Mitigation Monday

Defense Against Malicious E-mail Attachments, Release 1

March 2008 – Vulnerability Analysis and Operations Group, National Security Agency

Comments or feedback? ***mitigationmonday@nsa.smil.mil***

The networks that make up the critical infrastructure of the United States are under constant attack from sophisticated hacker groups in search of information. Reports to Congress state that foreign countries are interested in acquiring our intellectual capital to use for their own purposes. The DNI's *Annual Threat Assessment of the Intelligence Community* [February 2008] states that "Nation states and criminals target our government and private sector information networks to gain competitive advantage in the commercial sector." US Air Force Maj. Gen. William Lord, at the 2006 IT Conference in Montgomery, AL, said that China had already downloaded 10-20 terabytes of data from the DoD's NIPRNet.

Every network that contains important defense, political, or economic information is a target. These networks contain information representing billions of dollars in research and years of technical advantage. When military secrets are stolen, our troops face harm from adversaries who know too much. When economic data is stolen, our adversaries gain advantage over us in trade negotiations. When engineering designs are stolen, our adversaries can replicate or improve them—a resolute adversary could even modify our designs to make them malfunction at key moments.

Many CIOs are unaware of how vulnerable their networks are and how aggressively they are being targeted. This mitigation report presents a common attack scenario for Microsoft Windows networks and discusses how it can be prevented using a defense-in-depth strategy.

SNAC DoD 9800 Savage Rd Ft.Meade, MD 20755-6704 410-854-6632 DSN 244-6632 FAX: 410-854-6604 **www.nsa.gov/snac**

Scenario

A government official, we'll call him Bill, receives an innocent-looking e-mail from a colleague. He clicks on the e-mail and reads it. The message informs Bill of a meeting that will take place next month; it also contains an attached document with the meeting's agenda. "Great, another meeting," Bill thinks as he sips his coffee and peruses the agenda. Bill then closes the document, and proceeds to the next e-mail in his in-box. "Only 21 left," he sighs.

During his lunch break, Bill runs into the e-mail's author in the cafeteria and asks about the location of the meeting. Bill's colleague looks at him blankly and asks, "What meeting?"

Bill is actually luckier than most military officers, CEOs, and intelligence officials that are targeted daily, because his encounter with the e-mail's supposed author notified him that the e-mail was not authentic.

In reality, an attacker carefully forged an e-mail targeting Bill. E-mail is the most common vector for infiltrating systems. While the specifics vary, the general theme remains the same: A user reads or executes a file attached to an e-mail, and that file is designed by an adversary to exploit a vulnerability in the target system to gain control. When Bill opened the fake agenda attached to the e-mail, malicious code inside it took control of his computer by exploiting a "buffer overflow" vulnerability in his word processing program. The malicious code then connected to a server on the Internet and downloaded and installed a "backdoor" program. The backdoor created a persistent foothold on Bill's system and provided the adversary with access to it in the future.

One evening later that week, the adversary connected to Bill's compromised machine through the backdoor. The adversary rummaged through Bill's recently accessed files, stealing all new documents, presentations, spreadsheets, mail messages, and drawings of strategic or economic interest. The adversary closed the session, but left the backdoor installed, running surreptitiously in the background. "Just in case there is something of value we left on Bill's machine, or if we ever want to launch a Denial of Service attack against Bill's organization," the adversary mused.

SNAC DoD 9800 Savage Rd Ft.Meade, MD 20755-6704 410-854-6632 DSN 244-6632 FAX: 410-854-6604 **www.nsa.gov/snac**

Mitigation

In the scenario, the attack was a targeted, forged e-mail with a malicious attachment. The following diagram shows the steps that allowed the adversary to steal information from Bill, along with common, scalable defenses that could have prevented the attack and protected Bill's information. In the diagram, defenses shown with larger fonts in boxes farther to the left have greater Benefit vs. Cost. Note that some very effective defenses are not farther to the left because of their high cost to deploy and maintain.

Many networks implement some of these defenses. However, it is best to implement as many as possible: A "defense-in-depth" strategy must be used to defend against all stages of an attack.

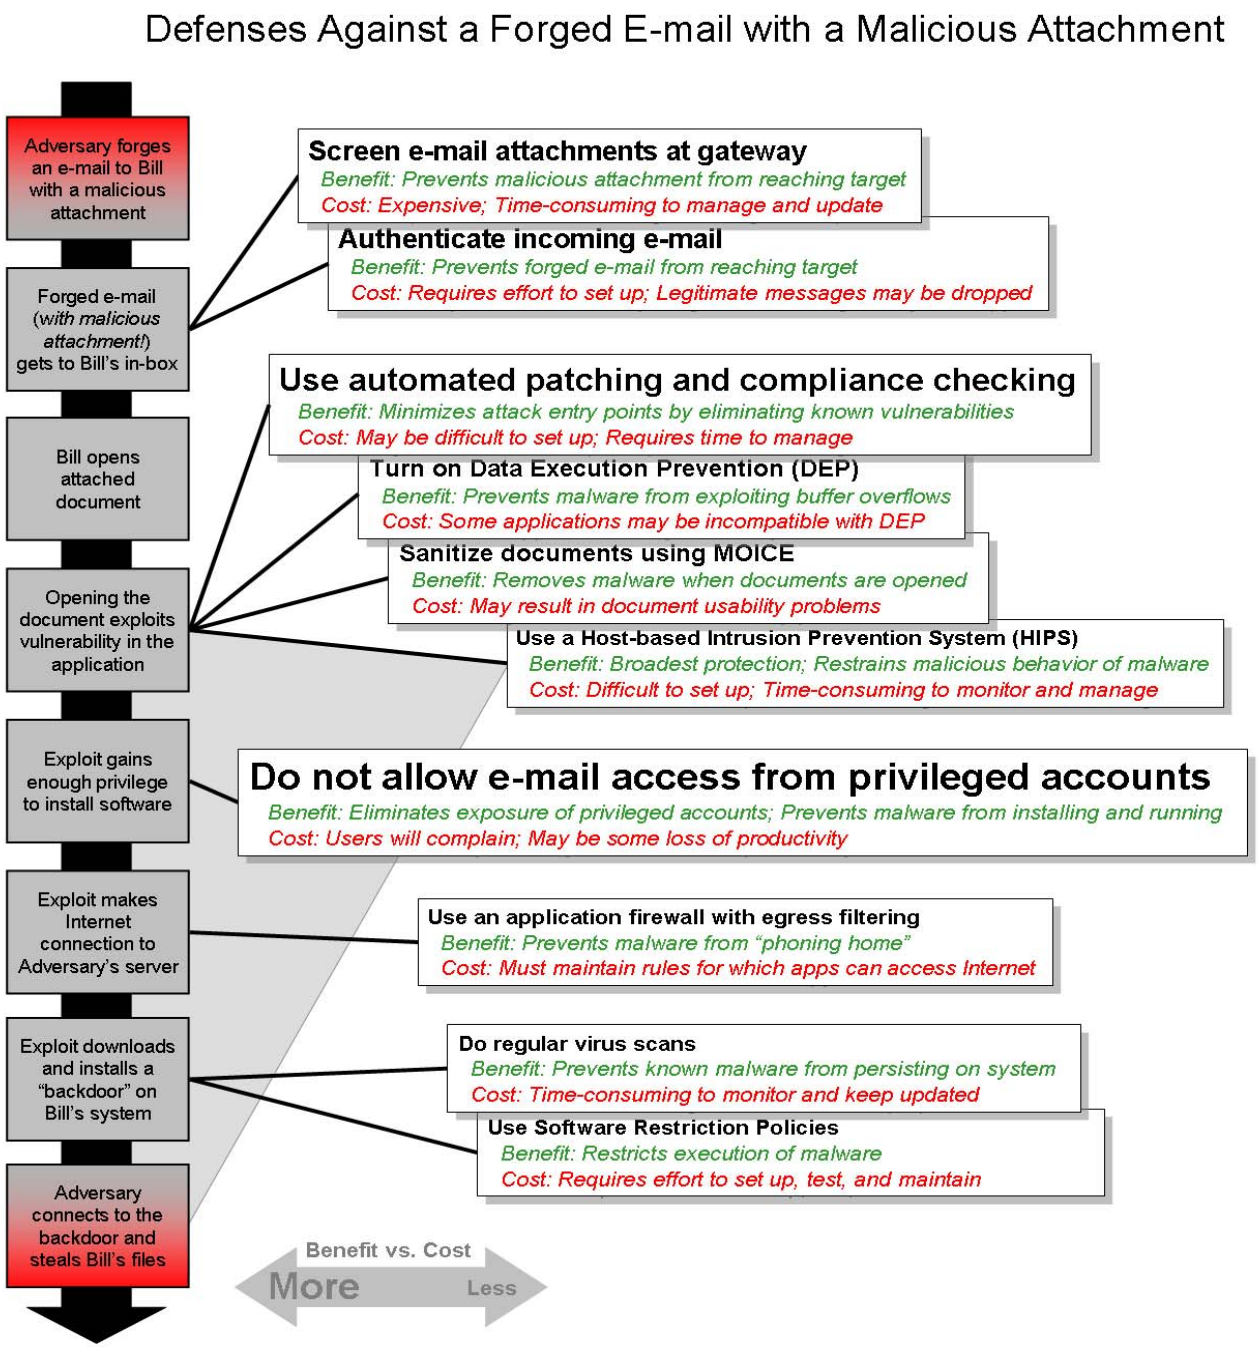

SNAC DoD 9800 Savage Rd Ft.Meade, MD 20755-6704 410-854-6632 DSN 244-6632 FAX: 410-854-6604 **www.nsa.gov/snac**

Defense Details

1. Do not allow e-mail access from privileged accounts

In order to install the "backdoor" on Bill's system, the adversary took advantage of the fact that Bill was reading his e-mail while logged in as a local administrator on his machine. If Bill had been logged in as a non-privileged user instead, the malicious e-mail attachment would not have been able to install the backdoor as part of the operating system on Bill's computer.

When an attack is successful, the adversary gains the privileges of the account from which the attack is run. Attackers prefer highly-privileged administrator accounts that allow them to install and run their malware. The goal of this mitigation is to confine the adversary to exploitation of less-privileged user accounts, so that when an attack does succeed, it does not have enough privilege to install or run code.

A vital first step in mitigating e-mail attacks is to stop users from reading e-mail and surfing the Web using privileged accounts. This includes all accounts that have local administrative privileges, that is, those in the local Administrators group, the Power Users group, and any domain privileged accounts.

Benefit: *Eliminates the exposure of privileged accounts and denies adversaries the ability to install and run malware.* An adversary will not have full administrative access to resources if he is only able to compromise a user's account.

Cost: There is no actual cost in prohibiting users from accessing e-mail and the Internet from privileged accounts. If users must be able to install software or devices, they should use an account specifically created for that purpose. Software developers should not use their privileged access accounts to read e-mail or surf the Web. Requiring users to switch accounts to access e-mail and the Web will result in complaints and may result in some loss of productivity.

2. Use automated patching and compliance-checking

In order to gain control of Bill's machine, the adversary took advantage of a publicly known flaw in Bill's document-reading software. Keeping document readers and other applications that take input from external sources fully patched is an essential step toward defending the network.

Sites should implement an automated patching mechanism, such as Microsoft's Windows Server Update Services (WSUS), and develop a separate strategy for patching applications that are not handled by an automated system. It is also vital to perform compliance checking to determine if the automated patching system is working. Besides getting reports from the automated patching system, use additional auditing tools to ensure that patches are being deployed successfully.

For more information on WSUS, go to http://technet.microsoft.com/en-us/wsus.

Benefit: *Eliminates known vulnerabilities to minimize attack entry points.* Patching can greatly reduce the attack surface that adversaries can use to enter a system. Also, from a manpower perspective, automated patching frees site Administrators from the labor-intensive manual process.

SNAC DoD 9800 Savage Rd Ft.Meade, MD 20755-6704 410-854-6632 DSN 244-6632 FAX: 410-854-6604 **www.nsa.gov/snac**

Cost: WSUS is available for free from Microsoft and is simple to set up and manage. Other software deployment tools are available to help patch applications that can't be handled by WSUS. It is inconvenient to patch applications that do not have an automated mechanism for updating, so consider limiting the number of different applications in the enterprise to reduce the patching cost.

3. Screen e-mail attachments at the gateway

The malicious document may not have gotten to Bill's in-box if his site screened or neutralized e-mail attachments at the gateway.

There are various application proxies that screen content and connection requests for the enterprise at the gateway. The malicious attachment or the backdoor program could have been blocked if a gateway solution contained a signature for them. Additionally, certain gateway proxies can stop web transactions to "blacklisted" sites, domains, and countries.

Most gateway proxies now support the Internet Content Adaptation Protocol (ICAP) protocol, which allows them to integrate with various engines for antivirus screening, malware detection, HyperText Markup Language (HTML) script examination, and content sanitization. For any host system or gateway screening solution, there should be a mechanism to ensure that the antivirus scanner is up-to-date and functioning properly.

Benefit: *Prevents malicious content from reaching the host system.* Known malicious files, known attacks, and evil web sites will be blocked before they can damage the user's system. Gateway proxies provide a scalable defense, in that they have visibility into all traffic and can redirect content off to any number of malware detection servers. They offer a point solution that scales to the enterprise.

Cost: These solutions can be costly, but offer a flexible and powerful way to defend your network's data. A gateway antivirus product, like any other security product, requires the management overhead of monitoring and updating.

4. Authenticate incoming e-mail

The forged e-mail to Bill appeared to come from one of his colleagues. Unfortunately, it is very easy to spoof e-mail addresses to make an e-mail seem to come from a legitimate source. The goal of this mitigation is to verify that the message actually comes from the user or the domain it says it comes from.

It is possible to detect or prevent forged e-mail with e-mail authentication technologies such as Sender ID or the Sender Policy Framework (SPF). Both of these technologies publish records in the Domain Name System (DNS) to validate that a received e-mail comes from the domain that it claims to originate from. Large Internet Service Providers (ISPs) and web-based email providers already implement these techniques to protect their users.

Other related technologies authenticate the content of e-mail, rather than its sender. They use cryptographic mechanisms to validate that the e-mail message has not been altered. These technologies include Secure MIME (S/MIME), Pretty Good Privacy (PGP) and DomainKeys Identified Mail (DKIM).

Benefit: *The integrity and legitimacy of incoming e-mail is ensured*. It becomes much more difficult for an adversary to spoof legitimate e-mail domains or modify e-mail content. It can also reduce the volume of SPAM received.

Cost: Sender ID and SPF records can be easily added to the DNS system to allow others to check e-mail that originates from your domain. It is also not difficult to set gateway e-mail servers to check for SPF or Sender ID records to determine if messages should be accepted. The e-mail sender must publish records in DNS that the recipient gateway can check; otherwise, his e-mail will not get through.

5. Turn on Data Execution Prevention (DEP)

If Bill's machine had had Data Execution Prevention (DEP) enabled, the attack may have been stopped when Bill opened the document containing the exploit.

DEP is a feature of the operating system that uses the processor hardware to prevent buffer overflows. Attackers exploit a buffer overflow to corrupt the systems memory and thereby get the attacker's code to run on the victim system. With DEP enabled, if a memory corruption is detected, the operating system terminates the application before the system can be compromised.

For more information on enabling DEP, go to http://www.nsa.gov/ia/ and search for "DEP".

Benefit: *Prevents malware from exploiting a buffer overflow to run on the host system.* DEP is a fundamental defense against most buffer overflows. It can protect most applications in your network and does not require signatures or regular updates.

Cost: DEP can be easily enabled on Window XP SP2, Windows Server 2003 SP1, and Windows Vista. Some applications may not function with DEP turned on; however, these applications can be omitted from DEP protection, if necessary.

6. Sanitize documents using MOICE

If the adversary's fake agenda document had been preprocessed by MOICE, it would have been rendered harmless when Bill opened it.

MOICE (Microsoft Office Isolated Conversion Environment) is a Microsoft-developed solution that translates Office documents into the Office Open XML format. Preprocessing with MOICE eliminates any illegally formed content, such as malicious code, that could enable an adversary to gain control of a machine. MOICE is installed and run on individual hosts.

For more information on implementing MOICE preprocessing, go to http://www.nsa.gov/ia/ and search for "MOICE".

Benefit: *Removes malware from Microsoft Office attachments when they are opened.* MOICE is a fundamental defense against Office documents that contain malicious code or macros. It can be deployed in a scalable way and requires little administration once enabled.

Cost: MOICE is available for free from Microsoft's web site. It can be enabled and completely administered via Group Policy. In some scenarios, enabling MOICE may result in minor usability problems for some users: For example, if they need to use documents containing macros.

SNAC DoD 9800 Savage Rd Ft.Meade, MD 20755-6704 410-854-6632 DSN 244-6632 FAX: 410-854-6604 **www.nsa.gov/snac**

7. Use an application firewall with egress filtering

The adversary that compromised Bill's machine needed a way to retrieve the interesting information from his machine. He also wanted to gain interactive access to Bill's machine and use it as a "beachhead" for a methodical attack against the network. When his malicious code exploited Bill's vulnerable word processing program, it contacted a server on the Internet and downloaded and installed a backdoor "Trojan horse" application. This backdoor periodically connected to an attacker-controlled server to receive instructions from the adversary.

Rogue programs that "phone home" can be blocked by host-based application firewalls. If properly configured, the firewall will allow only trusted applications (those on a "white list") to connect to the Internet. It can prevent both the malicious code that exploited the word processing program and the Trojan horse application from initiating communications with the attacker. Some Host-based Intrusion Prevention System (HIPS) products can also be configured to prevent applications from accessing the Internet.

Blackhole routing can also prevent non-proxy aware malware from communicating out of the network. Blackhole routing relies on there being one and only one way out of the network: through a proxy device. If the malware is not aware of the proxy being used, it will try and use the default route out of the network and its communication will be dropped (or routed to a collection point for analysis).

Benefit: *Prevents unapproved applications from communicating outside the network.* Host-based application firewalls and blackhole routing prevent backdoor programs from "phoning home." If the attacker's malicious program can no longer reach the Internet, he must use a more risky or technically difficult approach to remove the data from the network.

Cost: Application firewalls can be difficult to manage in a large enterprise. A user or malicious program may be able to modify the firewall rule set and allow an unwanted connection to the Internet. In addition, when you deploy a new application on your network, you may need to update the firewall rule set. Blackhole routing may require some changes to the network's configuration.

8. Perform regular virus scans with an up-to-date scanner

If Bill's system had been regularly scanned with an antivirus product, the backdoor Trojan might have been detected and removed before the adversary could use it to steal Bill's information—or at least before the adversary could do any more damage with it.

No single antivirus scanning solution offers full coverage of all known threats. A best practice approach is to screen the network using multiple virus scanners, to increase coverage. If possible, use different antivirus products on the gateways and the hosts to improve coverage.

Benefit: *Prevents known malware from persisting on systems.* Known malicious files can be detected and removed, ensuring that malware does not maintain a persistent presence on the network.

Cost: Antivirus products must be kept up-to-date and configured to run regularly. They must also be monitored to ensure that they are functioning as expected, and the logs they generate should be regularly reviewed.

SNAC DoD 9800 Savage Rd Ft.Meade, MD 20755-6704 410-854-6632 DSN 244-6632 FAX: 410-854-6604 **www.nsa.gov/snac**

9. Implement Software Restriction Policies (SRP)

The backdoor program that the attacker installed on Bill's system could have been prevented from running if Bill's administrators had implemented Software Restriction Policies (SRP).

Software Restriction Policies are built into Windows and allow administrators to control what programs can run on a machine. By allowing users to only run "white-listed" executables or executables in directories the users cannot write to, administrators can prevent untrusted software from running on the machine.

The main benefit of SRP is that it prevents unknown or untrusted code from executing and corrupting the integrity of the network. It can prevent exploits that require a downloaded file to execute, and social engineering techniques that trick a user into running a malicious file. For example, users that double-click an attachment that is executable content (not a document) will be prevented from executing the attachment, since it is not an approved application or in an approved location.

For more information on implementing SRP, go to http://technet.microsoft.com/en-us/library/bb457006.aspx.

Benefit: *Restricts execution of malicious code.* If configured to run only white-listed applications, SRP can prevent users from running programs that they download from the Internet or bring in on removable media.

Cost: SRP requires some overhead to set up and maintain. The policies must be developed and tested before they can be deployed operationally. If not thoroughly tested, some users may not be able to run applications that they need to perform their jobs.

10. Use a Host-based Intrusion Prevention System (HIPS)

If Bill's machine had been running a Host-based Intrusion Prevention System (HIPS), the harmful behavior of the malicious e-mail attachment could have been detected and prevented.

A HIPS provides security administrators with great flexibility in securing their networks. By monitoring at the application level, a HIPS can enforce specific policies of allowable program behaviors, where any harmful behaviors can be flagged and stopped. For example, a HIPS could be configured to restrict the behavior of a program such as Microsoft Word, so that it could not initiate Internet connections or write to certain folders. HIPS rule enforcement usually takes place in the operating system kernel, where it cannot be bypassed without high privilege.

Benefit: *Restrains and prevents malicious behavior of malware.* The benefits of a HIPS are its flexibility and strength in defending a network. HIPS solutions can wrap vulnerable applications so that they can't be exploited.

Cost: HIPS solutions can offer the broadest and most effective protection against malware. Unfortunately, the effectiveness of a HIPS is largely based on its configuration. Creating the appropriate configuration for a network can be a challenging task. A HIPS also requires monitoring to ensure that it is functioning correctly, and the logs they generate must be regularly reviewed.

SNAC DoD 9800 Savage Rd Ft.Meade, MD 20755-6704 410-854-6632 DSN 244-6632 FAX: 410-854-6604 **www.nsa.gov/snac**