to monitor each other. If one process is stopped, a second process automatically restarts it, and vice versa. Of the top ten security risks reported this period, only the adware program Websearch used watchdog processes to resist removal.

ISearch uses a slightly different anti-removal technique. It hooks kernel mode APIs (application program interfaces)[127] to check if the user is attempting to delete a file or registry key associated with it, and returns access denied, preventing removal of its components. ISearch was not included in the top ten security risks in the last half of 2005.

Stealth techniques

Some security risks use stealth techniques to hide from antivirus and antispyware scanners. Of the top ten security risks reported this period, IEFeats uses a stealth technique whereby it hides part of itself in an alternate data stream.

Alternate data streams were created by Microsoft to provide compatibility with Apple's HFS file system in order to allow Macintosh files to be copied to Windows fileshares without being corrupted.[128] Alternate data streams are not typically scanned by many security products. Attackers can use a simple technique to create an alternate data stream to hide content within otherwise innocuous files.[129] The adware program IEFeats uses this technique to hide some of its files.

Symantec has determined that while this technique is not in common use among security risk programs, it has been used for a number of Trojan horse programs, such asRustock.B,[130] Comxt.B,[131] and Fugif.[132] While it is possible that this technique could become more widespread, given the fact that it has been in the public domain for a number of years, a sharp increase in its usage seems unlikely.

[126] For more on run-time packers, please see the "Win32 viruses, worms, and bots" discussion in the "Malicious Code Trends" section of this report.
[127] Kernel mode APIs (application programming interfaces) are part of the Microsoft Win32 API. A detailed description of the Win32 API and of kernel mode is outside the scope of this report; however, suffice it to say that these are low-level system calls, which are associated with commands to delete files, which the security risk intercepts to prevent its deletion from the system.
[128] Alternate data streams were provided as part of the NTFS file system for Windows NT and later versions of Windows to provide compatibility with Apple's old Hierarchical File System (HFS). Files on HFS consist of a data fork, containing the contents of the file, and the resource fork, containing metadata, such as file type and other relevant details. A common problem when copying HFS files to the Windows FAT or FAT32 file system was that the resource fork information would be lost, thereby corrupting the file.
[129] More information on alternate data streams may be found at the following Web sites: http://www.symantec.com/avcenter/reference/ntfs.streams.a.primer.pdf and http://www.securityfocus.com/infocus/1822
[130] http://www.symantec.com/avcenter/venc/data/backdoor.rustock.b.html
[131] http://www.symantec.com/avcenter/venc/data/trojan.comxt.b.html
[132] http://www.symantec.com/avcenter/venc/data/pf/downloader.fugif.html

Self-updating

Programs that are used to detect and remove adware programs often do so by using signatures that are based on known characteristics of the adware. As a result, adware vendors will often update the program in order to alter those characteristics, thereby evading those signatures. If the software is updated, then signature-based antispyware products are less likely to recognize it and therefore may not be able to remove it. In some cases, the functionality of the adware program may also be updated. Table 12 lists the top ten most frequently updated security risks in the first half of 2006.

Risk name	Risk type	Updates per days
Dial Platform	Dialer	11.9
ZangoSearch	Adware	10.7
Aurora	Adware	8.3
Sfonditalia	Dialer	7.5
SpySheriff	Adware	6.2
Istbar	Adware	3.6
Lop	Adware	3.6
BetterInternet	Adware	2.9
SurfSideKick	Adware	2.9
DollarRevenue	Adware	2.7

Table 12. Top ten self-updating security risks
Source: Symantec Corporation

Top ten new security risks

Three of the top ten new security risks detected during the first six months of 2006 are what Symantec calls "misleading applications". 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 users to pay money to purchase software or upgrade to a version of security software that will purportedly remove the "threats" that were found. This is a becoming an increasingly common tactic. Misleading applications accounted for 50% of the volume of the top ten new security risks reported to Symantec in the first half of 2006.

Misleading applications can constitute a security risk for a number of reasons. First, the consumer will likely get little or no security protection from the upgraded "security software." The purchase of the upgraded software therefore may give users a false sense that their computer is secure, which may be worse than having no security at all. Second, in purchasing the upgrade, the user will likely have disclosed his or her credit card information to the owner of the misleading application, who may then be able to use it for further fraudulent purposes.

Third, the initial downloader program that installed the misleading application may download other security risks or malicious code onto the target system. Many downloaders can be reconfigured to download other programs from different locations, meaning that they could potentially open the doorto a wider variety of programs being installed on the target system. Finally, misleading applications also represent a threat to organizations because of the time and effort that may be wasted in removing such applications from users' systems.

Between January 1 and June 30, 2006, the most frequently reported new security risk was ErrorSafe,[133] a misleading application that accounted for 30% of the volume of the top ten new security risks (table 13). ErrorSafe gives exaggerated reports of threats on the computer and then prompts the user to purchase a registered version of the software in order to remove the reported threats.

Rank	Risk name	Risk type
1	ErrorSafe	Misleading Application
2	DesktopMedia	Adware
3	SpyFalcon	Misleading Application
4	NewWeb	Adware
5	AdvertMen	Adware
6	FCHelp	Adware
7	Caishow	Adware
8	BMCentral	Adware
9	ActivShopper	Trackware
10	MalwareWipe	Misleading Application

Table 13. Top ten new security risks
Source: Symantec Corporation

DesktopMedia was the second most common new security risk reported to Symantec in the first half of 2006.[134] This adware program installs a download manager toolbar for Internet Explorer and displays advertisements from a Chinese Web site. It accounted for 30% of the reports in the top ten new security risks in the first half of 2006.

SpyFalcon was the third most common new security risk that Symantec detected in the first six months of 2006.[135] Like ErrorSafe, it is a misleading application. It accounted for 19% of reports of the top ten new security risks.

In order to mitigate the threat posed by misleading applications, Symantec recommends that administrators and users follow the recommended best practices outlined in Appendix A of this report, and exercise caution when installing applications that purport to solve security issues. Enterprises should only install applications that have been reviewed and certified as legitimate applications. Any application should only be deployed as part of an approved security policy.

Security risks—prevention and mitigation

In order to protect against security risks such as adware, spyware, and misleading applications, Symantec recommends that all users continue to update their antivirus software regularly. Security administrators should also take extra measures to ensure that patch levels on all computers are up-to-date. Organizations can develop and implement “whitelists” of permitted applications that are known to be trustworthy.

Symantec recommends that users and administrators employ defense in-depth, including the use of a properly configured firewall, regularly updated antivirus, and IDS. Symantec also advises users to exercise caution when installing any software through a Web browser and to not download any software from sources that are not known and trusted.

[133] http://www.symantec.com/security_response/writeup.jsp?docid=2006-012017-0346-99&tabid=1
[134] http://www.symantec.com/security_response/writeup.jsp?docid=2006-050112-5838-99
[135] http://www.symantec.com/avcenter/venc/data/spyfalcon.html

Some security risks are installed using ActiveX controls. Symantec recommends that users either disable ActiveX or use a Web browser that does not support ActiveX. However, as was also stated earlier, some users may require ActiveX for some applications, in which case they should configure their browser to require a prompt for ActiveX controls to execute.

Symantec recommends that organizations implement and enforce acceptable usage policies. System administrators should regularly audit the system to ensure that no unauthorized software is installed or operating on the system. Furthermore, administrators and end users should read the end-user license agreements (EULAs) of all software programs before agreeing to their conditions.

One final note of caution should be raised. Symantec recommends that users exercise caution when removing spyware. Programs should be removed as non-intrusively as possible in order to minimize any problems that might result from the removal of the program. In order to avoid such problems, it may be necessary to ignore some non-critical aspects of these programs. Some components, such as registry keys, may also be used by other legitimate programs. Thus, if the artifacts are non-critical, it will not cause harm to leave them behind in the uninstall process.

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 clean them up before entering).
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. Both spyware and adware can be automatically installed on computers along with file-sharing programs, free downloads, and freeware and shareware versions of software, or by clicking on links and/or attachments in email messages, or via instant messaging clients. Ensure that only applications approved by the organization are deployed on the desktop.

Consumer Best Practices

1. Use an Internet security solution that combines antivirus, firewall, intrusion detection, and vulnerability management for maximum protection against malicious code and other threats.
2. Ensure that security patches are up-to-date and that they are applied to all vulnerable applications in a timely manner.
3. Ensure that passwords are a mix of letters and numbers. Do not use dictionary words. Change passwords often.
4. Never view, open or execute any email attachment unless the attachment is expected and the purpose of the attachment is known.
5. 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. All computer users need to know how to recognize computer hoaxes and phishing scams. Hoaxes typically include a bogus email warning to "send this to everyone you know" and/or improper technical jargon that is intended to frighten or mislead users. Phishing scams are much more sophisticated. Often arriving in email, phishing scams appear to come from a legitimate organization and entice users to enter credit card or other confidential information into forms on a Web site designed to look like that of the legitimate organization. Computer users also need to consider who is sending the information and determine if the sender is a trustworthy, reliable source. The best course of action is to simply delete these types of emails.
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. Be aware of the differences between adware and spyware. Adware is often used to gather data for marketing purposes and generally has a valid, benign purpose. Spyware, on the other hand, may be used for malicious purposes, such as identity theft.
10. Both spyware and adware can be automatically installed on a computer with the installation of file-sharing programs, free downloads, and freeware and shareware versions of software, or by clicking on links and/or attachments in e-mail messages, or via instant messaging clients. Therefore, users should be informed and selective about what they install on their computer.
11. Don't just click those "Yes, I accept" buttons on end-user license agreements (EULAs). Some spyware and adware applications can be installed after an end user has accept the EULA, or as a consequence of that acceptance. Read EULAs carefully to examine what they mean in terms of privacy. The agreement should clearly explain what the product is doing and provide an uninstaller.
12. 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. Both the Symantec DeepSight Threat Management System and the Symantec™ Managed Security Services refer to attacks in the same way, enabling analysts to combine and analyze attacks together. Symantec combines data derived from these sources for analysis. In some cases, only one data source is used if attributes required for a particular analysis are not available in the other.

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 enquiries

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.

Top Web browser attacks

Symantec identifies and ranks attacks that are detected being carried out against Web browsers across the Symantec Deepsight Threat Management System and Managed Security Services base. This ranking is representative of the distribution of attacks that the average Web browser user can expect to observe. Symantec derives this rank by determining the proportion of IP addresses that carry out each attack, as this gives the best insight into the popularity of the attack.

Wireless threats

Symantec identifies and ranks threats posed against a sample of wireless networks using wireless security tools. This ranking is representative of the relative popularity of each threat, and represents what a typical wireless network administrator is expected to observe. The threats identified against wireless networks for this metric do not deal with attacks against vulnerabilities on computers deployed on wireless networks, but threats against the wireless network infrastructure itself. The threats are ranked according to the number of threats observed, giving greatest insight into threatening activity posed against wireless networks.

Denial of service attacks

Although there are numerous methods for carrying out denial of service (DoS) attacks, Symantec derives this metric by measuring denial of service attacks carried out by flooding a target with SYN requests, 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, causing 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. Backscatter is only one method of obtaining DoS statistics and for the purposes of this report is only intended to provide a high-level overview of overall DoS activity. Although the values Symantec derives from this metric will not identify all DoS attacks carried out, it will provide insight into high-level DoS attack trends.

SYN flood attacks should not be confused with other types of DoS attacks. ICMP flooding is another method of carrying out a DoS attack.[136] This attack is carried out by bombarding a target computer with ICMP messages until it becomes overwhelmed by them, so that it cannot service legitimate requests. ICMP flooding is also employed when carrying out Smurf DoS attacks.[137] UDP flooding is another popular form of DoS attack. This type of attack is typically carried out by flooding a target with an excessive number of UDP packets in an attempt to tie up the network resources of the target computer so that it cannot service legitimate requests.

There are other types of DoS attacks, most of which are based on the exploitation of vulnerabilities in target services. In most cases, sending a malformed message to a target computer hosting a vulnerable service may cause it to crash or freeze, subsequently denying service to legitimate users.

To determine the countries targeted by DoS attacks, Symantec cross-referenced 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. Targeted sectors were identified using the same methodology as targeted countries; however, attackers considered were those carrying out a set of denial of service attacks that were detected by IDS and IPS software.

Bot networks

Symantec identifies certain scanning patterns and network traffic and cross-references this traffic with rules that define specific coordinated scanning behavior, which would indicate bot network activity. For this volume of the *Internet Security Threat Report*, Symantec implemented a new behavioral matching scheme to expand our view into potential bot threats.

For an originating computer to be flagged as participating in this coordinated scanning, it must fit into that scanning 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 coordinated way as a bot network. This behavioral matching will, however, identify many of the most coordinated and aggressive bot-infected computers and ultimately will give insight into the population trends of bot network computers.

Top bot network countries and cities

Using the data derived from the “Bot network” discussion of the “Attacks Trends” 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. The data produced is then used to determine the global distribution of bot-infected computers.

[136] Internet Control Message Protocol (ICMP) is employed by the TCP/IP stack to handle error and control messages. Its most commonly known functionality, and that exploited by ICMP Flood attacks, is the Echo Request, Echo Reply sequence used by ping utilities.
[137] http://securityresponse.symantec.com/avcenter/venc/data/smurf.dos.attack.html

Top originating countries

Symantec identified the national sources of attacks by automatically cross-referencing source 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.

Currently, Symantec cross-references source IP addresses of attacks against every country in the world. It is important to note that while Symantec has a reliable process for identifying the source IP address of the host that is directly responsible for launching an attack, it is impossible to verify where the attacker is physically located. It is probable that many of the sources of attack are intermediary systems used to disguise the attacker's true identity and location.

Top targeted sectors

For the purposes of the *Internet Security Threat Report*, a targeted attacker is one that is detected attacking at least three users or organizations in a specific sector, to the exclusion of all other sectors. Figure 41 represents the sector breakdown of the sensor distribution in the sample set in percentage terms. Sectors with less than ten sensors have been excluded from the resulting totals.

The targeted sector attack rate is a measure of the percentage of total attackers that target only organizations or users in a specific sector and is represented as a proportion of all targeted attacks. It can indicate which sectors are more frequently the targets of focused attacks. This metric may be affected by the overall attack rate experienced by each sector; nevertheless, it provides an indication of the interest that a sector holds for targeted attackers.

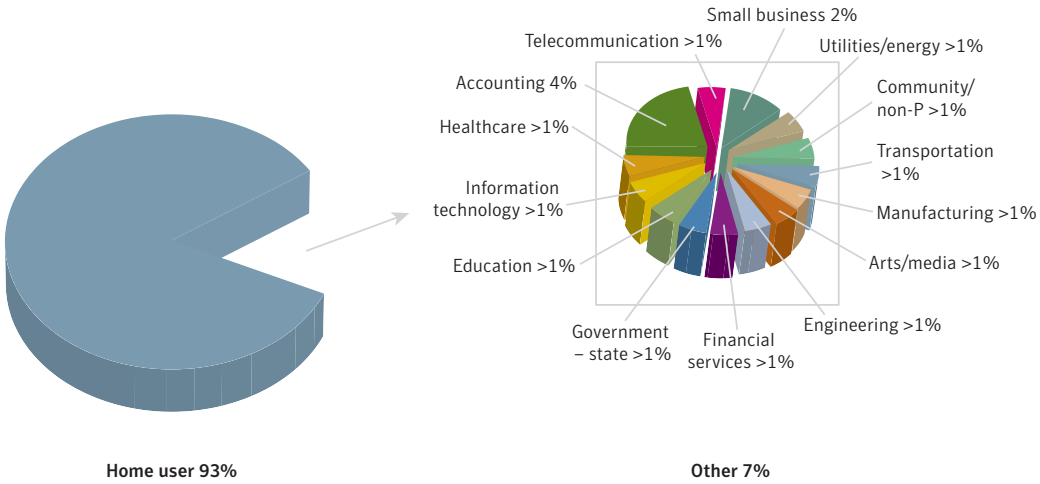

Figure 41. Industry representational breakdown
Source: Symantec Corporation

Appendix C—Vulnerability Trends Methodology

The "Vulnerability Trends" report of the Symantec *Internet Security Threat Report* discusses developments in the discovery and exploitation of vulnerabilities over the past six months. 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 18,000 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
- Web application vulnerabilities
- Easily exploitable vulnerabilities (Total, and breakdown by type)
- Patch development time (Enterprise, Operating System, Browser)
- Exploit development time (Enterprise, Operating System, Browser)
- 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),[138] to define classifications of vulnerabilities. Possible values are indicated below, and 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
- Race condition error
- Serialization error
- Atomicity error

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

- Environment error
- Configuration error
- Design error

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 from previous versions of the *Internet Security Threat Report*. This change was made to accommodate adoption of the exploitability rating in the Common Vulnerability Scoring System (CVSS).[139]

All vulnerabilities are classified into one of four possible categories defined by the CVSS, listed 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.
  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 the easily exploitable vulnerabilities by type. To provide further insight into the types of vulnerabilities that are considered easily exploitable, Symantec has categorized these vulnerabilities into several categories. They are as follows:

- **Browser vulnerabilities:** These vulnerabilities threaten Web browser applications through remote attack vectors.

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

- **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 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 to exploit. 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 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 Applications:** These vulnerabilities affect applications that are deployed on a Web server platform variety of some sort. Such applications are usually in a server side scripting language such as PHP or ASP.NET and accessed through the HTTP/HTTPS protocols.
- **Other:** There are vulnerabilities that do not discretely fall into the above categories. This can include applications for which the distinction is blurred between server and client, or hardware platforms where the affected component cannot be described by any of the other categories.

The specific categories themselves were devised so that the majority of vulnerabilities could easily be classified with little overlap between categories so that the total percentage of all categories equals 100%. These categories are defined in general by the attack vector and by the type of application that is threatened.

Operating system patch development time

This metric has a similar methodology to the “Operating system, enterprise vendors” metric, which was explained previously 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. This metric is incorporated when computing the “window of exposure”, which amounts to the patch development time average minus the exploit development time average.

Window of exposure, enterprise vendors

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

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 (these form the majority), only exploits for technologies from enterprise vendors (that is, those that generally have widespread deployment) are included. Those vendors are:

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

Patch development time, enterprise vendors

The time to patch metric measures the time lapse between the disclosure date of a vulnerability and the release date of a patch that is developed to repair that vulnerability. 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. An average from the aggregate of these is computed for each period. 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.

[140] 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.

Exploit code development time, 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 reference found (such as a mailing list post). The date of exploit code publication is the date of the first reference to the exploit code found. As 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 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 may materialize later that improve upon initial exploits. 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 average.

Window of exposure, Web browsers

This metric has a similar methodology to the “Window of exposure, 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:

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

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.[141] 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 average and the average exploit code development time average. During this time, the computer or system on which the affected application is deployed may be susceptible to attack, as administrators may have no

[141] 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.

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

Web browser patch development time

An average is calculated from the patch release time for each vulnerability affecting each Web browser during the reporting period. The patch development time average for each browser is then compared. This metric is incorporated when computing the "window of exposure," which amounts to the difference between the average patch development time and the average exploit code development time.

Web browser exploit development time

An average is calculated from the exploit release time for each vulnerability affecting each Web browser during the reporting period. The exploit development time average for each browser is then compared. This metric is incorporated when computing 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 were computed for this report. This includes vulnerabilities that have been confirmed by the vendor and those that are not vendor confirmed. This version of the *Internet Security Threat Report* differs from the previous version in that vulnerabilities that are not confirmed are also included in the data. These vulnerabilities were found to be statistically significant, especially given the disparity in patch times between vendors. This version of the report does not differentiate between vendor-confirmed and non-vendor-confirmed 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.

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 are detected on Symantec's honeypot computers for the first time before they are detected by other means. The proportion of previously unseen malicious code threats is derived by comparison with the total number of distinct malicious code threats observed.

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. This section will examine developments in these risks over the first six months of 2006. In particular, it will consist of three sub-sections, which will discuss:

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

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 attacks. 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 enquiries

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.

Six-month growth in phishing messages

Symantec maintains automated systems to identify new potential fraud 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 Security Response reviews events that are identified as attacks for the purposes of confirmation and filter development. The Symantec Brightmail Business Intelligence Department reviews phishing attacks in order to develop predictive filters known as Symantec Brightmail AntiSpam heuristics.

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 fraud 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 as a percent of email scanned

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 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. These sites are categorized according to the brand being phished and its industry sector. The Phish Report Network has senders that send in phishing attacks from many different sources. They include a client detection network that detects phishing Web sites as the clients visit various Web sites on the Internet. There is also server detection from spam emails.

The sender confirms all spoof sites before sending the address of the Web site into the Phish Report Network. After the spoof 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 24x7x365 and manually review all spoof sites sent into the Phish Report Network to eliminate false positives.

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.

Sample set normalization

Due to the numerous variables influencing a company's spam activity, Symantec focused 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 the number of email messages that trigger antispam 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 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. Following this logic, the region from which the spam originates may not correspond with the region in which the spammer is located.

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, or 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:

1. It is considered to be non-viral in nature;
2. It meets criteria for programmatic functionality having potential to affect security;
3. It has been reported to Symantec by a critical number of either corporate or individual users within a given timeframe. The timeframe 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 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" report is based on Symantec's analysis of these reports.

Any technical information that is made available by Symantec Corporation is the copyrighted work of Symantec Corporation and is owned by Symantec Corporation.

NO WARRANTY. The technical information is being delivered to you as is and Symantec Corporation makes no warranty as to its accuracy or use. Any use of the technical documentation or the information contained herein is at the risk of the user. Documentation may include technical or other inaccuracies or typographical errors. Symantec reserves the right to make changes without prior notice.

Copyright © 2006 Symantec Corporation. All rights reserved. Symantec, the Symantec logo, Brightmail, DeepSight, Digital Immune System, and Symantec AntiVirus are trademarks or registered trademarks of Symantec Corporation or its affiliates in the U.S. and other countries. Apple and Mac OS are registered trademarks of Apple Computer, Inc. Safari is a trademark of Apple Computer, Inc. Microsoft, ActiveX, Excel, MSN, Win32, Windows, and Windows Vista are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. Sun, JavaScript, and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc., in the U.S. or other countries. Other names may be trademarks of their respective owners.

About Symantec

Symantec is the world leader in providing solutions to help individuals and enterprises assure the security, availability, and integrity of their information. Headquartered in Cupertino, Calif., Symantec has operations in more than 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

Copyright © 2006 Symantec Corporation. All rights reserved. Any technical information that is made available by Symantec Corporation is the copyrighted work of Symantec Corporation and is owned by Symantec Corporation. 09/06 10756726