

Email Attacks: This Time It's Personal

Executive Summary.....	2
The Business of Cybercrime: The Role of Email.....	2
Reduction in Mass Attacks.....	2
Attack Classifications	3
Mass Attacks	3
Targeted Attacks.....	4
Economics of Attacks.....	5
Impact of Personalized Attacks	6
Impact of Spearphishing Attacks	6
Impact of Targeted Attacks	6
Overall Impact of Attacks	6
Conclusion	7
Solution: Cisco Security Intelligence Operations.....	8

Executive Summary

Cybercriminal business models have recently shifted toward low-volume targeted attacks. With email remaining the primary attack vector, these attacks are increasing in both their frequency and their financial impact on targeted organizations. Cisco Security Intelligence Operations' (SIO) research findings indicate that the annualized cybercrime business activity caused by mass, indiscriminate email attacks has declined by more than half. At the same time, the business activity caused by highly-personalized targeted attacks is growing rapidly, tripling in the last year. While the financial impact translates to monetary loss and stolen credentials, organizations that have been victimized by these attacks have to bear the cost of remediating infected hosts and the negative impact on their brand reputation.

The increasing prevalence of these attacks compounded by trends toward mobility and uncontrolled endpoints, underscores the need for today's organizations to implement a new approach to security that leverages the network. While many organizations train users to identify dangerous messages and avoid clicking on URLs that might lead to compromised websites or malware downloads, user education cannot completely protect organizations from these threats. Instead, organizations need a highly distributed security architecture that manages enforcement elements such as firewalls, web proxies, and intrusion-prevention sensors with a higher-level policy language that is context-aware.

This paper examines attack trends and explores the impact of these campaigns. The findings in this paper are based on research Cisco has conducted with organizations worldwide across a broad range of industries.

The Business of Cybercrime: The Role of Email

The shift in cybercrime business models has resulted in a prominent change in threat activity over the last year. Fewer mass attacks are launched, as evidenced by the 80 percent reduction in overall spam volumes. Instead, cybercriminals are focusing on higher-value endeavors, including increased scams and malicious attacks, spearphishing attacks, and targeted attacks.

Reduction in Mass Attacks

With more cybercriminals moving toward the use of targeted attacks, Cisco SIO estimates that the cybercriminal benefit resulting from traditional mass email-based attacks has declined more than 50 percent: from US\$1.1 billion in June 2010 to \$500 million in June 2011 on an annualized basis. This change reflects a reduction in spam volume from 300 billion to 40 billion spam messages daily from June 2010 to

June 2011. This reduction is consistent with low continued user conversion rates and is partially offset by increases in the average user spending on conversions.

This decline has been offset by a small subset of mass attacks: scams and malicious attacks, which make up about 0.2 percent of total mass attacks and have been providing greater cybercriminal benefit. By using more personalization tools, the user conversion rates for the better-crafted scams and malicious attacks have increased significantly in the last year. In addition, the average user loss caused by the malware or scam employed has increased because of the information shared.

In estimating total losses (see Table 1), Cisco SIO used the conservative estimate of US\$250 per victimized user. This amount is in line with the low-end estimate of recent publicly disclosed scams and malicious attacks. For instance, in June 2011, the U.S. Federal Bureau of Investigation (FBI) announced a scam email directing recipients to send \$350 to obtain a Clearance Certificate or else legal action would be taken against the recipient. Using these estimates, scams and malicious attacks (as a sub-category of mass attacks) have grown from US\$50 million to US\$200 million over the last year on an annualized basis.

Table 1: Cybercriminal Benefit from Mass Attacks

Cybercriminal Benefit (US\$ million)	1 Year Ago	Current
Spam Attacks	\$1,000	\$300
Scams and Malicious Attacks	\$50	\$200
TOTAL	\$1,050	\$500

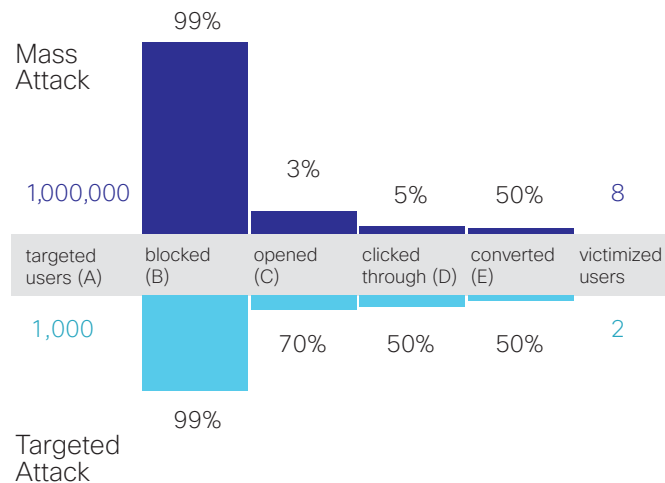
Starting in 2010 and continuing into 2011, the criminal ecosystem has been changing dramatically. Law enforcement authorities and security and industry organizations worldwide have been collaborating to shut down or limit the largest spam-sending botnets and their associates. SpamIt, a large spam-sending affiliate network, ceased operations in October 2010 after its database was leaked and Russian police pressed charges against its owner. Major botnets were severely curtailed or even shut down, including Rustock, Bredolab, and Mega-D. By disrupting the financial and technical business models of key cartels, threat volumes have declined in favor of more lucrative activities.

Let's look briefly at the differences in the conversion process and business models of mass attacks and targeted attacks.

Historically, the spam conversion pipeline started with lists of email addresses used by associated bots to deliver messages (see Stage A in Figure 1). Upon receipt, anti-spam engines correctly identify and block the vast majority of threat messages (Stage B). The messages that make it past the spam filters end up in the user's mailbox as supposedly legitimate mes-

sages. Knowledgeable users often ignore the spam messages and open only a small percentage of them (Stage C). Of these, only a fraction of users will click through (Stage D) and finally be “converted” (Stage E) when the unsuspecting user purchases products or downloads malware.

Figure 1: Threat Conversion Pipeline



This traditional spam pipeline still exists, but it has also evolved with increasing personalization, most acutely in targeted attacks. Targeted attacks typically hold much higher retention throughout the pipeline, as the email and website link are sent to valid users and appear legitimate to security engines and recipients. While the volumes are low, the conversion rates of targeted attacks are much higher. The higher conversion rates come at the cost of higher-value inputs:

- Lists of only valid email addresses with defined attributes
- Legitimate-appearing messages, often purportedly from a known contact with content specific to the recipient(s)
- Higher-quality and typically not-yet-discovered malware
- New websites often created specifically for an individual instance of a targeted attack (and not previously seen)

This is criminal Darwinism at work: Cybercriminals are adapting their campaigns to increase their staying power.

Attack Classifications

As cybercriminal activity continues to evolve, the specific attacks and their impact to organizations also change.

Mass Attacks

Mass attacks have been the basis of threats since the first days of distributed networks. Self-propagating worms, distributed denial of service (DDoS) attacks, and spam are some preferred methods for achieving financial gain or business disruption. The criminal creates a common payload and places it in locations that victims might access, often inadvertently. Examples include infecting websites, exploiting security vulnerabilities in file formats such as PDFs, sending emails to make a purchase, and mass phishing of banking credentials.

Traditional anti-threat methods rely on several factors, including quickly identifying the threat when first reported or seen in the network and then blocking similar threats in the future. If criminals infiltrate the security layers far enough to reach their targets, they'll achieve the desired result in sufficient quantities to make this business model lucrative.

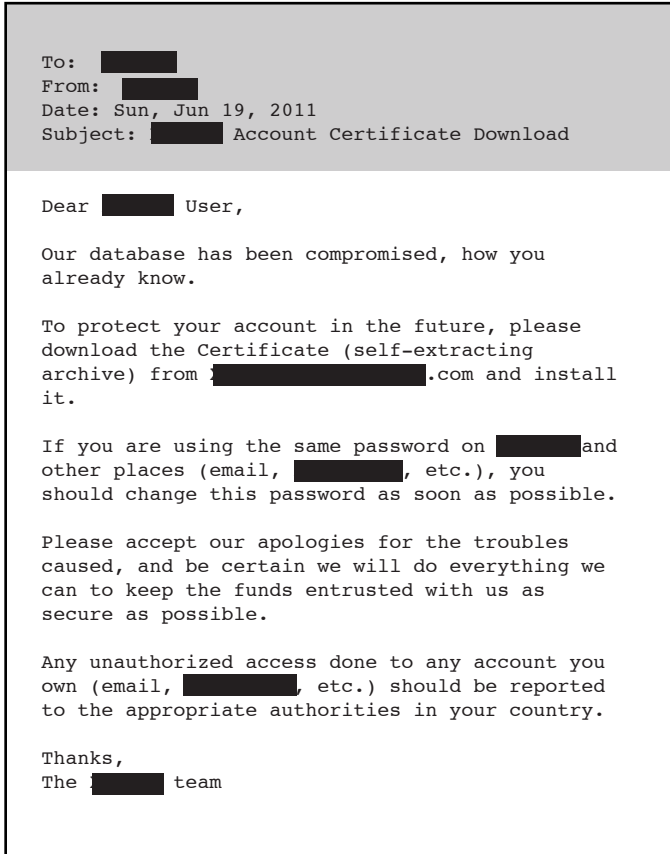
A significant segment of this type of attacks is the burgeoning number of scams and malicious attacks. As part of the evolution of the criminal ecosystem, these attacks are becoming highly focused. Regardless of the vector or delivery engine—including short message service (SMS), email, and social media—criminals are choosing their targets with greater care, using personalized information such as a user's geographical location or job position. Examples of these scams include:

- SMS financial fraud scams to specific locales
- Email campaigns that use URL shortening services
- Social media scams, where the criminal befriends a user or group of users for financial gain

When only a few threats are sent, these strategies may be effective in reaching the victims, but may not always prove cost effective to the criminals. Yet, for reaching high-value victims, this approach is increasingly being leveraged by smart, organized, and profit-driven criminals. When criminals are specific about their victim profiles, these threats are referred to as spearphishing attacks.

Spearphishing attacks are aimed at a specific profile of users, often high-ranking organizational users who have access to commercial bank accounts. Spearphishing attacks are typically well crafted; they use contextual information to make users believe they are interacting with legitimate content. The spearphishing email may appear to relate to some specific item of personal importance or a relevant matter at the company—for instance, discussing payroll discrepancies or a legal matter. According to Cisco SIO research, more than 80 percent of spearphishing attacks contain links to websites with malicious content. Yet, the linked websites are often specially crafted and previously unseen, making them complex to detect.

Figure 2: Spearphishing Message



Targeted Attacks

Targeted attacks are highly customized threats directed at a specific user or group of users typically for intellectual property theft. These attacks are very low in volume and can be disguised by either known entities with unwitting compromised accounts or anonymity in specialized botnet distribution channels. Targeted attacks generally employ some form of malware – and often use zero day exploits – in order to gain initial entry to the system and to harvest desired data over a period of time. With these attacks, criminals often use multiple methods to reach the victim. Targeted attacks are difficult to protect against and have the potential to deliver the most potent negative impact to victims.

While potentially similar in structure, the major differentiator of targeted attacks relative to spearphishing attacks is the focus on the victim. A targeted attack is directed toward a specific user or group of users whereas a spearphishing attack is usually directed toward a group of people with a commonality, such as being customers of the same bank. Targeted attackers often build a dossier of sorts on intended victims – gleaned information from social networks, press releases, and public

company correspondence. While spearphishing attacks may contain some personalized information, a targeted attack may contain a great deal of information which is highly personalized and generally of unique interest to the intended target.

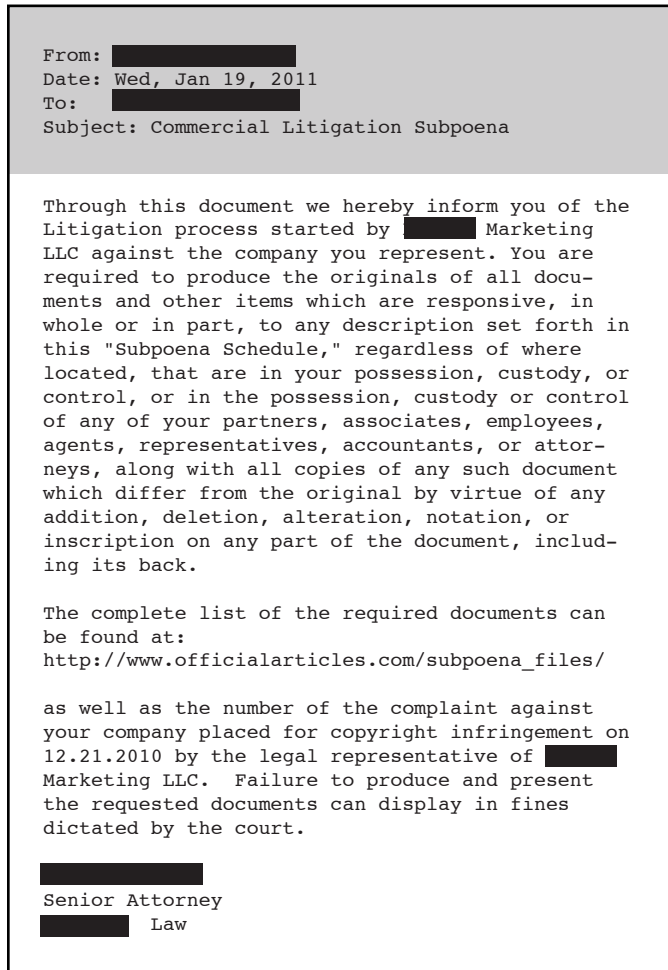
Table 2: Comparison Between Targeted and Spearphishing Attacks

Attributes	Targeted Attacks	Spearphishing Attacks
Intent	Intellectual Property Theft	Financial Gain
Malware	Yes, often with zero-day exploits	Possibly
Target Reconnaissance	Yes	No
Level of Personalization	Very High	Some

A well-publicized example of a targeted attack is the Stuxnet attack, a computer worm discovered in July 2010 which specifically targeted industrial software and equipment. Stuxnet exploited a vulnerability in the way that Windows handles shortcut files, allowing the worm to spread to new systems. The worm is believed to be purpose-built to attack Supervisory Control and Data Acquisition (SCADA) systems, or those used to manage complex industrial networks, such as systems at power plants and chemical manufacturing facilities. Stuxnet’s cleverness is in its ability to traverse non-networked systems, which means that even systems unconnected to networks or the Internet are at risk. Operators believed that a default Siemens password (which had been made public on the web some years earlier) could not be corrected by vendors without causing significant difficulty for customers. The SCADA system operators might have been laboring under a false sense of security—since their systems were not connected to the public Internet, they might have believed they would not be prone to infection. Federal News Radio’s website called Stuxnet “the smartest malware ever.”

In January 2011, Cisco SIO detected a targeted attack message sent to senior executives at a large corporation. This campaign was sophisticated, in that it used previously unseen resources. The message was sent by an unknown party through a legitimate but compromised server in Australia. The email message was seemingly legitimate (figure 3). The embedded action URL was hosted on a legitimate but compromised law blog. When clicked, the user’s browser was directed to a previously unknown copy of the Phoenix exploit kit. After the exploit was successful, it installed the Zeus Trojan on the victim’s computer.

Figure 3: Targeted Attack Message



Economics of Attacks

The economics of a typical campaign underscore the difference between mass and targeted attack business models. As a proxy, Table 3 compares the yield in the conversion pipeline and the relative economics to the cybercriminal for a sample mass phishing versus spearphishing attack:

Table 3. Economics of Mass Phishing vs. Spearphishing Attacks

Example of a Typical Campaign	Mass Phishing Attack (Single Campaign)	Spearphishing Attack (Single Campaign)
(A) Total Messages Sent in Campaign	1,000,000	1,000
(B) Block Rate	99%	99%
(C) Open Rate	3%	70%
(D) Click Through Rate	5%	50%
(E) Conversion Rate	50%	50%
Victims	8	2
Value per Victim	\$2,000	\$80,000
Total Value from Campaign	\$16,000	\$160,000
Total Cost for Campaign	\$2,000	\$10,000
Total Profit from Campaign	\$14,000	\$150,000

For an individual campaign, the economics of a spearphishing attack can be more compelling than for a mass attack. The costs are significantly higher, but so too are the yield and benefit. Cisco SIO estimates the costs of a spearphishing attack at five times the cost of a mass attack, given the quality of the list acquisition, botnet leased, email generation tools, malware purchased, website created, campaign administration tools, order processing back-end infrastructure, fulfillment providers, and user background research activity required. This significantly higher cost basis and greater effort requires highly specialized skills. It also requires higher yields to be effective.

Cybercriminals are balancing competing priorities: Infect more users or keep the attack small enough to fly under security vendors' radar? Spearphishing attack campaigns are limited in volume but offer higher user open and click-through rates. With these constraints, cybercriminals are increasingly focusing on business users with access to corporate banking accounts, to make sure they're seeing sufficient return per infection. This is why the average value per victim can be 40 times that of a mass attack. Ultimately, this approach is justified: Profit from a single spearphishing attack campaign can be more than 10 times that of a mass attack.

The potential returns are causing a shift in cybercriminal business models. Presently, the opportunity cost of spamming may not be worth the rate of return due to increases in both anti-spam efficacy and user awareness. Instead, cybercriminals are focusing more time and effort on different types of targeted attacks, often with the goal of gaining access to more lucrative corporate and personal bank accounts and valuable intellectual property.

To make their attacks more personalized, some cybercriminals have focused on infiltrating email marketing vendors, since they have valid names, email addresses, and other attributes. When used in scams and malicious attacks—whether on a mass scale or in spearphishing attacks—this personal information increases the likelihood of users opening an attack email.

The correlation of lower mass spam with recent data breaches is interesting, but the real takeaway is that attacks are becoming more personalized.

Impact of Personalized Attacks

Impact of Spearphishing Attacks

Spearphishing attacks, though lower in volume relative to other types of threats, have serious consequences for today's enterprises. The majority of spearphishing attacks ultimately lead to financial loss, making them incredibly dangerous to victims and incredibly valuable to cybercriminals.

Spearphishing uses customization methods superior than those used in mass scams and malicious attacks, resulting in significantly higher user open and conversion rates. These success factors have made spearphishing attack infections more effective, and hence more commonplace, which is corroborated by Federal Trade Commission estimates of 9 million Americans having their identities stolen each year.

The value per victim in spearphishing attacks can vary substantially, with the mean and median values being quite high. For example, according to primary consumer research conducted by Javelin Strategy & Research, the mean identity fraud amount per victim was \$4,607 in 2010. If we use a conservative estimate of user loss—\$400—the total cybercriminal benefit resulting from spearphishing attacks amounts to \$150 million in June 2010 on an annualized basis (see Table 4). This figure has tripled from \$50 million a year ago; it is expected to continue increasing in the coming months as cybercriminal activity returns to its prior business levels.

Impact of Targeted Attacks

The malicious nature of targeted attacks causes them to be very expensive to society in general and to individual organizations specifically. The cybercriminal benefit from a targeted attack, while substantial, is not easy to estimate because it is highly variable, based on the specific victim and

intellectual property compromised. However, the cybercriminal benefit is a subset of the overall cost to the victim organization, which also depends heavily on the organization's reputation and status.

The organizational costs resulting from targeted attacks can vary. According to the FBI, these costs can range from thousands to hundreds of millions USD. Similarly, the Ponemon Institute has estimated the potential cost per organizational data breach to range anywhere from US\$1 million to US\$58 million. As an example, a large gaming platform provider reported that the unauthorized access to its network that occurred in Q2 of 2011 has resulted in currently known associated costs of approximately US\$172 million. Costs include personal information theft protection programs, insurance to cover identity theft losses, costs of "welcome back" programs, customer support costs, network security enhancement costs, legal and expert costs, and the impact on profits due to possible future revenue decreases.

In another example, a public payments processor company experienced a data breach resulting in millions of compromised user account credentials. A year later, the company reported related expenses totaling US\$105 million. As per their 10-Q SEC filing, "The majority of these charges, or approximately \$90.8 million, related to: (i) assessments imposed by MasterCard and VISA against us and our sponsor banks, (ii) settlement offers we made to certain card brands in an attempt to resolve certain of the claims asserted against our sponsor banks (who have asserted rights to indemnification from us pursuant to our agreements with them), and (iii) expected costs of settling with certain claimants with whom settlement discussions are underway." During the same timeframe from the intrusion to the 10-Q results, the company lost 30% of its value relative to the Standard and Poor's 500 Index, or roughly \$300 million in shareholder value.

Ultimately, the corporate reputation is tarnished at a cost more significant than the costs of the monetary loss and remediation combined.

Overall Impact of Attacks

Table 4 aggregates these estimates and shows the the annual total monetary benefit to cybercriminals for different types of attacks.

Table 4: Total Annual Cybercriminal Monetary Benefit

Cybercriminal Benefit (US\$ million)	1 Year Ago	Current
Mass Attacks	\$1,050	\$500
Spearphishing Attacks	\$50	\$150
Targeted Attacks	Varies, see above	Varies, see above
TOTAL	\$1,100	\$650

It's clear that the shift in cybercriminal business models has provided an interim benefit from lower threat activity. Organizations are only partially able to appreciate the reduction in cybercriminal activity, though, as their costs can encompass far more than financial loss. To estimate these total losses, Cisco SIO conducted primary research with 361 organizations located globally to understand their perspectives.

The organizational impacts of attacks can be categorized as follows:

1. Financial
2. Remediation
3. Reputation



Financial: Financial loss directly to the cybercriminals can range widely based on the specific attack; as a result, organizations cannot estimate the loss.

Remediation: The remediation costs of spearphishing and targeted attacks are incurred by victim organizations. The administrative team must identify and remediate the compromised hosts; this can be challenging given the increasing use of surreptitious applications. Because of the complexity of current targeted attacks and the underlying malware, costs for remediation can be significant.

Remediation costs include the time required to address the infected host and the corresponding opportunity cost of that time. With the organizations surveyed, Cisco observed that infected hosts take an average of two hours of dedicated effort to resolve. The cost basis of two hours of effort per resolution is specific to each organization, as is the corresponding opportunity cost of that time.

Based on Cisco SIO research, organizations estimated that the direct remediation cost per infected user is \$640, or 2.1 times that of the direct monetary loss.

Reputation: The negative reputation impact of attacks can be experienced over time by victim organizations and users. For example, building a brand typically takes years, but a negative event or news story, especially one that is highly visible, can quickly tarnish a company's image. The direct impact can be a significant decline in business, sometimes even leading to the organization's demise.

Determining the true costs of adverse reputation impact can be challenging, as is estimating the value of an organization's brand. Nevertheless, organizations have made it clear that adverse events can impact their reputation, which in turn can create a significant decline in business and shareholder value.

Based on Cisco SIO research, organizations estimated that the reputation cost per infected user is \$1,900, or 6.4 times that of the direct monetary loss.

Combined Impact: The overall costs of spearphishing and targeted attacks to organizations are substantially more than their direct monetary loss to cybercriminals. Table 5 provides results from the 361 organizations Cisco SIO researched.

Table 5: Overall Organizational Costs per Attack

Size of Organization	Monetary Loss*	Remediation Cost*	Reputation Cost*
Up to 1,000 users	\$327	\$558	\$2,346
Between 1,000 and 5,000 users	\$233	\$484	\$1,436
More than 5,000 users	\$290	\$833	\$1,553

*Per Infected User

While the costs can vary widely depending on the specific organization and attack, one point is clear: The overall costs to organizations can be significant. In addition, reputation management and remediation efforts can create a strain on the organization.

Conclusion

The increased number of low-volume targeted attacks has impacted users in many organizations, regardless of industry, geography and size. Their prevalence has caused both a related increase in criminal financial benefit and impact on victimized organizations. Organizations have to bear the burden of not only the monetary loss but also the cost of remediating infected hosts and the negative impact on their brand reputation. With the number of targeted attacks expected to increase, cybercriminal activity will continue to evolve, as will its impact.

Solution: Cisco Security Intelligence Operations

Traditional security, which relies on layering of products and the use of multiple filters, is not enough to defend against the latest generation of malware, which spreads quickly, has global targets, and uses multiple vectors to propagate.

Cisco stays ahead of the latest threats using real-time threat intelligence from Cisco Security Intelligence Operations (SIO), the world's largest cloud-based security ecosystem. Cisco SIO uses SensorBase data from almost one million live data feeds from deployed Cisco email, web, firewall, and intrusion prevention solutions.

Cisco SIO weighs and processes the data, automatically categorizing threats and creating rules using more than 200 parameters. Security researchers also collect and supply

information about security events that have the potential for widespread impact on networks, applications, and devices. Rules are dynamically delivered to deployed Cisco security devices every three to five minutes. The Cisco SIO team also publishes security best practice recommendations and tactical guidance for thwarting threats.

Cisco is committed to providing complete security solutions that are integrated, timely, comprehensive, and effective—enabling holistic security for organizations worldwide. With Cisco, organizations can save time researching threats and vulnerabilities, and focus more on taking a proactive approach to security.




Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

 Cisco and the Cisco Logo are trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and other countries. A listing of Cisco's trademarks can be found at www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1005R)