



Mitigation Monday #2

Defense against Drive-By Downloads

June 2009 – Vulnerability Analysis and Operations Group, NSA



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.

Scenario

Bill, a government official, is reading his e-mail. One particular e-mail contains a hyperlink to a webpage about a product his division has been interested in. “Great!” Bill thinks as he’s about to click the link. Then he briefly pauses, remembering the last time he clicked on something in an e-mail. “But it’s not an attachment—it’s just a link,” he reasons to himself. “And it is work-related,” he continues, remembering his security training. He clicks on the link.

The webpage comes up and Bill peruses it for a bit. It gives some general information, but nothing particularly useful. It has some links to other information and Bill checks those out, but they’re nothing that Bill hasn’t seen before. Disappointed, Bill closes his browser and goes back to his e-mail.

He has no idea his computer has just been hacked.

In reality, an adversary carefully crafted an innocent-looking webpage that he knew would interest Bill. The adversary then sent Bill a spoofed e-mail with a link to this page, hoping to trick Bill into clicking on the link. Bill falls for this targeted “spear-phishing” attack. When Bill clicked on the link, malicious content in the webpage exploited a vulnerability in his Web browser and took control of his computer. The malicious content then made an Internet connection to the adversary’s server and downloaded and installed a “backdoor” program. This backdoor is a persistent foothold on Bill’s system, which the adversary can use to access Bill’s files whenever he wants to. This “drive-by download” of the backdoor was all done without Bill ever knowing that anything had happened.

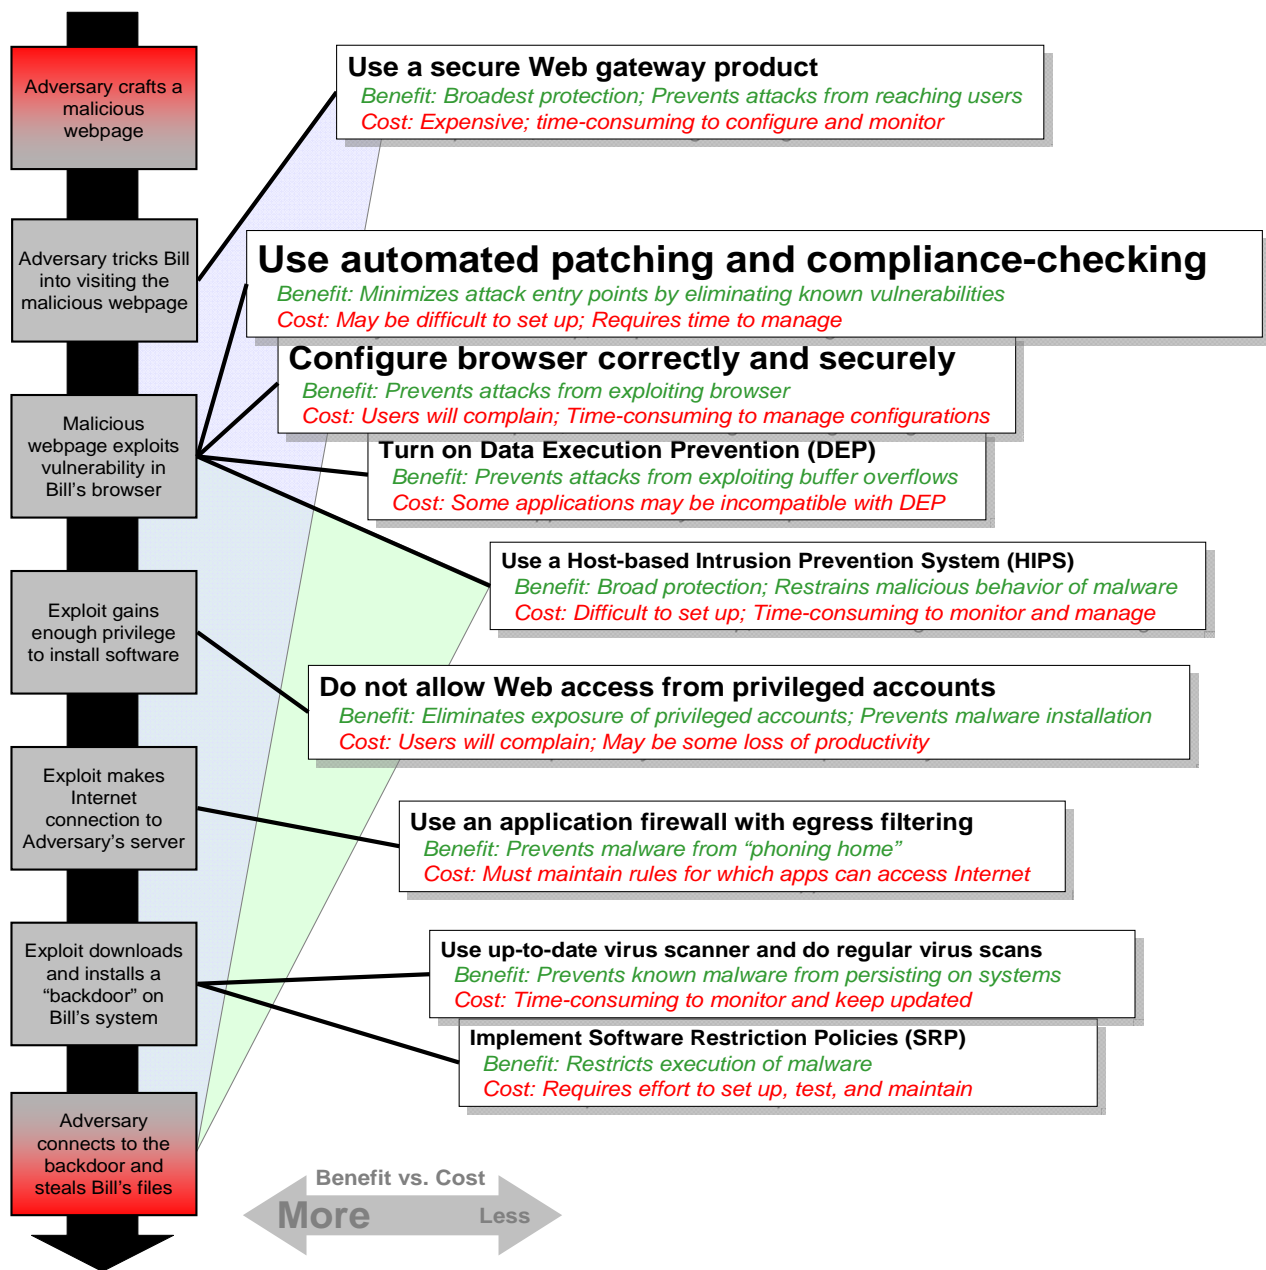
“Poor Bill,” the adversary muses, smiling, as he steals all Bill’s important files later that month—and uses Bill’s computer to compromise the rest of the network...

Mitigation

In the scenario, the attack was a drive-by download. 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.

Defenses Against a Drive-By Download



Defense Details

1. Use automated patching and compliance-checking

In order to gain control of Bill's machine, the adversary took advantage of a flaw in Bill's Internet browser or one of its plug-ins. Keeping browsers (including all their plug-ins, add-ons, and helper objects) and other applications that take input from external sources fully patched is an essential step toward defending the network.

Organizations 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.

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 (such as Mozilla's Firefox). 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. In rare cases, patching applications may break their functionality.

2. Configure browser correctly and securely

If Bill's Web browser had been configured more securely, the malicious content in the adversary's webpage may not have been able to compromise Bill's system.

Active content in webpages (JavaScript, Java, and ActiveX controls) is used to increase the functionality of the pages and to make them more interactive and visually interesting. Unfortunately, attackers can also use active content in a webpage to perform malicious actions. If possible, Web browsers should be configured to only accept active content from trusted sites; this is known as "white-listing" the sites. In general, white-listing (specify trusted and deny everything else) is more effective than "black-listing" (specify untrusted and allow everything else). This is because it's impossible to list *everything* that's untrusted in your black list.

Browser plug-ins can also be used by an adversary to compromise a system. All plug-ins should also be correctly and securely configured.

For suggestions on configuring the various Web browsers, go to http://www.us-cert.gov/reading_room/securing_browser/. For Internet Explorer, especially note the guidance on Security Zones. For Firefox, strongly consider using the NoScript add-on that is mentioned.

Some Web browsers, such as Mozilla's Firefox, cannot be centrally managed. This makes it difficult for your organization to push out configuration policy to these browsers. More seriously, there is no way to keep your users from changing (accidentally or otherwise) their browser preferences to something less secure. Third party applications can be used to manage and push policy to Firefox. Locking Firefox preferences, while not a complete solution, can at least protect novice users from themselves: see http://kb.mozillazine.org/Locking_preferences.

Benefit: *Prevents attacks from exploiting browser.* Web browsers that are correctly and securely configured to accept active content only from trusted sources will be much less vulnerable to exploits.

Cost: Properly configuring your Web browsers goes a long way toward improving your security, but refining and managing those configurations can be time-consuming, especially if you have multiple browsers or browsers that are difficult to manage. Restricting the content of websites for security reasons will very likely result in user complaints and some loss of productivity. Users will also want to add special plug-ins to customize their browsers; this could become a nightmare to manage.

3. Do not allow Web access from privileged accounts

In order to install the "backdoor" on Bill's system, the adversary took advantage of the fact that Bill was surfing the Internet while logged in as a local administrator on his machine. If Bill had been logged in as a non-privileged user instead, the malicious webpage 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 and run code.

A vital step in mitigating attacks is to stop users from surfing the Web and reading e-mail 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. Note however that an adversary does not *need* to have high privilege in order to steal information. After all, the information he wants is most likely in the user's documents folder anyway! This is why a defense-in-depth approach is very important.

Benefit: *Eliminates the exposure of privileged accounts and denies adversaries the ability to install malware as part of the operating system.* 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. Users can also use the "Run as" feature to access administrative functions. 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.

Another option would be to only allow user Web browsing from inside a virtual machine (VM). A VM can be used to run programs in an isolated environment where they will not affect the real host. When the user is done browsing, the VM can be shut down and any persistent threat would be eliminated.

4. Use a secure Web gateway product

If Bill's Internet traffic had been directed through a secure Web gateway, the adversary may have been prevented from even getting close to Bill's machine, much less exploiting it and stealing information from it.

A secure Web gateway (or Web security gateway) is an Internet gateway that employs multiple methods (URL filtering, anti-malware, policy enforcement, Web application control, authentication, etc.) to offer the broadest protection for your Internet users. It also functions as a proxy, which allows it to scan encrypted content and monitor all Web activity. This "all-in-one" protection is also easier to manage than trying to coordinate and manage multiple security products.

Benefit: *Prevents attacks from reaching users' systems.* A secure Web gateway is an "all-in-one" protection solution. It can block access to malicious websites and stop malware from being downloaded. A secure Web gateway offers a point solution that scales to the enterprise.

Cost: A secure Web gateway offers the broadest protection against Web attacks. However, they are expensive, time-consuming to configure correctly, and may require some changes to your network. Due to the scanning and proxy functionality, a secure Web gateway may slow down your Internet access. In addition, it will require monitoring to ensure that it is functioning correctly, and the logs it generates should be regularly reviewed.

If you can't spend the money on a secure Web gateway product, consider at least directing all your Web traffic through a proxy. This will allow you to implement simple filters on all your Internet traffic, including your encrypted traffic. For added security, use an authenticating proxy.

With a proxy in place, "blackhole routing" can be used to 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 your proxy device. No network router knows the path to the proxy; an application must explicitly contact the proxy's IP to get out to the Internet. So if the malware is not aware of the proxy being used, it will try to use the default route out of the network. No router will know how to route the malware's traffic and it will be dropped (or routed to a collection point for analysis).

Note that regularly reviewing the logs generated by your network defenses (gateways, proxies, firewalls, virus scanners, HIPS, etc.) is extremely important. This may be the only indication you get that your network has been compromised by an adversary. For example, check your gateway/firewall logs for anything out of the ordinary: communication with unexpected IP addresses, non-business hours activity, excessively large data transfers, etc. Consider implementing a system so that your network administrators are automatically informed when anomalous events occur. But don't get too overzealous! Bad guys generally continue to do bad

things over time, so look for multiple indications of malicious activity. Make sure that your “adversary” is not just a regular user who made a one-time mistake!

5. Turn on Data Execution Prevention (DEP)

If Bill’s machine had had Data Execution Prevention (DEP) enabled, the attack may have been stopped before the adversary could gain a foothold on Bill’s system.

DEP is a feature of the operating system that uses the processor hardware to prevent buffer overflows. Attackers can exploit a buffer overflow to corrupt memory and get their code to run on the victim system. With DEP enabled, if 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”.

Internet Explorer 8 and Firefox 3 are protected by turning on DEP, but older versions of these browsers are not. This is a security risk, as the browser is exactly what the adversary is trying to exploit! Upgrade to the newest versions of these browsers as soon as possible. DEP can be enabled (via a browser setting) for IE7 running in Vista, but IE7 users running Windows XP are out of luck, as are all Firefox 2 users.

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. DEP is an effective way to protect against zero-day exploits that try to take advantage of a buffer overflow vulnerability in an application before it can be patched.

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

6. Use an application firewall with egress filtering

The adversary that compromised Bill’s machine needed a way to retrieve the interesting information from it. 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 the adversary’s malicious webpage exploited Bill’s vulnerable Web browser, 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. The firewall can be configured to allow only trusted (white-listed) applications to connect to the Internet. It can prevent 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.

Benefit: *Prevents unapproved applications from communicating outside the network.* Host-based application firewalls 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.

7. Use up-to-date virus scanner and do regular virus scans

Most antivirus products have realtime protection, or “guard,” functionality that will automatically scan files when they are written to or accessed from disk. If Bill’s system had had this turned on, the backdoor Trojan may never have made it onto his computer. Even if it was not caught that early, if the antivirus product had been kept up-to-date and if Bill’s system had been regularly scanned, 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. Do *not* try installing multiple virus scanners on each of your hosts! Antivirus products generally do not “play well” together on a single machine.

Benefit: *Prevents known malware from persisting on systems.* Known malicious files can be detected and either kept off of systems or removed later, 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. Turning on realtime protection may slow down file access times slightly.

8. 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 application whitelisting solutions such as SRP is that they prevent unknown or untrusted code from executing and corrupting the integrity of the system. They can prevent exploits that require execution of a downloaded file, and social engineering techniques that trick a user into running a malicious file. For example, the backdoor Trojan in the scenario may have been prevented from executing, since it was 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>.

Consider carefully whether you should check the option to apply SRP for the local administrator accounts on your machines. If you do, this may make administering those machines more

difficult. If you do not and your users are still surfing the Web from their local administrator accounts, then they (and thus the exploit) can circumvent any SRP restrictions. If you do not apply SRP for the local administrators, SRP will only be effective if your users are not allowed Web access from privileged accounts.

Benefit: *Restricts execution of malicious code.* If configured to run only white-listed applications, SRP can prevent users from running programs that they download (knowingly or not) 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.

9. 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 content in the adversary's webpage 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, and any harmful behaviors can be flagged and stopped. 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 broad and 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 it generates must be regularly reviewed.

Additional Mitigation Considerations

First, it is worth noting that an adversary has other options to direct unsuspecting users to his malicious webpage besides phishing. With a DNS poisoning attack, the adversary can direct users who are trying to go to a legitimate webpage to instead go to his malicious webpage. Be sure to keep your DNS servers upgraded to prevent this. If your organization serves Web content, the adversary could gain access to your Web servers and add his malicious content to *your* webpages, thereby compromising everyone who surfs to your sites. Make sure that your Web servers are properly hardened and that your Web applications are developed with security as a primary consideration.

Second, although user education cannot be relied upon to prevent compromises—because users will make mistakes or just plain not listen—it is still an important part of your network defense. There are numerous places on the Internet to find “safe surfing” guidance (for example, see the “Safe Browsing” section of <http://www.us-cert.gov/cas/tips/>). Make sure that your users know

how to browse the Web securely. Also make sure that your users know what's in your *current* security policy concerning Web browsing.

Third, there is a class of solutions known as Data Loss Prevention (DLP) or “Extrusion Prevention” systems. These systems can mark important data and prevent it from leaving your network. Currently these systems are challenging to set up and of limited effectiveness, but this could be a promising technology in the future.

Finally, if you find that implementing (and verifying the implementation of) these suggested defenses is too difficult on your network, then your network may not be manageable. If so—if you can't verify that your security protects *every single device* on your network—then your network is insecure. Your first priority should be to get your network manageable; this will automatically improve its security. For more information, request a copy of the *Manageable Network Plan* at manageable@thematrix.ncsc.mil.