Traps Advanced Endpoint Protection

Technology Overview

March 2015

Dear Reader,

Just three weeks before sitting down to write this letter, I was the chief information security officer for a large multi-national corporation that handles corporate and customer data across twenty countries and 400 branch offices.

If you're like I was, your company, customers, investors, and the public rely upon you to keep that information secure. You undoubtedly have multiple layers of security in place to ensure your data is protected. The problem is that the existing technologies have not been adequate. Our endpoints still become infected. We spend much of our time trying to keep up with patches, and then detecting, remediating, and often re-imaging systems when we should be preventing these breaches.

This whitepaper will introduce you to a better way. Traps Advanced Endpoint Protection is a product like no other I've seen. Over the years, I've had just about every vendor come in and demonstrate its anti-malware, APT prevention, ETDR, and cybersecurity tools. Traps was the only product I saw that actually prevents exploits that have been bypassing all of my other controls — even Zero Day. I was so excited about this that I decided to join Palo Alto Networks to help spread the word to my industry peers.

Traps not only describes a new product, but an entirely new product category that is going to change the way we think about endpoint protection. Read on and I'll tell you all about how it works.

Sincerely,

Sebastian E. Goodwin
Director, Traps Advanced Endpoint Protection

Introduction

Cyberattacks are attacks performed on networks or endpoints in order to inflict damage, steal information, or achieve other goals that involve taking control of computer systems that belong to others. Cyberattacks are perpetrated either by causing a user to unintentionally run a malicious executable, or by exploiting a weakness in a legitimate executable in order to run malicious code “behind the scenes” without the knowledge of the user.

Despite a plethora of endpoint security products on the market purporting to solve this problem, one can pick up any newspaper and realize that endpoints are still being infected at an alarming rate. A new approach was needed to protect endpoints from advanced threats.

Why Endpoints Still Get Infected

Traditional endpoint protection solutions use methods that simply cannot keep up with the rapidly evolving threat landscape. The most advanced threats these days leverage vulnerabilities in software that we use on a regular basis. They often come in the form of commonly used data files (e.g., pdf, rtf, doc, ppt, xls), or they can be individually crafted to target proprietary software used in various industries. These files open just fine in their native applications and display content that looks normal, but there is malicious code embedded in the file. This code exploits a vulnerability in the native application causing the attacker’s code to run. All of this can happen as your endpoint anti-malware suite stands idly by looking for a bad executable that it’s seen before, or some other indication that something bad is happening. The problem is, there may not be any indication that something bad has happened. Until now, patching was the only way to ensure protection from known vulnerabilities, and there was no reliable method to protect systems from unknown vulnerabilities.

With vulnerabilities existing long before patches are released, it makes for an inevitable delay in installing patches, and increases the security risk.

Stop Detecting, Start Preventing

There was a time in the not-too-distant history of our profession where we as information security professionals decided that it was virtually impossible to prevent these kinds of advanced threats on the endpoint. We focused our attention on detection and response. What we’ve learned since then is that we are detecting malware that has been stealing our data for months, even years. Most of us are not willing to stand behind this strategy. We require an approach that can rebuild lost confidence in endpoint security, one that is built on the principle that attacks on the endpoint can not only be detected, but actually prevented. We call this Advanced Endpoint Protection.

Advanced Endpoint Protection should deliver on the following:

1. It must be able to prevent all exploits, including those utilizing unknown zero-day vulnerabilities.

2. It must be able to prevent all malicious executables, without requiring any prior knowledge.

3. It must provide detailed forensics against prevented attacks to strengthen all areas of the organization by pinpointing the target and techniques used.

4. It must be highly scalable and lightweight to seamlessly integrate into existing operations with minimal to no disruption.

5. It must integrate closely with network and cloud security for quick data exchange and cross-organization protection.

How Traps Works

Traps™ is an advanced endpoint protection solution that prevents advanced attacks originating from either exploits or malicious executables before any malicious activity can successfully run, regardless of software patches in place.

If an attack attempt is made, Traps will immediately block the technique or techniques, terminate the process, and notify both the user and the administrator that an attack was thwarted. Whenever a block does occur, Traps will collect detailed forensics, including the offending process, the memory state when it was prevented, and many other details that are reported to the Endpoint Security Manager.

Multiple Attack Types, Total Protection

Attacks come in different forms and can arrive via multiple vectors including Web, email, and external storage. Most traditional endpoint security products protect endpoints from malicious executable files, which are the least sophisticated form. The most advanced and targeted attacks arrive in the form of seemingly harmless data files that are opened by legitimate applications. For example, malicious code can be implanted in a Microsoft Word or PDF document. Once the file is opened, the malicious code takes advantage of a vulnerability in the legitimate application being used to view the file, allowing it to execute code and take full control of the endpoint. We call this an exploit.

Traps protects endpoints by preventing malware in the form of executables, and exploits in the form of data files or network-based attacks.

Preventing Exploits

Many advanced threats work by placing malicious code in a seemingly innocuous data file. When the file is opened, the malicious code leverages a vulnerability in the native application used to view the file, and the code executes. Because the application being exploited is allowed by IT security policy, this type of attack will bypass whitelisting controls. What sets Traps apart is the fact that it focuses on the core techniques used by all exploits. It turns out that, although there are many thousands of exploits, they all rely on a small set of core techniques that change infrequently. Furthermore, each exploit needs to use a series of those techniques in order to be successful. Traps renders these techniques completely ineffective, which means the application is no longer vulnerable.

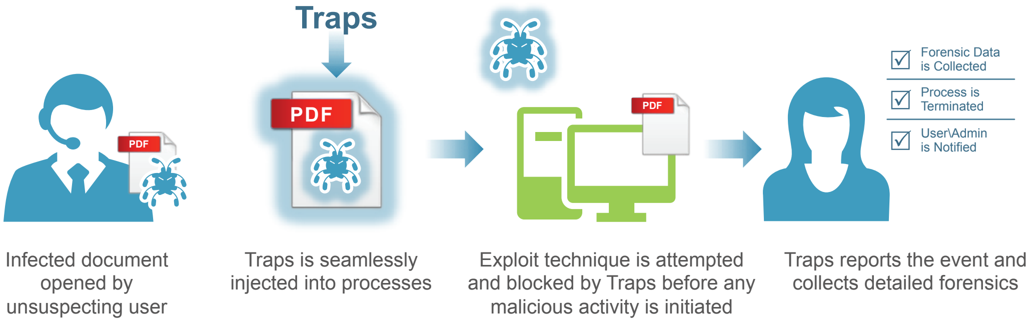
Traps
PDF
PDF
PDF
☑ Forensic Data is Collected
☑ Process is Terminated
☑ User\Admin is Notified
Infected document opened by unsuspecting user
Traps is seamlessly injected into processes
Exploit technique is attempted and blocked by Traps before any malicious activity is initiated
Traps reports the event and collects detailed forensics

Figure 1: Exploit prevention — user experience.

The Traps agent injects itself into each process as it is started. When an attacker attempts to exploit a software vulnerability, the Traps protection modules cause the exploit attempt to fail because Traps has already made the process impervious to those techniques. When the attempt is prevented, the Traps agent kills the process and reports all of the details to the Endpoint Security Manager (ESM).

By default, Traps policy is configured to protect over 100 processes — each one with dozens of proprietary EPMs. But unlike other products, Traps is not limited to protecting only those processes or applications. Our customers use Traps to protect all manner of processes and applications by simply adding them to the policy configuration. This is especially useful for those customers running industry-specific applications. Traps can protect point-of-sale (POS) systems, ATM machines, SCADA, and any other application from exploitation.

If for some reason an application conflicts with one of the EPMs, this is not a problem. Simply disable that EPM for the specific application and computer. The application is still protected by dozens of other EPMs. Since exploits rely on a chain of techniques to successfully run, the other EPMs will continue protecting that application and will block at least one of the techniques, thus breaking the chain.

Examples of attacks that the EPMs can prevent include:

- DLL hijacking — replacing a legitimate DLL with a malicious one of the same name
- Hijacking program control flow
- Inserting malicious code as an exception handler

Preventing Malicious Executables

In addition to preventing exploits hiding in data files or launched over the network, Traps employs a comprehensive approach to the prevention of malicious executables. Malicious executables, also known as malware, can be inadvertently downloaded and run by users without their knowledge. Traps malware prevention engine uses a combination of policy-based restrictions, WildFire™ analysis, and malware protection modules to prevent the execution of malicious executables. When combined, these methods offer unparalleled malware prevention. The process works as follows:

Policy-based Restrictions: Policy-based restrictions dramatically reduce the attack surface by preventing execution in high-risk scenarios. For example, you may want to prevent the execution of a particular file type directly from a USB drive or of files in directories where applications should not reside.

Advanced Execution Control: Prevention can only be achieved with a multilayered approach to continuously reduce the attack surface at every stage of the attack. Traps provides granular execution control capabilities that significantly mitigate the possibility of malware execution. These controls include:

1. Execution Restrictions – Traps robust restrictions provide granular control of policies related to local and network folder locations, child processes, external media, and unsigned executables. These restrictions increase the business flexibility while minimizing the security risk.

2. Granular System Hardening – For relatively static or special purpose systems, execution control policies can be used to statically define which executables are allowed or prohibited. These are specified by file hash. While WildFire integration can be used to analyze files dynamically, these policy settings can be used to override the WildFire verdict locally.

WildFire Executable Analysis and Prevention: WildFire integration provides the ability to have both the security of granular execution control and the manageability of a dynamic security policy driven by automated analysis of unknown executables. If an executable file has never been seen before on the endpoint, Traps can submit the file hash for immediate identification by WildFire. If WildFire identifies the file as malicious, Traps will prevent execution before any damage is done. With nearly two million samples analyzed each day, there is a good chance WildFire has seen the file and can alert Traps if it is malicious. If the file has not been seen by WildFire, it can be automatically uploaded for rapid analysis in order to determine if it is malicious. Since both Traps and the next-generation firewalls can submit files to WildFire, this integration allows for seamless sharing of threat intelligence between the next-generation firewall and the endpoints.

Malware Prevention Modules: If the file is allowed to execute, malicious activity can still be blocked by a Malware Prevention Module (MPM). Like the EPMs, MPMs focus on core techniques leveraged by many types of malware. For example, they will prevent malicious code from being injected into trusted applications.

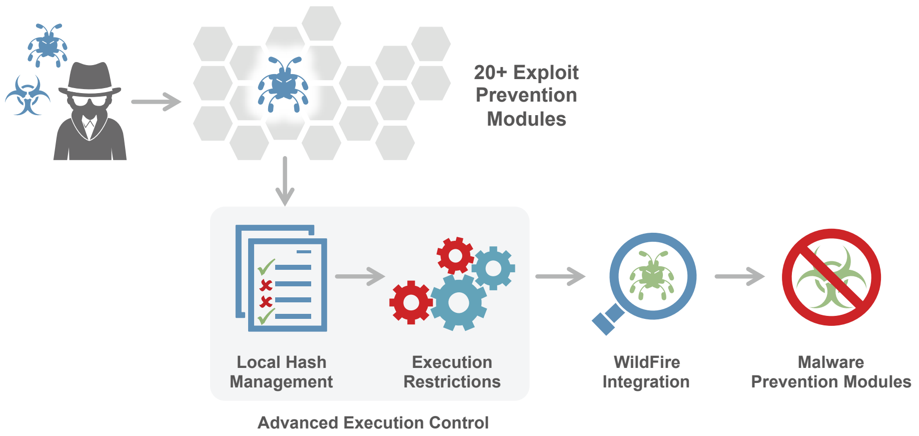
20+ Exploit Prevention Modules
Local Hash Management
Execution Restrictions
WildFire Integration
Malware Prevention Modules
Advanced Execution Control

Figure 2: The right way to prevent malicious executables

Forensic Data Gathering

Extensive data is gathered from the Traps agent. On an ongoing basis, the agent records and sends to the ESM server information about each process that is run. The agent will also alert if there are any attempts to stop, remove, or otherwise tamper with Traps. When an attack is prevented, further detail can be gathered from the endpoint, including a full memory capture and information about the activities attempted by the malicious code. After memory is captured, Traps will then perform a secondary analysis surrounding the nature of the event and search the memory for any traces of malicious activity.

It is important to recognize that that scope of forensic information available after an attack has been prevented is unavoidably less than the information available about an attack that has succeeded and done damage.

Benefits of the Traps Approach

Coverage for zero-day vulnerabilities and unknown malware: Rather than waiting for signatures or indicators of compromise to be released, you remain protected from the newest, most advanced threats.

Install patches on your own schedule: Vulnerabilities exist long before patches are released, and deploying patches can be a lengthy and cumbersome process. IT teams struggle to ensure patches are thoroughly tested and deployed to all endpoints within a reasonable timeframe. Furthermore, nearly every organization has those legacy systems that, for one reason or another, cannot be patched. With advanced endpoint protection, endpoints are protected regardless of patch levels.

Protect any application from exploits: The focus on exploit techniques rather than application-specific characteristics means this protection can be extended to any application. While many endpoint security products protect only a few commonly used applications from exploitation, our approach is used by customers to protect hundreds of proprietary applications.

Minimal performance impact: This approach does not rely on system scanning, virtualization, or any other bloated technology. The agent is lightweight and nonintrusive. It can be completely invisible to the end user.

Saves time and money: When you prevent attacks, you no longer have to deal with the costs of remediation and end-user downtime that often result from malware infection, especially when systems have to be re-imaged.

Ease of management, no frequent updates: One of the problems with traditional endpoint-protection products is the need to constantly deploy signature updates. Traps focuses on a small set of techniques that do not require frequent updates.

Threat intelligence through WildFire integration: With WildFire-enabled customers benefit from the threat-intelligence ecosystem with over one million samples submitted daily by the community. Automated upload and analysis of unknown executables ensures that every new executable launched on an endpoint can be analyzed.

Traps Deployment Architecture

Endpoint Security Manager Console

The Traps infrastructure supports various architectural options to allow for scalability to large distributed environment. Installation of the ESM creates a database on a Microsoft SQL server and installs the administrative console within IIS. Microsoft SQL 2008 and 2012 are supported and the SQL server may be dedicated to ESM or a database can be created on an existing SQL server.

The Endpoint Server can be installed on Windows Server 2008 R2, Windows Server 2012, or Windows Server 2012 R2 on physical or virtual machines.

Endpoint Security Manager Servers

ESM servers essentially act as proxies between Traps agents and the ESM database. Communications from Traps agents to ESM servers occur over HTTPS. ESM servers do not store data and, therefore, can be easily added and removed from the environment as needed to ensure adequate geographic coverage and redundancy.

To ensure global connectivity, customers who do not use a mobility solution like Palo Alto Networks® GlobalProtect™ may opt to put an ESM server in the DMZ or in a cloud-based environment with external connectivity. ESM servers can be installed on Windows Server 2008 R2, Windows Server 2012, or Windows Server 2012 R2 physical or virtual machines.

Traps Agent

The Traps agent installer is a ~9 MB MSI package that can be deployed using your software deployment tool of choice. Subsequent updates to the agent can be deployed via the ESM. The agent consumes less than 25 MB on disk and less than 40 MB while running in memory. Observed CPU utilization is less than 0.1 percent. The agent also employs various tamper-proofing methods that prevent users and malicious code from disabling protection or tampering with agent configuration.

The lightweight structure allows for the Traps environment to scale horizontally and support large deployments of up to 50,000 agents per ESM, while still maintaining a centralized configuration and database for policies. Traps can co-exist with most major endpoint security solutions, and the CPU utilization and I/O remains incredibly low. With such minimal disruption, this makes Traps optimal for critical infrastructures, specialized systems, and VDI environments.

Traps currently supports the following Windows-based operating systems:

OPERATING SYSTEM:

- Windows XP (32-bit, SP3 or later)
- Windows 7 (32-bit, 64-bit, RTM and SP1; all editions except Home)
- Windows 8 (32-bit, 64-bit)
- Windows 8.1 (32-bit, 64-bit)
- Windows Server 2003 (32-bit, SP2 or later)
- Windows Server 2003 R2 (32-bit, SP2 or later)
- Windows Server 2008 (32-bit, 64-bit)
- Windows Server 2012 (all editions)
- Windows Server 2012 R2 (all editions)
- Windows Vista (32-bit, 64-bit, and SP2)

VIRTUAL ENVIRONMENTS:

- VDI
- Citrix
- VM
- ESX
- VirtualBox/Parallels

PHYSICAL PLATFORMS:

- SCADA
- Windows Tablets

External Logging

The ESM can write logs to an external logging platform, such as SIEM, SOC or syslog, in addition to storing its logs internally. For an organization that deploys multiple ESMs, an external logging platform allows for an aggregated view of log databases.

Integrated Security Platform

In 2005 we set out to deliver a next-generation security platform designed from the ground up to reduce an organization's attack surface, and prevent the most sophisticated cyberattacks from achieving their objectives. To accomplish this, it was clear that a disruptive new approach was required, one that increases a customer's ability to see all network traffic; establish positive controls for applications and users; and prevent all attacks through a tightly integrated system that links network, cloud and endpoint security into a single common architecture.

We have delivered on that strategy and today offer a complete Enterprise Security Platform that is comprised of three core elements: a Next-Generation Firewall, Threat Intelligence Cloud, and Advanced Endpoint Protection. This natively integrated platform eliminates the need for point tools and other disjointed technologies, streamlines day-to-day operations, and significantly boosts an organization's security efficacy through a multi-layered model that prevents threats at each stage of the attack kill chain.

Our Enterprise Security Platform protects every corner of your organization — from your mobile workers to the core of your cloud-enabled data center. Automation within the platform eliminates the need for expensive, manual processes, and improves an organization's ability to quickly respond to new global threats.

paloalto networks®
the enterprise security company™

4401 Great America Parkway
Santa Clara, CA 95054

Main: +1.408.753.4000
Sales: +1.866.320.4788
Support: +1.866.898.9087

www.paloaltonetworks.com