Targeted Phishing

Email is the medium most organizations have come to rely on for communication. Unfortunately, most incoming email is unwanted - or even malicious. Today's modern spam-blocking appliances have little problem weeding out the vast majority of unsophisticated spam campaigns, leaving end-user inboxes filled with only legitimate email. That's in spite of the fact that more than 85 percent of incoming mail consists of spam or "abusive messages," according to the Messaging Anti-Abuse Working Group.

To get around advanced antispam technology, online criminals are becoming more dangerous and sophisticated. In addition to enticing a spam recipient to buy a dubious product, more lucrative "phishing" attacks seek to glean users' personal information, such as names and addresses - and even login information for their banks. Although the number of such phishing emails being sent is still relatively low, it is increasing, and the danger for intended victims is high. As Internet users become more adept at detecting clumsy attempts to phish personal information, spammers are selectively phishing smaller and smaller demographics with content that appeals specifically to each group. This form of highly targeted, socially engineered email is called "targeted phishing" or "spear phishing," and can fool even the savviest of Internet users.

Trends and Solutions

Since the late 1990s, "phishing" emails (messages designed to fool the recipient into handing over personal information, such as login names and passwords) have been flooding email inboxes. The "phishers" - the online criminals who create emails that mimic messages from well-known online services or legitimate companies - typically send out millions of emails at a time, in hopes of stealing the online banking or other login names and passwords of even just a few recipients.

The Growth and Payoff of Targeted Phishing

A growing percentage of email-borne attacks are targeted phishing attacks, where a specific organization or group of individuals is singled out. The targets receive cleverly crafted phishing messages that are designed to solicit a deeper level of personal data, such as login and password information that could grant access to corporate networks or databases filled with sensitive information. In addition to soliciting login information, targeted phishing emails can also deliver malware - for instance, keystroke logging programs to track everything the victim types.

Targeted phishing costs online criminals more time and money than traditional phishing campaigns. The scammers need to rent or steal lists of valid email addresses for a target organization or group, and then create plausible emails that are likely to lure their recipients into supplying personal data. However, when targeted phishing succeeds, it has the potential for a bigger payoff, making the investment worthwhile.

Currently, targeted phishing messages represent about 1 percent of all phishing campaigns. However, since targeted phishing is often aimed at just a few well-placed individuals in an organization, it can potentially do a great deal of damage - from financial, data security, and customer relations standpoints. Additionally, the personalized approach of targeted phishing makes it more difficult to weed out these emails via standard antiphishing technologies, leaving organizations vulnerable.

Figure 1. Traditional spam campaigns are sent in high volumes with low expected click-through and sales conversion rates. New targeted attacks are more dangerous in nature and are relying on low volume to get through traditional spam filters.

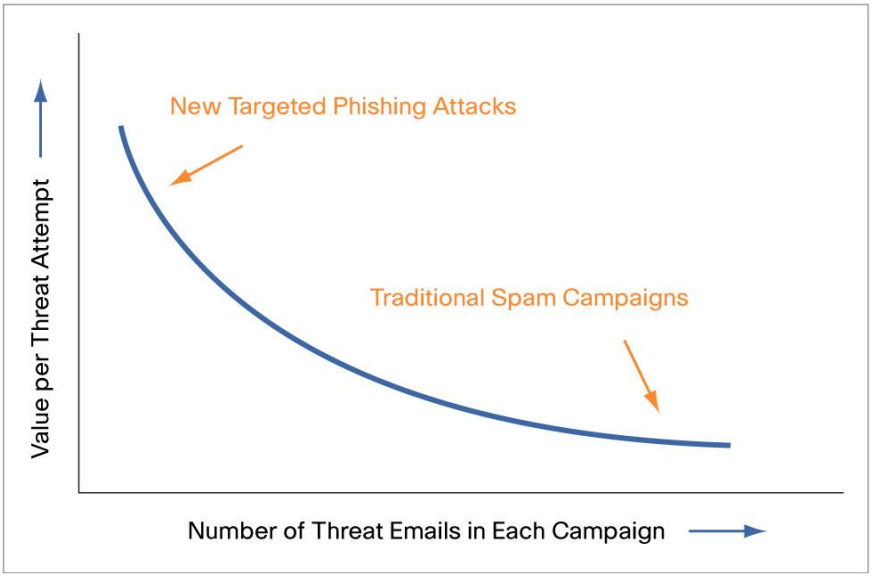

Why Targeted Phishing Works

Techniques for getting victims to click through to websites - where they either unwittingly submit sensitive information to scammers or download malware on their computers - are becoming increasingly sophisticated. Most spam now includes URLs directing recipients to malicious websites. These days, the fraudulent websites that victims are directed to often look and feel extremely similar to legitimate sites.

According to a UC Berkeley study, even longtime, frequent Internet users are sometimes fooled by malicious websites. To avoid being taken in by phishing websites, users had to use a strategy of consistently checking the content's apparent level of legitimacy, the address bar and its security settings, the padlock images in the browser frame, and the security certificate of any website they were directed to.

Phishing emails aimed at broad distribution lists today depend on social engineering techniques, such as content that demands an action from the recipient and referrals to legitimate-looking websites (like fraudulent online banking sites). But these types of emails rarely use any personal data within the message.

Meanwhile, targeted phishing emails take social engineering to a new level. By addressing a recipient by name and sending the message directly to his or her email address, scammers ratchet up the credibility of the malicious email and the fake websites to which the victim is directed.

In the example below, business executives received a phishing email purporting to be from the Internal Revenue Service, which claimed that a criminal tax fraud investigation into their company was underway. The email was sent to a specific person, and cited the company name in the body of the message.

A URL in the email launched an executable file for a Trojan that would steal all interactive data sent from the recipient's email browser and would access form data before it was SSL-encrypted. Another targeted phishing email to executives mimicked messages from a U.S. district court, ostensibly subpoenaing the recipient for appearance in a civil court case.

Figure 2. Targeted phishing attacks require criminals to efficiently build appropriate resources and trick victims into revealing valuable private information.

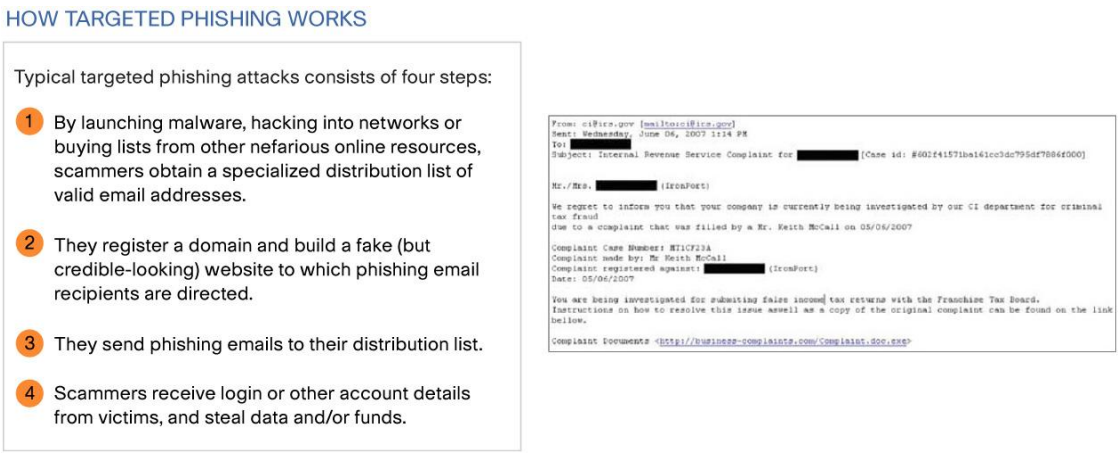

Social Engineering for Success

Targeted phishing attacks are not aimed at business executives only. Many recent campaigns involved emails, supposedly from local banks to their customers, asking customers to renew their online accounts or log in to their accounts to read special messages. Other targeted attacks, seemingly from university IT departments, directed email users to reply with their webmail credentials in order to retain their university email account or take advantage of a security upgrade. These compromised accounts are then often used to send large-scale spam campaigns.

Figure 3. Examples of targeted phishing messages, purportedly from university IT departments, designed to harvest email credentials.

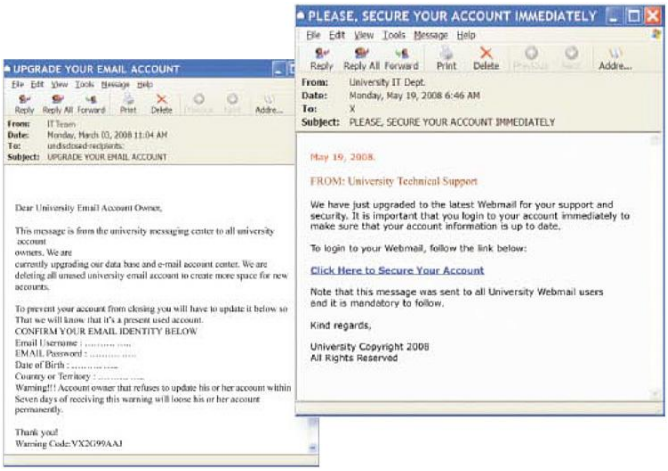

The scammers sending out targeted phishing campaigns continue to refine their tactics for luring victims to fraudulent or compromised websites. Online criminals have even been known to send text messages to mobile phones. One such attack targeted mobile numbers in the same area as a local bank, informed customers that their accounts had been closed due to suspicious activity, and then directed them to call a phone number to re-activate the account. The call-in number was set up by the scammers to collect account numbers and login credentials.

Figure 4. Criminals have also set up automated systems to collect banking login information from unsuspecting customers.

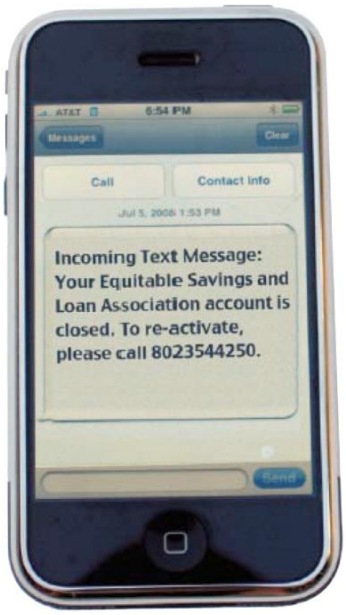

No matter how targeted phishing attacks are perpetrated, their goal is to glean personal data that allows online criminals to steal money or information. In 2007, senders of a successful targeted phishing campaign came close to scamming a major grocery store chain, Supervalu Inc., out of US$10 million.

In the Supervalu case, online scammers fraudulently obtained wire-transfer instructions from American Greetings and Frito-Lay, two of the grocery chain's suppliers. The criminals (masquerading as employees of the suppliers) then sent emails to individual Supervalu employees with "updated" wire-transfer instructions, directing them to transfer $10 million over several days to the criminals' bank account.

Fortunately for Supervalu, alert employees at American Greetings and Frito-Lay realized they did not receive payment and contacted Supervalu. Together, they were able to notify law enforcement authorities and have the fraudulent accounts frozen.

Other recent phishing campaigns also demonstrate the threats associated with these messages, and the fact that criminals are seeking to steal more than just banking information. One campaign involved emails sent to American Airlines frequent flyers, offering them $50 to fill out an online survey - which, unbeknownst to the recipients, was designed to obtain personal information. American Airlines had to email all its customers, warning them about the phishing email. Another campaign asked Google AdWords customers to confirm their accounts. When they did so, not only was their Google AdWords account information (including financial data) stolen by the scammers, but the victims' AdWords traffic was also redirected to spammers' websites.

Figure 5. Messages that appear to be related to their Google AdWords accounts trick victims into providing login information.

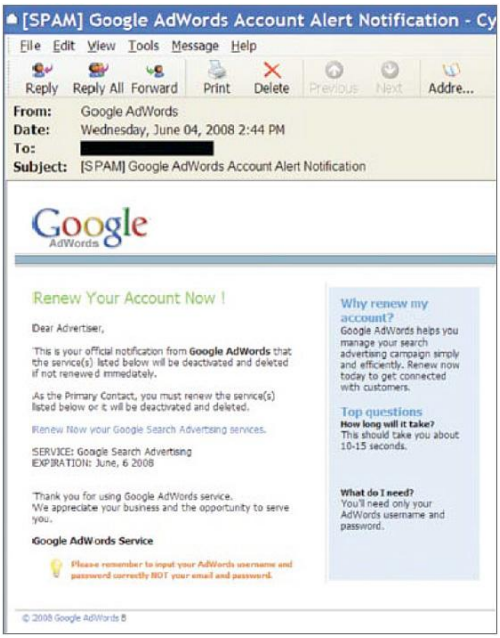

Even when recipients aren't fooled into responding to phishing emails, targeted phishing attacks adversely affect companies - and their relationships with their customers.

According to Forrester Research, executives who receive targeted phishing messages are losing confidence in email. In its "Phishing Concerns Impact Consumer Online Financial Behavior" report, Forrester notes that 26 percent of U.S. consumers will not use online financial products, and 20 percent of consumers won't open emails from their financial provider or enroll in online banking or bill payment, all because they fear falling victim to a phishing attack.

How IronPort Thwarts Targeted Phishing

Now part of Cisco, IronPort provides a sophisticated array of ever-expanding technologies to ensure that end users don't have to make decisions about whether the email they are viewing is an illicit bid for account information. With IronPort® solutions, customers can safely use email and the web and are protected from new types of attacks such as targeted phishing. IronPort enables a multilayered approach that involves monitoring worldwide email and web traffic and using sophisticated web reputation filters and advanced email authentication technologies.

SenderBase: IronPort's SenderBase® network constantly monitors more than 30 percent of worldwide email and web traffic. Using IP address information, SenderBase tracks more than 150 parameters, such as email sending and website traffic volume, complaint levels, "spamtrap" accounts, DNS resolution, country of origin, and blacklist presence. The system then uses the data collected to create a Reputation Score to indicate the threat level for every email message coming into an organization, as well as the URLs contained in each email. Because 90 percent of malicious emails contain URLs, SenderBase's unique capability of monitoring both web and email traffic is a key component in IronPort's ability to effectively identify and block targeted phishing attacks.

Figure 6. Over 100,000 organizations participate in SenderBase, enabling the world's largest email and web traffic monitoring system.

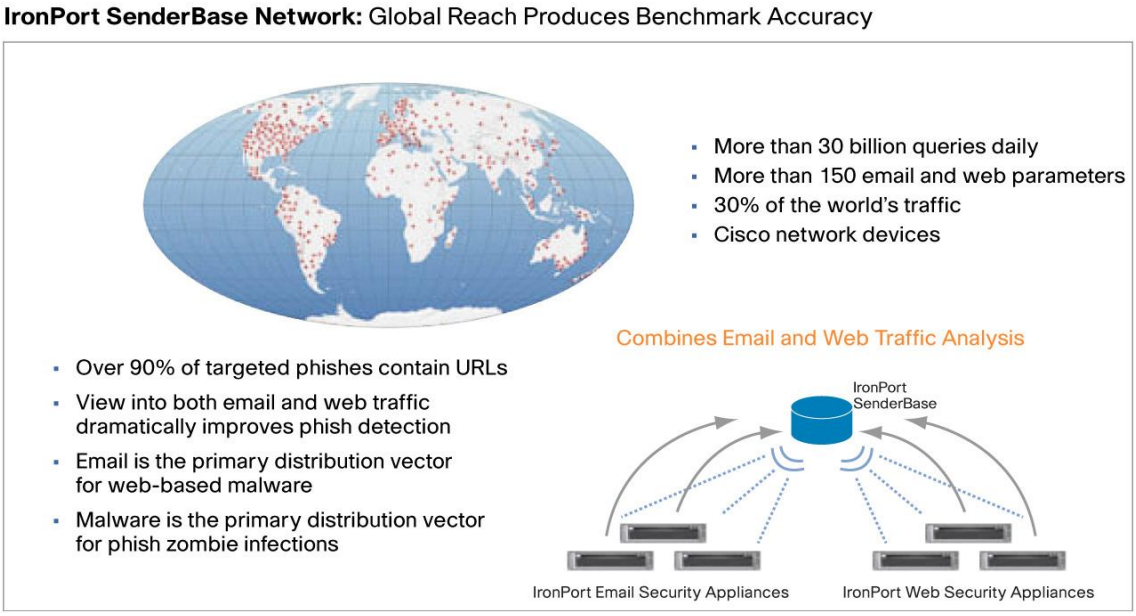

IronPort Web Reputation Filters: IronPort Web Reputation Filters assign web reputation scores to URLs in all emails, based on each URL's likelihood to host malicious content. Based on these reputation scores, IronPort email and web security appliances then allow, flag, or block emails from certain senders and traffic with certain websites.

Reputation scores are based on SenderBase data and additional analysis of hard-to-spoof IP address data, such as how long a domain name has been registered, in which country the site is hosted and/or how frequently that changes, and whether a domain purporting to be hosted by a Fortune 500 company is in fact hosted by that company.

Cisco IronPort's SenderBase Network and Reputation Filters, which together block 99 percent of phishing emails, primarily depend on in-depth insight into IP addresses and activities around them.

SPF and DKIM email verification: Advanced email authentication techniques enable the matching of purported and actual sender identities. Sender Policy Framework (SPF) and DomainKeys Identified Mail (DKIM) are two popular, complementary email authentication methods that can detect fraudulent emails. These authentication technologies are currently gaining broader acceptance among industry groups, email providers, and enterprises, and their wider use enhances their efficacy in combating phishing email. SPF and DKIM each address specific issues and can be used in combination to provide a layered approach to security.

SPF is a form of sender path authentication that helps recipients identify the authorized mail servers for a particular domain, and validate that emails they received did in fact originate from these authorized sources. Mail senders (such as ISPs and corporations) that use this technology publish SPF records that specify which hosts are permitted to use their names. SPF-compliant mail receivers use the published SPF records to test the authorization of the sending Mail Transfer Agent's identity during an email transaction.	**DKIM** is a cryptographically based authentication method that helps verify and determine the authorization of email from a given domain. DKIM provides a "cryptographic signature" (or "key") of multiple email header fields and the body of a message. In its DNS record, a web domain protected by DKIM publishes the public key (or "domain key") that corresponds to its self-generated private signing key. Email recipients can use that key to verify that the message header and body match the identity of the sending domain, helping them determine whether the email is likely to be a phishing or other malicious message.

SPF and DKIM, especially used together, can be very effective at detecting targeted phishing messages. According to the Authentication and Online Trust Alliance (AOTA), about half of all legitimate email worldwide is currently authenticated. This level of adoption means that high-profile executives, who receive extremely high levels of spam and phishing emails, may benefit from additional email filtering and blocking tools based on email authentication failures.

HTML sanitization: HTML sanitization (also known as HTML-Convert) offers additional protection for emails that meet predetermined criteria, such as when SPF and DKIM are not able to authenticate a message. When HTML sanitization is enabled, URLs are made non-clickable and converted to plain text, exposing hidden, potentially malicious content to the recipient. This does add some burden for the recipient when they wish to visit a legitimate URL cited in a message, as they have to copy and paste the plaintext link into their browser to visit the website. But for those individuals targeted by scammers, it offers an excellent layer of additional protection because it enables them see which website they are attempting to visit.

The IronPort S-Series Web Security Appliance: For organizations looking for in-depth defense against targeted phishing and other malware attacks, the IronPort S-Series provides an integrated, layered, and easy-to-manage platform for web security. It addresses the entire spectrum of web traffic, protecting against both known and unknown sites through the use of powerful reputation filters and antimalware defense technologies.

Summary

A new threat in phishing is becoming a dangerous problem: targeted phishing. These messages use advanced social engineering techniques - such as addressing recipients by name (and identifying their companies) - to convince carefully selected victims to unwittingly pass sensitive data or money to online criminals.

As targeted phishing emails take more resources to set up, they currently comprise only a small portion of phishing email worldwide. However, their payoff can be enormous, which means their number will undoubtedly increase.

To help organizations avoid falling victim to targeted phishing campaigns and other malicious attacks, Cisco IronPort offers a layered, integrated approach for email and web security - combining Internet traffic monitoring, reputation filters, and authentication technologies.

Contact

Cisco IronPort sales representatives, channel partners, and support engineers are ready to help you evaluate how IronPort products can make your email infrastructure secure, reliable, and easier to manage. If you believe that your organization could benefit from Cisco IronPort's industry-leading products, please call 650-989-6530 or visit us on the web at http://www.ironport.com/try.

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.

Printed in USA C11-701843-00 02/12