THREAT INTELLIGENCE CLOUD

Leveraging the Global Threat Community to Prevent Known and Unknown Threats

Palo Alto Networks
4401 Great America Parkway
Santa Clara, CA 95054

www.paloaltonetworks.com

paloalto NETWORKS®

Executive Summary

No organization today is immune to cybercrime. Cyber criminals are ramping up activity across the globe and utilizing new methods to evade traditional security measures. An effective network security solution must not only prevent known threats from entering and damaging the network, but also leverage global threat intelligence to protect the network from unknown threats. Traditional detection-focused solutions typically focus on a single threat vector across a specific section of the network, leaving multiple areas vulnerable to attack. In addition, these legacy solutions are made up of a "patchwork"

of point products that make it very difficult to coordinate and share intelligence among the various devices. By employing the Palo Alto Networks® Threat Intelligence Cloud, businesses can reduce their attack surface, block all known threats, and leverage the global threat community to detect unknown threats and convert them into known, stoppable threats.

Keeping Up with Advanced Threats: A Futile Effort?

The speed at which new security threats appear and the growing sophistication of hackers' techniques make fighting cybercrime a constant challenge. In 2014, 71 percent of the security professionals polled in the Cyberthreat Defense Report said their networks were breached, up significantly from the previous year (62 percent). And a majority (52 percent) of respondents felt that a successful cyberattack against their network was likely in the next 12 months, compared to just 39 percent in 2013[1]. To say that keeping up with attackers' evolving techniques and advanced threats is difficult is an understatement. New attack methods and malware pop up all the time, each more evasive than the last, and data breach headlines are in the news so often that they're becoming commonplace. The cat-and-mouse game between attackers and defending organizations is no longer a competition — the attackers have not only pulled ahead, but they've gained so much distance that most security teams have given up the fight and are now focused on detection and remediation, abandoning traditional "preventive" approaches because they have failed.

CYBERATTACKS ON THE RISE

In 2014, 71 percent of respondents' networks were breached with 22 percent of them victimized six or more times. Phishing/spear-phishing, malware, and zero-day attacks are perceived as posing the greatest risk to organizations.

— 2015 Cyberthreat Defense Report, which polled more than 800 IT security decision makers representing 19 different industries across North American and Europe

The problem with traditional detection-focused solutions is that they hone in on one, specific section of the network, leaving multiple areas exposed and vulnerable to attack. As a result, security teams are faced with the time-consuming process of piecing together logs from different devices, combing through them to discover unknown threats, and then manually creating and deploying protections. By the time this happens — oftentimes days or weeks later — it's too late because minutes or hours are all an attacker needs to accomplish his or her end goal.

Global Community Fighting Cybercrime

Palo Alto Networks has developed an innovative solution to this problem. Instead of relying on the tedious, manual efforts of detection-and-remediation approaches, Palo Alto Networks Threat Intelligence Cloud proactively discovers unknown threats and automatically turns them into known, preventable quantities so that organizations around the world can be made aware of newly discovered threats in as close to real-time as possible. By combining threat intelligence from thousands of organizations around the globe, the Threat Intelligence Cloud provides automatic prevention and data correlation that is very powerful in minimizing threats and breaches.

[1] The CyberEdge Group 2015 Cyberthreat Defense Report, March 11, 2015.

PALO ALTO NETWORKS THREAT INTELLIGENCE CLOUD

192,000 Anti-Malware Protections per Day

24,000 URL Protections per Day

13,500 DNS Protections per Day

Data Sourced from **6,100+** Global WildFire Customers

Protections Delivered Automatically in **15 Minutes**

As of 4/7/15

In order to protect against today's far-reaching cyberthreats, we must leverage the global community and combine threat intelligence from a variety of sources to help "connect the dots" among various cyberthreats. Real-time, global intelligence feeds help security teams keep pace with threat actors and easily identify new security events. Because the Threat Intelligence Cloud collects meticulous data around every threat, the stack of forensic work involved in incident response — sifting through multiple logs, searching for correlated events, tracking down users and compromised hosts — is significantly reduced. This data can be fed into practically any SIEM tool by using the open API integration to further support security operations and incident response.

Threat Prevention Is Automatic

Beyond detection and alerts, new threats — both known and unknown — are automatically discovered and prevented. Unknown files are analyzed and, if a new threat is identified, protections are automatically developed and delivered to customers in the form of an update. The Threat Intelligence Cloud virtually eliminates the time spent responding to successful attacks by providing easy access to correlated data using a few mouse clicks.

For example, a new piece of malware only has to be seen once by any single customer before protections are automatically created and applied within 15 minutes to the arsenal of prevention techniques employed by the Palo Alto Networks Enterprise Security Platform[2].

What Is the Threat Intelligence Cloud?

The concept behind the Threat Intelligence Cloud is to leverage threat intelligence collected from thousands of globally deployed security devices in order to automatically discover unknown threats, and deliver protections back to the community.

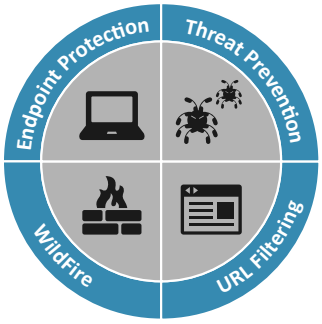

Threat intelligence is analyzed both manually and automatically. The Palo Alto Networks threat research team, Unit 42, analyzes the accumulated data to determine how attackers are evolving their methods. Most importantly, this data is analyzed automatically and can rapidly produce protections against those previously unknown advanced persistent threats — malware, command-and-control (CnC) servers, and malicious DNS entries and URLs — to block them before they gain a foothold within the networks of the global customer base.

Palo Alto Networks Threat Intelligence Cloud is composed of four distinct subscription-based product components that are all purpose-built to work closely with each other:

1. Threat Prevention incorporates intrusion prevention, network anti-malware, and anti-command-and-control (CnC) features, which scan all traffic for known threats hiding within incoming and outgoing traffic and provide a layered defense against known attacks when combined with URL Filtering.

2. URL Filtering categorizes websites by content type, including malicious activity, and prevents harmful Web pages from administering malware by blocking them. It stops users from inadvertently navigating to malicious URLs, exploited Web pages, and watering holes where legitimate sites become compromised. In combination with Palo Alto Networks App-ID™ functionality, URL filtering is an invaluable tool to secure Web traffic.

[2] Palo Alto Networks Enterprise Security Platform https://www.paloaltonetworks.com/products/platforms.html

3. WildFire™ provides protection against advanced malware and threats. WildFire analyzes files, URLs and DNS requests from thousands of customers in every industry across the globe, and then generates content-based protections, which are delivered back to the Threat Prevention and URL Filtering profiles deployed by every device within the global customer base.

4. Traps™ Advanced Endpoint Protection prevents advanced attacks on the endpoint, originating from either exploits or malicious executables. By focusing on a core set of exploit and malware techniques common to all attacks, Traps prevents the attack regardless of patches in place, before any damage can be done.

How Does the Threat Intelligence Cloud Work?

Designed to work in tandem, the four components of the Threat Intelligence Cloud create a closed loop of continuous detection and prevention that protects customers' networks from known and unknown threats at every stage in the Attack Kill Chain, whether it be delivery, exploitation or command-and-control.

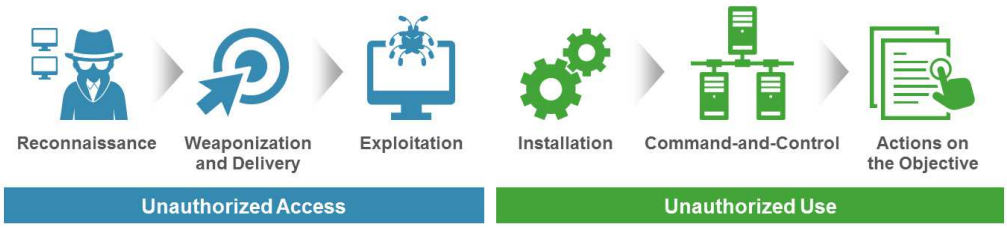

Figure 2: The Palo Alto Networks Attack Kill Chain Model

Prevent Known Threats

Much of today's malware is designed to be completely undetectable. Palo Alto Networks prevents known threats by first reducing the network's exposure to unnecessary risk, and then fully inspecting all allowed traffic for threats.

URL Filtering categorizes URLs by their content, with special categories denoting sites delivering malware, and allows users to block websites in specific categories or downloaded files from websites in specific categories. This cuts off the delivery method for Web-based attacks, protecting users from inadvertently infecting their devices and, by extension, the network. PAN-DB, the database that drives URL filtering, is constantly updated by WildFire analyses as new malware sites and compromised pages are seen.

Anti-malware signatures within the threat prevention suite prevent known malware from entering the network by comparing byte sequences found within all traffic against payload-based signatures. When a signature match occurs, the packet is dropped and the connection reset, ensuring that the threat is never delivered. The anti-malware signatures are fed by WildFire analysis and comprise protections against previously analyzed malware samples and their variants.

For all allowed traffic, including permitted Web traffic, the threat prevention suite uses CnC (Command and Control) signatures to stop requests to known malicious DNS entries and IP addresses, blocking any outbound communication channels to attacker-controlled servers at the DNS resolver before they're established. This effectively cuts off the CnC channel with the attacker and obstructs his or her ability to control infected devices or extract data. All unknown files and Web content are sent to WildFire for analysis. Threat prevention signature libraries receive new anti-malware and anti-CnC protections, created and dispatched by WildFire, and apply them immediately to prevent advanced threats across the entire network — not just at the perimeter.

Automatically Detect and Prevent Unknown Threats

The cloud-based WildFire service simplifies an organization's response to the most dangerous threats — those that are unknown — by detecting new malware and quickly implementing preventive measures automatically within threat prevention and URL filtering tools.

WildFire uses static and dynamic analysis — looking at both the code itself and how it instructs the file to behave — to examine data within the file header and body across multiple operating system environments. It specifically investigates runtime behaviors for indicators of malicious activity or intent, like changes to the registry or code that detects security-scanning software. Individual behaviors and combinations of behaviors are both analyzed for harmful outcomes. All network activity is analyzed, including backdoor creation, visitation to low-reputation domains, and network reconnaissance to determine where the file came from and with whom it's programmed to establish outbound communication channels (CnC activity).

REAL-TIME THREAT INFORMATION

"WildFire gives us great, actionable information to show management, plus it runs right on our Palo Alto Networks boxes.

You benefit from the collective intelligence of a huge global user base. It doesn't matter if you're a 50-person dry cleaning company or a Fortune 100 firm, you get to share in the collective wisdom and experience from thousands of companies to help you fight advanced threats. WildFire is a top notch tool that gives us incredible functionality without much more cost."

— Blake Wofford,
Senior Security Engineer,
Exeter Finance Corp

When a file is deemed malicious, session data associated with the delivery, the original malware sample, and Packet Captures (PCAPs) of the dynamic analysis session are logged. An anti-malware signature is automatically generated based on the file payload — not based on hash, which is easily mutated. The signature is designed to match byte sequences within the file body to protect against both the original sample's hash as well as any mutation, so a single new signature covers multiple variations of the file. In addition, Domain Generation Algorithms within the payload are examined, from which DNS signatures are created. The new anti-malware and CnC protections are then incorporated into the threat prevention signature library within 15 minutes of the original submission to WildFire. Protections are also derived from websites from which the malware was delivered by designating the site to one of the "malicious" URL categories in PAN-DB within 30 minutes. The previously unknown threat — the malware, its CnC channels, and its delivery source — becomes known and is preventable almost immediately after it is first encountered.

Extend Zero Trust to the Endpoint

Zero Trust[3] is intended to promote a "never trust, always verify" policy as its guiding principle. With exploit kits readily available to attackers, even "good" applications can go "bad," so the same rigor must be applied on the endpoint, on the OS, on connected devices, and in memory. This is particularly important, as most resources an attacker might be interested in — data and applications — will live on the endpoint.

Until new threat behaviors are discovered and protections against them are created, the endpoint remains vulnerable. The challenge organizations face is how to prevent "patient-zero" infections: stopping zero-day attacks from compromising their endpoints, even when those threats are quickly discovered by network detection solutions.

The Traps Advanced Endpoint Protection agent injects itself into each process as it is started. When an attacker attempts to exploit a software vulnerability, the exploit prevention modules obstruct individual exploit techniques by making the process impervious to those techniques, regardless of the order in which they are carried out. 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. Policy-based controls allow administrators to further reduce the attack surface by controlling execution

[3] Zero Trust Model of information security, Forrester Research

and files. Highly granular restrictions are also available to define trusted processes or file types, locations and registry paths these processes can read from and write to.

When Traps encounters a new executable file that's never been seen by any endpoint within the network, it sends the file to WildFire before allowing it to execute. WildFire checks the file against its library of known malware samples and replies with a verdict as to whether the file has been identified as clean, malicious, or truly an unknown sample. 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. Based on that verdict and customized policies, Traps either cuts off any file execution attempt or allows it to run, protecting endpoints from being compromised by unknown malware. 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.

Leverage Global Threat Intelligence

Protections against unknown threats are extremely valuable to more than just the organization that first detected them. Those protections are also shared with every device deployed within the customer community, so customers around the globe can begin preventing attacks on any network protected by Palo Alto Networks that may be a target.

Figure 3: Threat Intelligence Collected Globally; Protections Sent Back to Community

Intelligence gleaned from malicious file samples is shared with all devices within the customer network. Data that may be related to the discovering organization is anonymized to keep customer privacy intact. For example, the Palo Alto Networks passive DNS network uses anonymized customer DNS queries and compares them to the global passive DNS dataset to predict future malicious domains based on patterns within components, like shared records and name servers, but customer-specific data like the source IP address is removed. The threat-relevant information is distributed to the customer network in the form of CnC protections and global context. It guarantees that the community is immediately protected against new malware as soon as any one customer sees it, and includes protection against both the file itself and any malicious DNS entries used as CnC communication channels.

The Threat Intelligence Cloud gathers an abundance of data around each identified threat. This data is then connected to other threats and vectors that share a likeness, contextualizing attacks on the individual enterprise network within the greater, global threat landscape. Integrated logs and analysis of WildFire submissions, Traps obstructions, signature matches, and third-party research are used by customers to easily recognize patterns so that most attacks are blocked, and host- and network-based indicators of compromise become immediately recognizable and actionable.

Beyond the Average Intelligence Cloud

To keep up with attackers, defensive tactics must evolve in parallel with new threats in both detection and prevention methods. By discovering unknown malware, and making it known and preventable within minutes, Palo Alto Networks Threat Intelligence Cloud does both and more:

- Content-based AV Protections: Because metadata from the header is used, Palo Alto Networks auto-generated signatures are based on payload — this is very different from signatures based on hash. Hashes are easily mutated and hash-based signatures bypassed. A single Palo Alto Networks AV signature covers multiple variations of the malware file — potentially thousands of mutations, including variants that haven't been created or discovered yet.
- Encrypted Data Is Analyzed: The Threat Intelligence Cloud encompasses threat data from all traffic, including anything sent with encryption. This is increasingly important, as nearly 35 percent of all enterprise traffic today is sent and received over SSL.
- Identifies Mobile Threats: The mobile device is a largely popular attack vector, with total infected devices in 2014 estimated at 16 million[4], and Android phones accounting for 50 percent of that total. Palo Alto Networks Threat Intelligence Cloud analysis includes APK files, mobile browsers, and links within text messages, and extends its layered protections to all mobile devices.
- All Data Remains Private: Submissions are secured by an encryption certificate that Palo Alto Networks signs on both sides, making sure all data remains safe and well-guarded.
- Professional Threat Analysis: Palo Alto Networks in-house threat research teams, including Unit 42, analyze data amassed by the Threat Intelligence Cloud to identify and investigate cutting-edge attack methods and malware, and report on unfolding trends within the black hat space.

The Threat Intelligence Cloud makes prevention automatic and detection of security events much simpler to identify and remediate, in part because most potential security events have already been prevented. Palo Alto Networks gathers threat data from customers around the globe and leverages that intelligence to build smarter protections, which are delivered back into the prevention tools deployed by the global customer base. In this way, the Threat Intelligence Cloud is a powerful tool that gives network defenders an advantage over advanced attacks.

For more information regarding the Palo Alto Networks Threat Intelligence Cloud solution and its component technologies, please visit **www.paloaltonetworks.com**.

[4] Alcatel-Lucent Motive Security Labs Malware Report H2 2014, February 12, 2015.

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