TREND MICRO™

Countering the Advanced Persistent Threat Challenge with **Deep Discovery**

A Trend Micro White Paper I February 2013

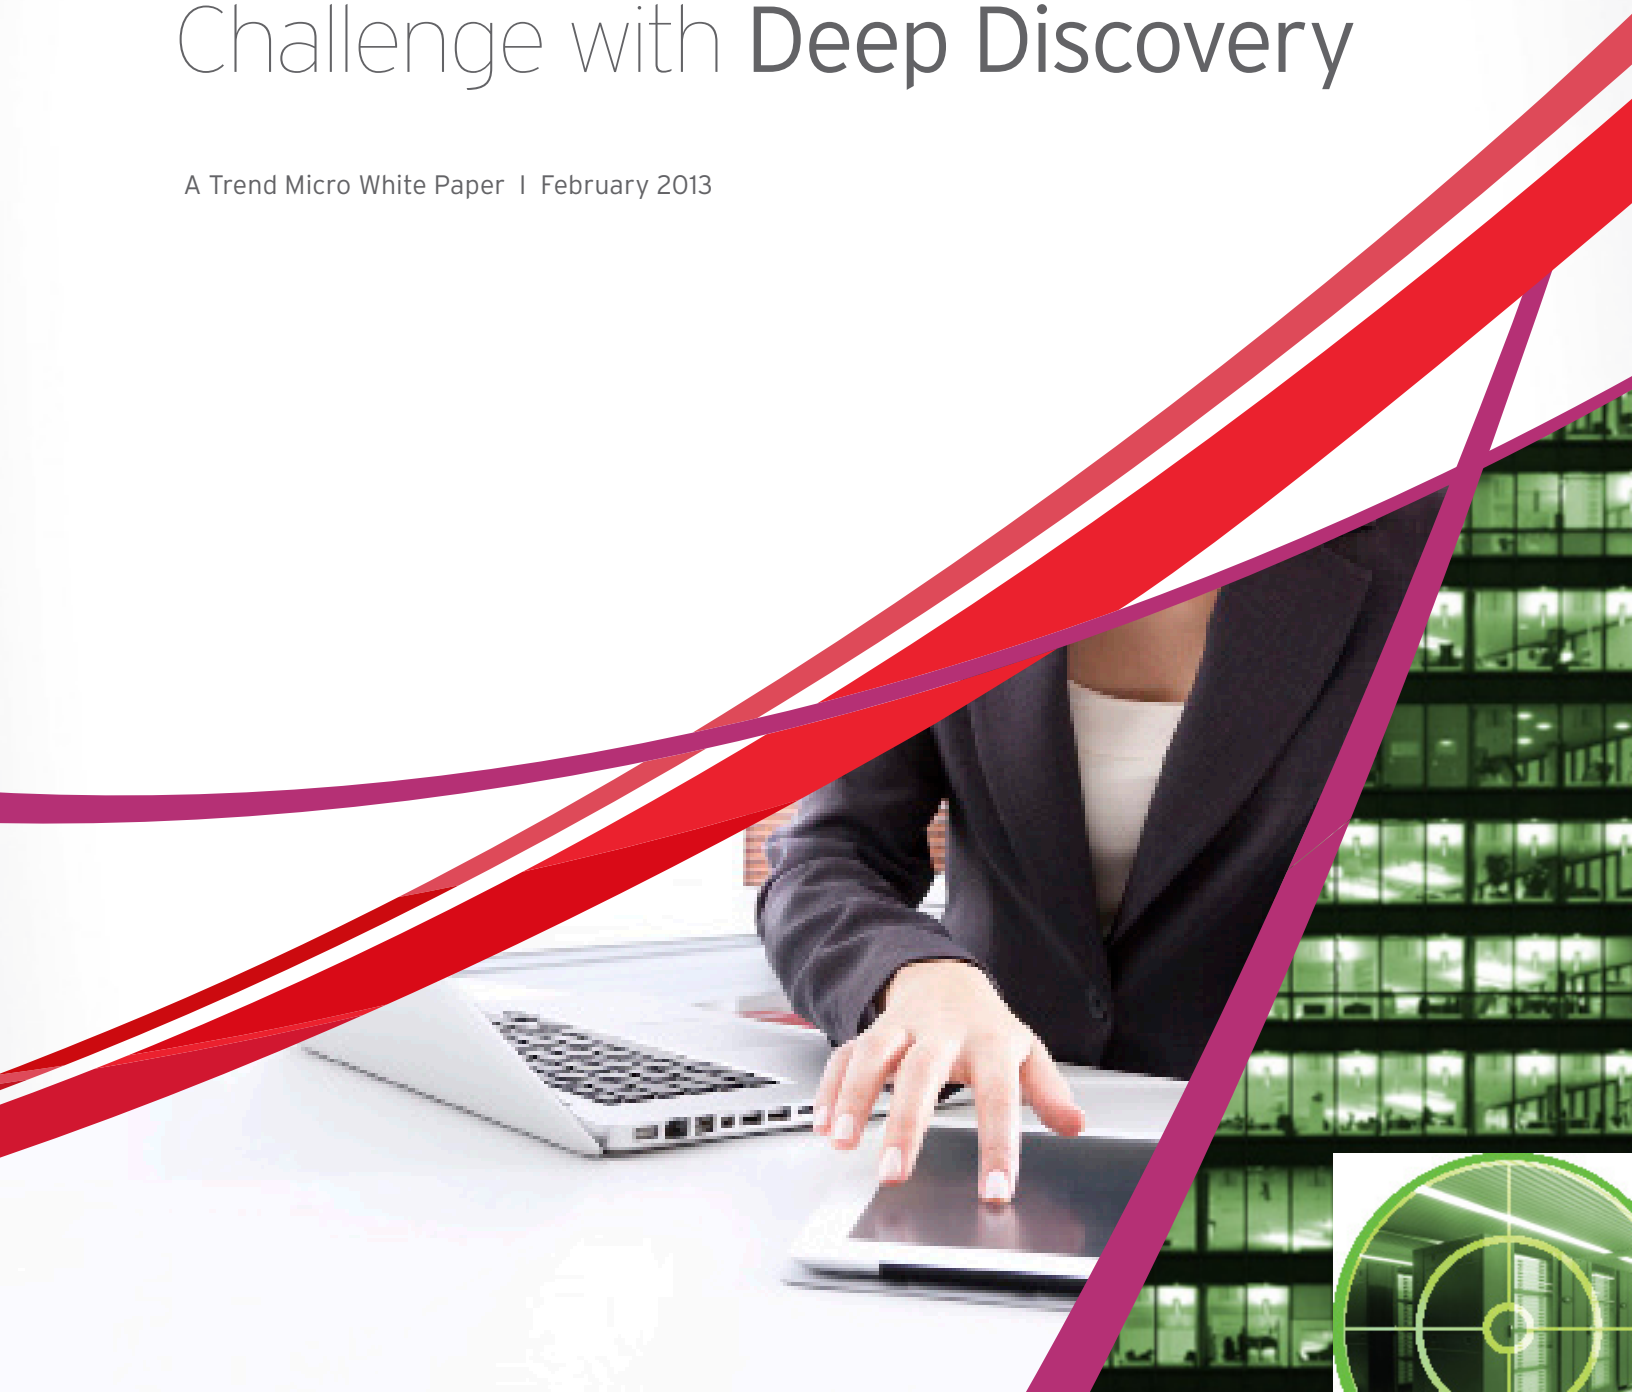

Contents

Executive Summary

Targeted attacks and advanced persistent threats (APTs) are quickly becoming the new norm of cyber security threats—encompassing organized, focused efforts that are custom-created to penetrate enterprises and government agencies for valuable data, trade secrets, and access to internal systems. Significant breaches at RSA, Citibank, and Global Payments have made headlines, and according to a recent ISACA member survey, 21% of respondents reported that their enterprise has already been victimized by an APT, and 63% think it is only a matter of time before their enterprise is targeted.[1]

While traditional security products can defend against malware and other known vulnerabilities, they are ineffective against this new era of custom, targeted, never-been-seen-before, "slow and low" attacks. According to Gartner Research, "There is widespread agreement that advanced attacks are bypassing our traditional, signature-based security controls and persisting undetected on our systems for extended periods of time. The threat is real. You are compromised; you just don't know it."[2]

Combating these custom attacks requires a new approach—one that closes the security gap created by this new generation of stealthy, sophisticated targeted attacks. An organization's strategy against APTs should utilize an approach that takes into account how targeted attacks infiltrate and work inside an organization, and it should also provide custom detection and intelligence adapted to the organization and its attackers. In addition, an ideal solution should integrate advanced detection technology into an organization's existing endpoint and gateway defenses to help strengthen detection of targeted attacks.

This whitepaper will explore the anatomy of targeted attacks: the inner workings of the APT lifecycle. It will also provide an in-depth overview of Trend Micro Deep Discovery advanced threat protection solution, and how it enables enterprise IT to adopt a custom defense strategy that modernizes its risk management program to defend against targeted attacks. Deep Discovery is at the heart of the Trend Micro Custom Defense solution against targeted attacks.

The Anatomy of a Targeted Attack

APTs are highly sophisticated and are a reality for both small and large organizations. Already, we have seen the likes of ShadowNet, Flame, Global Payments, and the recently discovered Red October campaign as successful examples of carefully crafted attacks focused on specific goals in targeted entities. While cyber-attacks previously employed a mass scale approach that readily enabled the creation of security signatures, advanced malware not only disguises its presence but also goes to great lengths to hide its communications within seemingly legitimate network traffic—making it virtually impossible to defend against using traditional, signature-based approaches.

The goal of this "one-to-one," stealthy technique is to steal valuable intellectual property, money, and other personally identifiable information (PII). Given the customized nature of these advanced malware attacks, they have a high rate of success and have resulted in extensive cost to organizations. As recently as January 2013, Global Payments, Inc. reported an updated figure of $93.9 million in costs associated with the data breach they discovered in April 2012.

In order to understand how APTs are so successful, it is important to take a deeper look at real examples to gain insight into the attack sequence and techniques used. Here, we'll review the **RSA**, **Diginotar,** and **Luckycat** attacks.

[1] ISACA, Advanced Persistent Threats Awarness, February 2013
[2] Gartner Inc., Best Practices for Mitigating Advanced Persistent Threats, Report G00224682, 18 January 2012

RSA

In March 2011, when EMC disclosed an attack against its RSA division that successfully stole SecureID data, it quickly made national headlines—especially due to the millions of RSA SecureID tokens in use at the time, providing protection to corporate networks and smartphones. It was subsequently discovered in June 2011 that targeted attacks against Lockheed Martin, L-3 Communications, and Northrop Grumman were made possible from the SecureID data obtained in the successful RSA breach.

ATTACK OVERVIEW:

1. Two spear phishing emails were sent over a two-day period targeted at low to mid-level employees with subject "2011 Recruitment Plan" and .xls attachment with the same title
2. .xls file contained an exploit through an Adobe Flash zero-day vulnerability that installed a backdoor using a Poison Ivy RAT variant set in a reverse-connect mode
3. Attackers moved laterally to identify users with more access and admin rights to relevant services and servers of interest
4. Access was then established to staging servers at key aggregation points
5. Data of interest was moved to the internal staging servers, aggregated, compressed, and encrypted for extraction
6. FTP was then used to transfer password protected RAR files to a compromised machine at a hosting provider
7. Files were subsequently removed from the host to cover up traces of the attack

DigiNotar

In August 2011, a network compromise was discovered at DigiNotar, a former Dutch certificate authority (CA), which led to the issuing of fraudulent digital certificates—used to make malicious Web sites and malware look legitimate. In particular, valid certificates were obtained for a number of high-value domains, including Yahoo, Mozilla, and Google, who discovered the fraudulent certificates in use in a large-scale, Man-In-The-Middle (MITM) attack on 300,000+ of its Gmail users—who were being eavesdropped on for weeks before detection. By September, DigiNotar filed for bankruptcy and shut down, and all major browser and operating system vendors revoked all DigiNotar signed certificates. The DigiNotar breach and subsequent Google MITM-attack resulted in an erosion of public trust in the existing Public Key Infrastructure.

ATTACK OVERVIEW:

1. Attacker located compromised network access—Web servers in external DMZ were breached
2. Diginotar used a highly segmented, ringed network defense methodology to protect its eight CA servers, and attackers methodically compromised one ring at a time to gain access
3. With access to the CA servers, more than 531 (identified) rogue certificates were issued
4. Rogue certificates were transmitted to attacker's external IP server, using a proxy tunneling tool

Luckycat

The Luckycat APT, active since June 2011, is a broad attack tied to a gang of Chinese cybercriminals aimed at more than 90 targets in Asia, including executive level employees in the aerospace, energy, and engineering industries with the initial intended goal of remaining on infected systems, covertly monitoring activity over an extended time frame. Subsequently, in July 2011 research uncovered malware in the form of two, unfinished and undelivered Android apps that communicate with Luckycat's C&C server.

ATTACK OVERVIEW:

1. Attackers researched intended targets, including sensitive entities in Japan and India as well as Tibetan activists
2. Spear phishing emails were predominantly used as point of entry—one aimed at a Japan target used the confusion after the Great East Japan Earthquake to lure potential victims into opening a malicious .PDF attachment and another aimed at an entity in India lured victims into opening a .DOC file on India's ballistic missile defense program
3. Once opened, both example targets were displayed a decoy document while zero-day vulnerabilities in .PDF and in .DOC enabled the malware, TROJ_WIMMIE, to install
4. Exfiltrated data, including attack campaign codes, victim's identity details and contents of the compromised computers and servers were sent back to Luckycat C&C servers

From these attack examples, the level of sophistication in targeted attacks is readily apparent. It is also notable that initial attacks are often linked to subsequent attack targets (e.g. the RSA attack linked to Lockheed Martin) or repeatedly executed by an inter-connected, criminal network that uses or even shares the same infrastructure or malware components. While these successful attacks were all custom tailored to their attack target, each followed a carefully staged lifecycle to enter the intended organization and retrieve the desired data before detection. With further analysis of the anatomy of these example APTs, six distinct stages become apparent:

The APT Attack Sequence

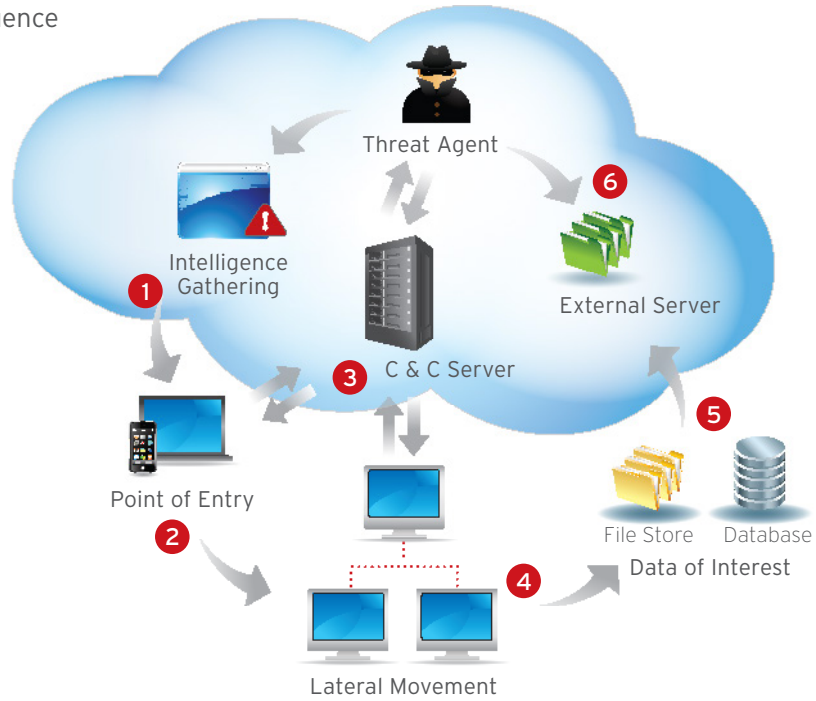

Stage 1: Intelligence Gathering

In this stage, cyber criminals have their attack targets in mind and conduct research to identify target individuals within the organization—most likely leveraging social medial sites, such as LinkedIn, Facebook, and MySpace. With the wealth of personal information provided on these sites, attackers arm themselves with in-depth knowledge on individuals within the organization—for example, their role, hobbies, trade association memberships, and the names of those in their personal network. With this information in hand, attackers prepare a customized attack in order to gain entry into the organization.

As seen in the successful attack against RSA, the criminal's intelligence and gathering phase focused on identifying a small group of employees within two groups to target with a well-crafted and compelling email. According to RSA, the targeted employees weren't considered "particularly high profile or high value targets." This research approach has become commonplace, whereby employees within a certain department or with a desired management level are targeted, which also demonstrates the importance in educating employee about security awareness.

Stage 2: Point of Entry

Once cyber criminals have gathered the intelligence on their intended target, they begin work on designing their point of entry into the organization. While in the Diginotar attack example, a network compromise was found, initial compromises are typically from attacks that exploit zero-day vulnerabilities to place malware on the system, which are most often delivered via social media (email, IM or drive-by download) as seen in the RSA and Luckycat examples. In fact, recent Trend Micro research has discovered 91% of targeted attacks involve spear phishing.

In the RSA example, the attack began with spear phishing emails sent to targeted employees with an excel attachment titled, "2011 Recruitment Plans." When the employee opened the spreadsheet, it ran malware that exploited a previously unknown Adobe Flash zero-day vulnerability (CVE-2011-0609) to install a Poison Ivy Remote Administration Tool (RAT).

In the broad Luckcat APT campaign, a variety of spear phishing emails were sent across the spectrum of targets that exploited Adobe and Microsoft vulnerabilities to penetrate networks. As seen in the Japan target example, the attackers used the 2011 Japan earthquake disaster to lure potential victims into opening a malicious .PDF attachment with information containing radiation dose measurement results, and once opened, it exploited a vulnerability in Adobe Reader—CVE-2010-2883, to install malware (TROJ_WIMMIE) onto the target's system.

Stage 3: Command and Control Communication

Once the malware is successfully installed on a compromised machine, it is able to communicate back to the cyber criminal's command and control (C & C) servers for further instructions or download additional malware and attacker tools. This C & C connection allows the attacker to instruct and control the compromised machines and malware throughout all subsequent phases of the attack and to establish long-term access to the network. Most commonly, this includes the download of additional malware executables following the initial infection, such as, key loggers, Trojan backdoors, and password cracking tools.

The Luckycat campaign extensively used free hosting services for its C & C servers. The domains used were considerably diverse, and all were available from three, free hosting services. As such, the attackers had extensive resources to continue creating diverse domain names for their C&C servers.

DOMAIN	EMAIL ADDRESS
cattree.1x.biz	lindagreen56@rediffmail.com
charlesbrain.shop.co	yamagami_2011@mail.goo.ne.jp
footballworldcup.website.org	ajayalpna@hotmail.com

Sampling of free web-hosting service domains attackers used for Luckycat C&C servers

In the Diginotar APT example, the attacker clearly had frequent C&C communications as efforts were made to steal fraudulent certificates over the span of the attack from June 17 to July 24, 2011, and in the RSA breach, attackers used a Poison Ivy RAT set in reverse-connect mode to remotely manage the attack from their external location.

Stage 4: Lateral Movement and Persistence

Once inside the network, the attacker moves laterally within the organization to compromise additional machines in order to harvest credentials and gain escalated privilege levels. The attacker will also acquire strategic information about the IT environment—operating systems, security solutions and network layout—to maintain persistent control of the target organization. In the attack against Diginotar, several tools were used to increase the intruder's level of access in the network, including, port redirectors, scanning tools, and remote process executor tools.

In the RSA breach, attackers obtained login credentials from the first compromised accounts, including usernames, passwords, and domain information, and then pursued higher-value accounts with more access privileges. According to Uri Rivner, former Head of RSA New Technologies and Identity Protection, "This is one of the key reasons why, having failed to prevent the initial social engineering phase, detecting [an APT] quickly is so important. In many APT [examples] the attackers had months to do digital 'shoulder surfing' on the attacked users, map the network and resources, and start looking for a path to the coveted assets they desired. Then they use the compromised accounts, coupled with various other tactics to gain access to more 'strategic users.' In the RSA attack, the timeline was shorter, but still there was time for the attacker to identify and gain access to more strategic users."

Stage 5: Asset/Data Discovery

In an advanced malware attack, cyber criminals are in pursuit of a high valued asset. This could be anything from financial data, trade secrets, or source code, and most noteworthy, attackers know the intended data of interest when a target organization is selected. As seen in the RSA breach, attackers pursued the company's SecureID two-factor authentication data, and in the attack against Diginotar, rogue certificates were created and stolen. The attacker goal is to identify the data of interest as quickly as possible without being noticed.

In the asset and data discovery phase, the attacker uses several techniques to identify the noteworthy servers and services that house the data of interest, for example they will:

- Check the configuration of the infected host's email client to locate the email server
- Locate file servers by checking the host for currently mapped network drives
- Obtain the browser history to identify internal Web services, such as CMS or CRM servers
- Scan the local network for folders shared by other endpoints

Stage 6: Data Exfiltration

In this final stage of a targeted attack, sensitive information is gathered and then funneled to an internal staging server where it is chunked, compressed, and often encrypted for transmission to external locations.

In the Diginotar attack, once the Certificate Authority servers were compromised, the criminals used their complete access to create fraudulent but valid certificates and exfiltrated them from Diginotar's network to their own location. In the RSA attack, once the criminals located the data they wanted to steal, they gathered it in a staging area, compressed it, and then exfiltrated it via FTP.

Trend Micro Deep Discovery: Custom Defense against APTs

Standard protection products' signature-based, one-size-fits-all approach cannot deal with the custom nature of targeted attacks and their dedicated perpetrators. The malware, communications, and attacker activities used in targeted attacks are invisible to standard endpoint, gateway, and network security measures.

Security analysts and experts recommend a new type of network monitoring that uses specialized detection and analysis techniques designed specifically to discover the telltale signs of these attacks. Trend Micro Deep Discovery is a leading product in this movement, enabling organizations to deploy a full Detect-Analyze-Adapt-Respond lifecycle to protect themselves from these attacks.

The Deep Discovery solution is comprised of two components: Deep Discovery Inspector and Deep Discovery Advisor. Deep Discovery Inspector provides network threat detection, custom sandboxing, and real-time analysis and reporting. Its capabilities are the primary topic of this section.

The optional Deep Discovery Advisor provides open, scalable, custom sandbox analysis that can augment the protection capabilities of an organization's existing security products, such as email and Web gateways. It also provides visibility to network-wide security events and security update exports—all in a unified intelligence platform. Deep Discovery Advisor is the gateway to the full power of the Trend Micro Custom Defense solution, described in the next section.

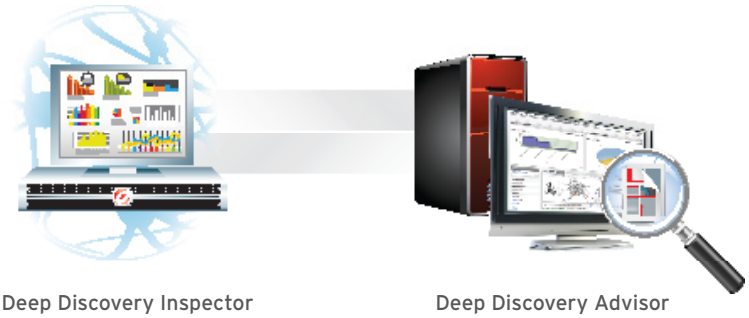

Deep Discovery Inspector

- Network traffic inspection
- Custom threat detection
- Real-time analysis & reporting

Deep Discovery Advisor

- Open custom sandboxing platform
- Deep investigation & analysis
- Adaptive Protection against attack

HOW DEEP DISCOVERY INSPECTOR WORKS

	ATTACK DETECTION	DETECTION METHODS
Malicious Content	• Emails containing embedded document exploits • Drive-by-downloads • Zero-day and known malware	• Decode and decompress embedded files • Sandbox simulation of suspicious files • Browser exploit kit detection • Malware scan (Signature and Heuristic)
Suspect Communication	• Command-and-control communication for all malware: bots, downloaders, data stealing, worms, and blended threats • Backdoor activity by attacker	• Destination analysis (URL, IP, domain, email, IRC channel, etc.) via dynamic blacklisting, white listing • Smart Protection Network™ Web reputation • Communication fingerprinting rules
Attack Behavior	• Malware activity: propagation, downloading, spamming, etc. • Attacker activity: scan, brute force, service exploitation • Data exfiltration	• Rule-based heuristic analysis • Identification and analysis of usage of 100's of protocols and applications including HTTP-based apps

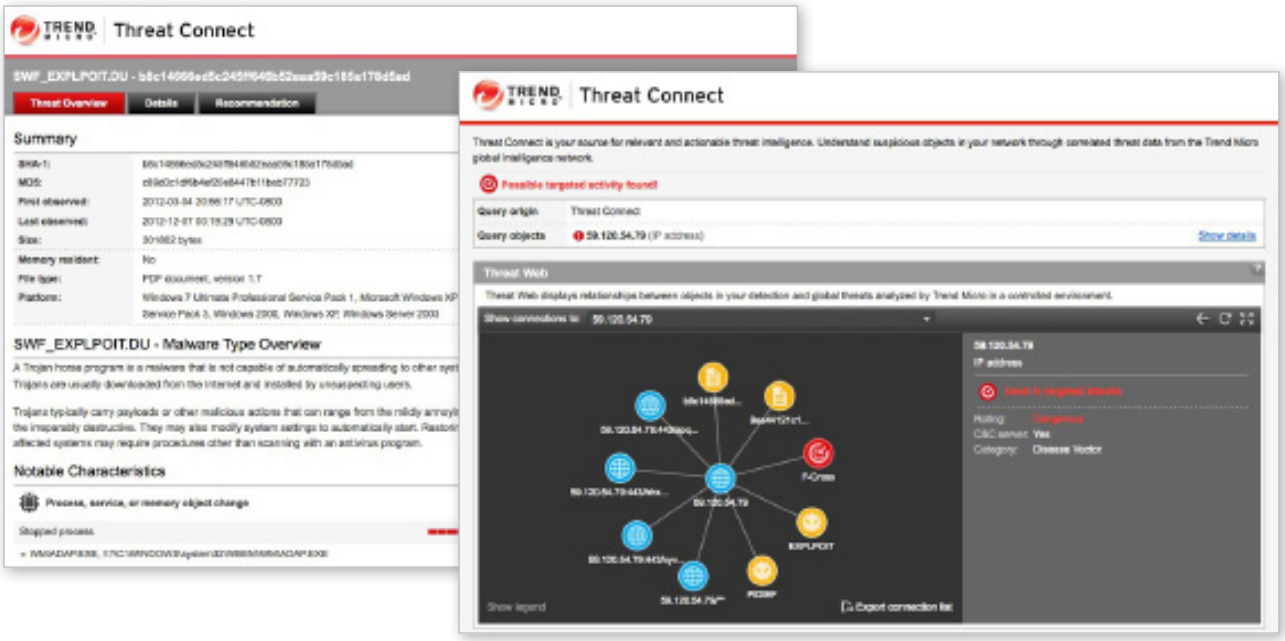

THREAT PROFILE
What are the characteristics, origins and variants of this malware?

RELATED IPS/DOMAINS
What are the known C&C comms for this attack?

ATTACK GROUP/CAMPAIGN
Who and what is behind this threat?

CONTAINMENT AND REMEDIATION
What to look for, how to remediate and eradicate?

Deep Discovery Inspector is purpose-built for detecting APT and targeted attacks—identifying malicious content, communications, and behavior that may indicate advanced malware or attacker activity across every stage of the attack sequence.

Deep Discovery Inspector uses a three-level methodology to perform initial detection, simulation, correlation, and, ultimately, a final cross-correlation to discover targeted attacks discernible only over an extended period of time. Specialized detection and correlation engines provide the most accurate and up-to-date detection aided by global threat intelligence from the Trend Micro™ Smart Protection Network™.

Dedicated threat researchers continually update the detection rules that correlate events and define the behavior and communication fingerprinting that detects targeted attacks. The result is a high detection rate with low false positives and in-depth incident reports that help speed up containment of an attack.

CORE TECHNOLOGIES

Network Content Inspection Engines

Deep packet inspection engines perform protocol detection, decoding, decompression, and file extraction across all ports and hundreds of protocols

Advanced Threat Scan Engines

Malware file scanning combines with aggressive heuristic scanning techniques to detect both known and unknown malware and document exploits

Custom Sandbox Analysis

A virtualized threat sandbox analysis system uses customer-specific images that exactly match the target environments—ensuring accurate detection and reducing false positives

Threat Detection and Correlation Rules

Behavior and communication techniques used by attackers are detected based on identification and correlation rules developed and continually updated by Trend Micro's 1,200 threat researchers and extensive Smart Protection Network intelligence

Trend Micro Smart Protection Network

The first and most extensive global threat intelligence and reputation service processes over 16+ billion requests daily ensures Deep Discovery detects any threat variant seen worldwide

Threat Connect Intelligence Portal

Full breadth of relevant threat intelligence about a customer's targeted attack that includes malware characteristics, origins and variants, related C&C IPs, attacker profile, and suggested remediation procedures

Going Beyond Detection: Enabling a Complete Custom Defense

An ideal solution against APTs should not only perform custom detection and analysis of attacks at the network level but should also integrate advanced detection technology into an organization's existing endpoint, messaging, and gateway defenses to strengthen the protection levels of current security investments. Detection of an attack at any one protection point would automatically update other protection points to defend against further attack—all working in a multi-vendor security environment. An ideal solution should leverage the global intelligence of a major security vendor to aid in detection, and use it to provide threat profile information relevant to an organization's particular attack. Finally, an ideal solution should pair this profile with network-wide event analysis to guide rapid containment and remediation.

In short, an ideal solution against APTs goes beyond detection and provides a complete custom defense employing a comprehensive Detect–Analyze–Adapt–Respond lifecycle unique to each, specific organization and the threats against it.

Trend Micro believes the attributes of a custom defense strategy make it the best choice to combat targeted attacks—and that belief is being put into action by delivering a complete Trend Micro Custom Defense—with Trend Micro Deep Discovery serving as the foundation. The TrendMicro Custom Defense weaves an organization's entire security infrastructure into a tailored and adaptable defense that is tuned to an organization's particular environment and particular attackers. Using custom sandbox analysis, custom intelligence and custom security updates, Trend Micro Custom Defense enables an organization not only to detect and analyze APTs and targeted attacks, but also to rapidly adapt its protection and response to these attacks.

THE TREND MICRO CUSTOM DEFENSE SOLUTION—HOW IT WORKS

DETECT: *what standard defenses can't*

In the previous section you read about Trend Micro Deep Discovery, which provides advanced threat protection that performs network-wide monitoring to detect zero-day malware, malicious communications and attacker behaviors that are invisible to standard security defenses. Deep Discovery sandbox simulation is also integrated with other Trend Micro products including Messaging Security products, giving them the power to block the spear phishing and social engineering exploits commonly used by attackers in the initial phase of a targeted attack. And, Deep Discovery supports an open Web Services interface so that any security product can integrate with the custom sandbox detection.

ANALYZE: *using real-time global and local intelligence*

Upon detection, Deep Discovery analytics and attack-relevant intelligence from the Smart Protection Network and Threat Connect portal create a rich threat profile that enables an organization to gain an in-depth understanding of the risk, origin and characteristics of the attack that help prioritize and guide containment and remediation plans. The depth of these threat profiles also enables the adaptive protection capability of the Trend Micro Custom Defense solution.

ADAPT: *security protection points to block the new threat*

To immediately adapt and strengthen protection against further attacks, the Trend Micro Custom Defense solution uses in-depth threat profiles to update the Smart Protection Network and to generate automated, custom security updates (IP/Domain blacklists and security signatures) to existing Trend Micro products in a customer's environment, including endpoint, gateway, and server enforcement points. Built using an open and extensible platform, the solution can also export security updates to non-Trend Micro security products that may already be an important part of the organization's defense in-depth strategy.

RESPOND: *with rapid containment and remediation*

Finally, the Trend Micro Custom Defense solution delivers 360-degree contextual visibility of the attack by combining the rich threat profile with results from employing specialized attack response tools and intelligence gathered from network-wide security event collection and analysis. Alternatively, the threat profile and other findings can be shared with a SIEM system already in place. Armed with this information organizations gain the insight needed to expedite the containment and remediation process and to contact authorities, as may be appropriate.

Conclusion

Targeted attacks are successfully bypassing traditional security defenses, and the majority of IT professionals now believe their organizations have been targeted. According to an Information Week Security article by Mathew Schwartz, "[APTs take a] low-and-slow approach that's difficult to detect, but which has a high likelihood of success. Attackers only need to trick a single employee into opening a piece of malware that exploits a zero-day vulnerability, thus giving them access to not just the employee's PC, but potentially the entire corporate network."

As seen with the review of successful APT attacks, these threats are highly sophisticated and take a highly specialized and customized approach to gain access to their targets. These are next-generation attacks that require a next-generation security approach to close the gaps these threats exploit. While each attack is tailored to its target, they consistently follow key lifecycle phases:

- intelligence gathering
- entry point
- command and control communication
- lateral movement
- asset/data discovery
- data exfiltration.

A strong defense against APTs must have in-depth detection and analysis capabilities across all phases of the attack lifecycle.

As organizations plan their IT security projects for 2013, it is critical to include the defense strategy against APTs as part of the project scope. Trend Micro Deep Discovery should be considered as a key solution to defend against targeted attacks. Deep Discovery uniquely detects and identifies evasive threats in real-time and provides the in-depth analysis and relevant, actionable intelligence specific to each organization's environment.

Have Sales Contact Me >>

TREND MICRO INC.
U.S. toll free: +1 800.228.5651
phone: +1 408.257.1500
fax: +1408.257.2003