VERDASYS

ADVANCED PERSISTENT THREAT (APT), DETECTION, MITIGATION, AND PREVENTION

OCTOBER 2011

WHITE PAPER

"The United States faces a significant and ongoing cyber security threat today; one that presents issues of national and economic security… These cyber espionage attacks result in massive losses of private sector intellectual property and sensitive government information… I don't believe that there is a precedent in history for such a massive and sustained intelligence effort… to blatantly steal commercial data and intellectual property."

Rep. Mike Rogers (R-MI), Chairman,
House Permanent Select Committee on Intelligence,
Open Hearing: Cyber Threats and Ongoing Efforts to Protect the Nation,
October 4, 2011

Attack Phase 1:

A research scientist at a chemical manufacturer receives an email from a colleague with a file named "presentation" attached. The message simply says "You did a great job here", so the scientist opens it to find content he had presented at an industry event the prior year, that had been posted online by the event organizer. Thinking nothing of it, the scientist closes the email and goes about his work.

But, this was no ordinary attachment, and the email was not really from someone he knew. As a senior researcher cleared to work on classified government projects the scientist had been targeted for a highly-personalized "spear phishing" campaign designed to be the entry point for a multi-phased cyber attack launched halfway around the world. Detailed personal and professional information about the scientist easily found on his Facebook and LinkedIn accounts and other web searches allowed the attacker to create a believable ruse to exploit a common weakness of corporate networks — the unsuspecting user.

Using a common spear phishing tactic that causes someone to unknowingly infect their own machine, the email was "spoofed" to look exactly like an email from a person the scientist trusted forwarding an innocuous attachment. The attacker exploited a previously undiscovered vulnerability in the application used to open the attached file to surreptitiously embed malware that would launch, undetected, the moment the scientist opened the attachment. If successful, it would ultimately allow a remote user to silently gain access to any data on his company's network.

A mission to steal the company's intellectual property was underway.

Sound incredible? Not to anyone responsible for defending sensitive data from the escalating cyber warfare being waged by sophisticated and well-funded attackers against governments, businesses, and citizens worldwide. The real-world scenario detailed in this paper illustrates a new type of targeted and sustained effort by hackers to compromise "secure" networks. This next-generation of online attacks is so different and more dangerous to companies relative to "classic" malware it has earned its own categorization, dubbed Advanced Persistent Threats (APT).

APT is often associated with government espionage, but it actually encompasses a wide range of missions, from corporate data theft to outright destruction of infrastructure. No organization is totally safe from all APT; attacks have been confirmed in manufacturing, high tech, oil & gas, banks, pharmaceutical, critical infrastructure and public utilities — and, of course, military & government networks.

Attack Phase 2:

The malware begins its next phase by executing itself in the infected laptop's memory. As is typical with APT, this initial stage of attack operates at the kernel ("root") level of a user's machine, subverting the operating system's security features and defensive applications. Once it has gained root access, the APT program hijacks an existing system account with "super user" privileges and hides its presence on the local machine to evade detection by typical antivirus and intrusion prevention technologies that rely on the operating system. The attacker now has absolute authority over a trusted insider's machine to search for sensitive data on it, or from which to spread and embed new malware for the next stages of the APT mission.

An APT mission successfully combines social engineering with a mix of old and never-before-seen hacking techniques to achieve its objectives. First, it must remain unnoticed in order to penetrate and operate within a network indefinitely; second, it must be able to adapt to and exploit secure IT environments; and finally, it must be able to covertly communicate with an external source to receive additional commands or to exfiltrate information. Persistence is key to an APT mission's success, and if one combination of techniques fails along the way, a new one will be tried.

Beyond stealth, the ability to overcome challenges during a mission by applying updated instructions is a necessary feature for APT attacks to succeed. To do this, a "command and control" channel is established from which to send and receive information from outside the network, allowing malware code to be revised based on new network intelligence. The ability to adapt on the fly significantly increases the chances of an APT mission's success; however, this communication link can also expose an attack to anomaly detection by the right counter-technology.

Attack Phase 3:

Now operating as a rogue agent from a privileged user's system, the APT attack enters its next stage: discover the location of top secret data for a topical anthrax vaccine in Phase III clinical trials, and acquire the credentials needed to access it. The malware first creates an encrypted index of files and emails on the scientist's machine which contains combinations of the keywords "vaccine", "anthrax", "confidential", "secret", "restricted", "defense", and "military"; it then searches for regular expressions (file content, buffers, and headers) that identify machine names, network addresses, and applications associated with any keyword "hits". It also installs another application to record keystrokes ("keylogger") whenever someone logs onto the network, or connects to a specifically identified resource (e.g. network address or application).

Eventually, the malware begins systematically connecting to target systems using words recorded by the keylogging utility. In order to remain unnoticed, it attempts to emulate normal traffic usage so as not to present as a suspicious pattern to a network scanner. Once all application/system/credential permutations are exhausted without success, the malware begins spreading itself to other machines using a domain administrator's credentials that were captured when he accessed the scientist's machine to perform a routine file backup.

Defending against an APT attack requires equally skillful counter-measures that can identify threatening behavior by predicting the tactics necessary for it to succeed. This can be done by analyzing the legitimacy of abnormal system and network activities that can be traced to a suspicious origin. But, what is the difference between abnormal and malicious behavior? The variance of this answer on any given network is what makes APT so difficult to detect without crippling the business process with too many false alarms. However, depending on the APT's goal (which is based on the company and its data) it is possible to foresee certain threatening events because they will be required to complete the mission.

For instance, if the APT attack is designed to steal sensitive data then it *must* first infiltrate the network. As well, it *must* find the targeted data without knowing where it might be among the thousands of potential systems (if a large organization). And, unless the data can be found on the first machine infected, APT *must* spread itself to as many systems as necessary to locate its target's location and acquire the proper credentials to access it. Finally, this type of APT attack *must* be able to "exfiltrate" stolen data off the network once it has been found.

Of course, the methods and tactics may vary immensely, but *all* APT stages must be accomplished this type of the mission to succeed. Therefore, this requirement inherently provides several opportunities to deploy effective APT countermeasures.

Attack Phase 4:

Now possessing the ability to access virtually any system on the network, the malware quickly gains control over other machines and replicates the investigative process to locate its target data and gather access credentials using content matching and keylogging. After several months of silent reconnaissance the malware finally discovers a file, created by a customized application, which connects to a restricted database. Although encrypted, the malware accesses the database using credentials recorded from the infected workstation, and runs a series of SQL queries to map the schema. The APT then sends the encrypted query results to its external master confirming a potential target acquisition, and awaits further instructions.

The first clue beyond the initial compromise that an APT attack has begun is often abnormal communications between user workstations. Other suspicious events could include any unusual access of sensitive applications or restricted network shares during off hours hours, or encrypted outbound communications using atypical port/protocol combinations.

Attack Phase 5:

The malware receives updates from its master with specific SQL strings to extract the sensitive data, which is then staged on yet another compromised machine within the network. Over the next several weeks, the extracted files are covertly sent to several web addresses over standard ports and protocols to blend in with typical network traffic patterns. Finally, after each phase of the operation is completed the malware deletes evidence of the attack – including accounts, files, settings, and system logs – which could be used to trace its presence or activities, except for a minimal "backdoor" capability in order to re-install later.

The APT's mission has been accomplished without detection.

Detecting, Mitigating, and Preventing APT Attacks with Digital Guardian

The most effective defense against APT's requires a unified and layered approach that can detect, segregate, and neutralize malicious software at multiple mission stages. Digital Guardian is a data-centric technology platform for Enterprise Information Protection (EIP) that integrates endpoint and network agents into a coordinated and multi-layered sentry capable of stopping an APT attack by challenging its adaptability and stealth. Digital Guardian agents can be relied upon to confirm malicious activities on every level of network or system operations without using signature or heuristic techniques relied upon by antivirus technologies that APT's can easily defeat.

Digital Guardian detects potential APT-related activities on the endpoint based on policy rules defining what a certain user is allowed to do with certain data under certain operating conditions. It then monitors and responds to policy violations on workstations, laptops, and servers using hardened kernel-level agents with stealth and tamper-resistant capabilities.

The Digital Guardian security model operates autonomously from the user and independently from the operating system. Digital Guardian agents alone determine which system and data-level transactions are allowed by policy based on a file's sensitivity and the employee's respective usage rights, regardless of his or her administrative privileges. This provides several layers of systematic defense against an APT attack without interrupting normal business processes, including preventing the use of compromised administrative accounts to access encrypted data and maintaining encryption if a sensitive file is moved off a compromised workstation or server.

Digital Guardian endpoint agents use a powerful combination of technologies which provide early warning and tactical responses to thwart APT at first contact. These capabilities include advanced memory scanning and forensics to detect suspicious executable code present in memory; policy-based prompts that require a user to manually acknowledge or justify a risky data transaction before it can occur; policy-based access controls (including encryption) for files in-use or in-motion based on user privilege and file sensitivity; and application monitoring to identify and block the launch of dangerous executables.

Digital Guardian network agents also provide session-level analysis of incoming and outgoing traffic to detect potential APT activities. APT's often use obfuscated payloads, encrypted messaging, and vulnerabilities to pass by perimeter defenses undetected. In response, Digital Guardian network agents provide deep session inspection at line speed across all ports and protocols that can detect and block events like connections from suspicious IP addresses, encrypted traffic, or application port-hopping; they can also deconstruct payloads to see inside embedded content (e.g. JavaScript within a PDF file) that could be used as an attack vector.

The Digital Guardian technology platform provides a comprehensive and integrated defense-in-depth security model to stop APT without disrupting unaffected systems or network segments. Digital Guardian combines unmatched visibility on systems and sessions to identify multiple suspicious activities which, once correlated, can expose an active APT attack that can be monitored and mitigate through a common management interface. If and when an APT event is confirmed, Digital Guardian assures an organization can protect itself with actionable intelligence and real time responses capable of stopping the mission at multiple stages, while ensuring critical data remains contained and used productively throughout the business process.

ABOUT VERDASYS

Verdasys (www.verdasys.com) provides Enterprise Information Protections (EIP) solutions to manage proprietary data and assure the integrity of business processes in highly collaborative and mobile environments for the world's most secure organizations. Verdasys, a leader in the 2011 Content Aware Data Loss Prevention (DLP) Magic Quadrant, offers the industry's most complete data security platform for identifying sensitive information; auditing and controlling its uses by privileged insiders; and preventing its compromise by targeted cyber attacks. Verdasys Digital Guardian protects sensitive data across mobile and virtual environments, file servers, emails, and networks for government agencies and global leaders in manufacturing, automotive, technology, financial services, insurance, biopharmaceuticals, and healthcare. Companies serious about information protection choose Verdasys.

VERDASYS

Corporate Headquarters
404 Wyman Street
Waltham, MA 02451 USA
info@verdasys.com
781-788-8180

www.verdasys.com