Next-Generation Firewall Overview

Recent changes in application behavior and usage patterns have steadily eroded the protection that the traditional firewall once provided. Users are accessing any application, from any location, often times to get their job done. Many of these applications use non-standard ports, hop ports, or use encryption to simplify and streamline user access and to bypass the firewall. Cybercriminals are taking full advantage of this unfettered application usage to spread a new class of highly targeted modern malware. The result is that the traditional firewall, with its reliance on port and protocol, can no longer identify and control the applications and threats traversing the network.

Attempts to regain control over application usage and protect digital assets for all users have produced duplicate local and remote security policies supported by an industry of stand-alone or sheet metal integrated firewall helpers. This approach introduces policy inconsistency and does not solve the visibility and control problems due to inaccurate or incomplete traffic classification, cumbersome management, and multiple, latency-inducing scanning processes. Restoring visibility and control requires a new, fresh, from-the-ground-up approach to safe application enablement that is only delivered by a next-generation firewall.

Key Next-Generation Firewall Requirements:

- Identify applications, not ports. Identify the application, irrespective of protocol, encryption, or evasive tactic and use the identity as the basis for all security policies.
- Identify users, not IP addresses. Employ user and group information from enterprise directories for visibility, policy creation, reporting, and forensic investigation—no matter where the user is located.
- Block threats in real-time. Protect against the entire lifecycle of an attack including dangerous applications, vulnerabilities, malware, high-risk URLs, and a wide array of malicious files and content.
- Simplify policy management. Safely and securely enable applications with easy-to-use graphical tools and a unified policy editor.
- Enable a logical perimeter. Secure all users, including traveling or telecommuting users, with consistent security that extends from the physical to the logical perimeter.
- Deliver multi-gigabit throughput. Combine purpose-built hardware and software to enable low-latency, multi-gigabit performance with all services enabled.

Palo Alto Networks next-generation firewalls enable unprecedented visibility and control of applications, users, and content using three unique identification technologies: App-ID™, User-ID, and Content-ID. These identification technologies, found in every Palo Alto Networks firewall, enable enterprises to safely and securely enable application usage, while significantly reducing total cost of ownership through device consolidation.

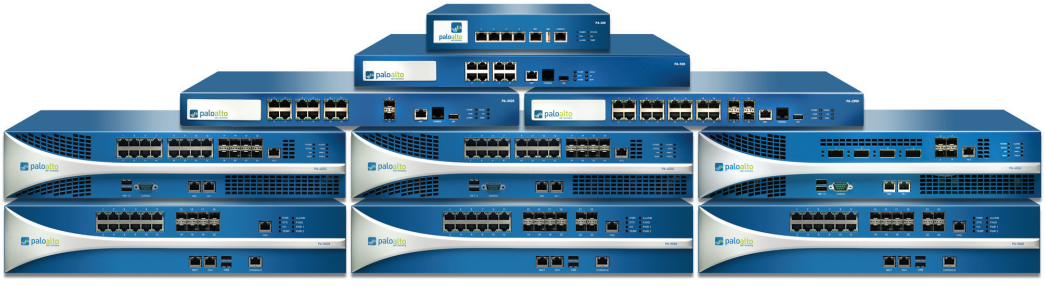

paloalto NETWORKS
the network security company™

App-ID: Classifying All Applications, All Ports, All the Time

Accurate traffic classification is the heart of any firewall, with the result becoming the basis of the security policy. Traditional firewalls classify traffic by port and protocol, which, at one point, was a satisfactory mechanism for securing the network. Today, applications can easily bypass a port-based firewall; hopping ports, using SSL and SSH, sneaking across port 80, or using non-standard ports. App-ID addresses the traffic classification visibility limitations that plague traditional firewalls by applying multiple classification mechanisms to the traffic stream, as soon as the firewall sees it, to determine the exact identity of applications traversing the network.

Unlike add-on offerings that rely solely on IPS-style signatures, implemented after port-based classification, every App-ID automatically uses up to four different traffic classification mechanisms to identify the application. App-ID continually monitors the application state, re-classifying the traffic and identifying the different functions that are being used. The security policy determines how to treat the application: block, allow, or securely enable (scan for, and block embedded threats, inspect for unauthorized file transfer and data patterns, or shape using QoS).

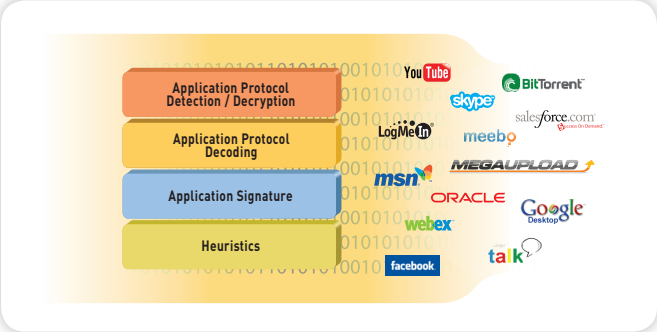

User-ID: Enabling Applications by Users and Groups

Traditionally, security policies were applied based on IP addresses, but the increasingly dynamic nature of users and computing means that IP addresses alone have become ineffective as a mechanism for monitoring and controlling user activity. User-ID allows organizations to extend user- or group-based application enablement polices across Microsoft Windows, Apple Mac OS X, Apple iOS, and Linux users.

User information can be harvested from enterprise directories (Microsoft Active Directory, eDirectory, and Open LDAP) and terminal services offerings (Citrix and Microsoft Terminal Services) while integration with Microsoft Exchange, a Captive Portal, and an XML API enable organizations to extend policy to Apple Mac OS X, Apple iOS, and UNIX users that typically reside outside of the domain.

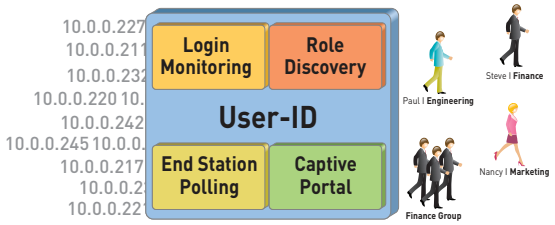

Content-ID: Protecting Allowed Traffic

Many of today's applications provide significant benefit, but are also being used as a delivery tool for modern malware and threats. Content-ID, in conjunction with App-ID, provides administrators with a two-pronged solution to protecting the network. After App-ID is used to identify and block unwanted applications, administrators can then securely enable allowed applications by blocking vulnerability exploits, modern malware, viruses, botnets, and other malware from propagating across the network, all regardless of port, protocol, or method of evasion. Rounding out the control elements that Content-ID offers is a comprehensive URL database to control web surfing and data filtering features.

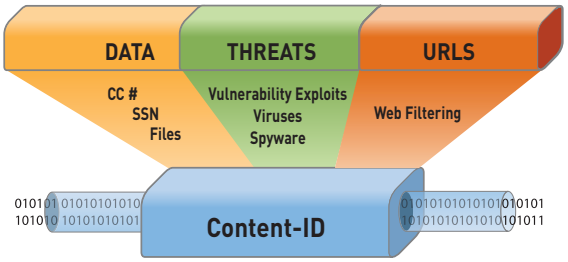

Secure Application Enablement

The seamless integration of App-ID, User-ID, and Content-ID enables organizations to establish consistent application enablement policies, down to the application function level in many cases, that go far beyond basic allow or deny. With GlobalProtect™, the same policies that protect users within the corporate headquarters are extended to all users, no matter where they are located, thereby establishing a logical perimeter for users outside of the network.

Secure enablement policies begin with App-ID determining the application identity, which is then mapped to the associated user with User-ID, while traffic content is scanned for threats, files, data patterns, and web activity by Content-ID. These results are displayed in Application Command Center (ACC) where the administrator can learn, in near real-time, what is happening on the network. Then, in the policy-editor, the information viewed in ACC about applications, users, and content can be turned into appropriate security policies that block unwanted applications, while allowing and enabling others in a secure manner. Finally, any detailed analysis, reporting, or forensics can be performed, again, with applications, users, and content as the basis.

Application Command Center: Knowledge is Power

Application Command Center (ACC) graphically summarizes the log database to highlight the applications traversing the network, who is using them, and their potential security impact. ACC is dynamically updated, using the continuous traffic classification that App-ID performs; if an application changes ports or behavior, App-ID continues to see the traffic, displaying the results in ACC. New or unfamiliar applications that appear in ACC can be quickly investigated with a single click that displays a description of the application, its key features, its behavioral characteristics, and who is using it.

Additional visibility into URL categories, threats, and data provides a complete and well-rounded picture of network activity. With ACC, an administrator can very quickly learn more about the traffic traversing the network and then translate that information into a more informed security policy.

Policy Editor: Translating Knowledge into Secure Enablement Policies

The knowledge of which applications are traversing the network, who is using them, and what the potential security risks are, empowers administrators to quickly deploy application-, application function-, and port-based enablement policies in a systematic and controlled manner. Policy responses can range from open (allow), to moderate (enabling certain applications or functions, then scan, or shape, schedule, etc.), to closed (deny). Examples may include:

- Protect an Oracle database by limiting access to finance groups, forcing the traffic across the standard ports, and inspecting the traffic for application vulnerabilities.
- Enable only the IT group to use a fixed set of remote management applications (e.g., SSH, RDP, Telnet) across their standard ports.
- Define and enforce a corporate policy that allows and inspects specific webmail and instant messaging usage but blocks their respective file transfer functions.
- Allow Microsoft SharePoint Administration to be used by only the administration team, and allow access to Microsoft SharePoint Documents for all other users.
- Deploy web enablement policies that that allow and scan traffic to business related web sites while blocking access to obvious non-work related web sites and “coaching” access to others through customized block pages.

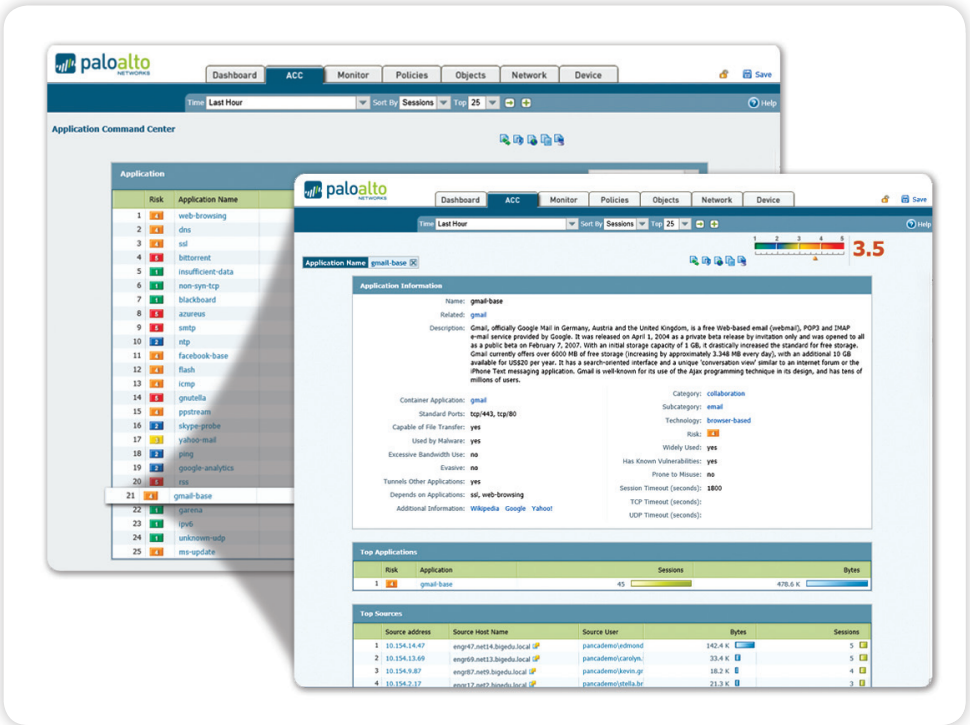

Application Visibility
View application activity in a clear, easy-to-read format. Add and remove filters to learn more about the application, its functions and who is using them.

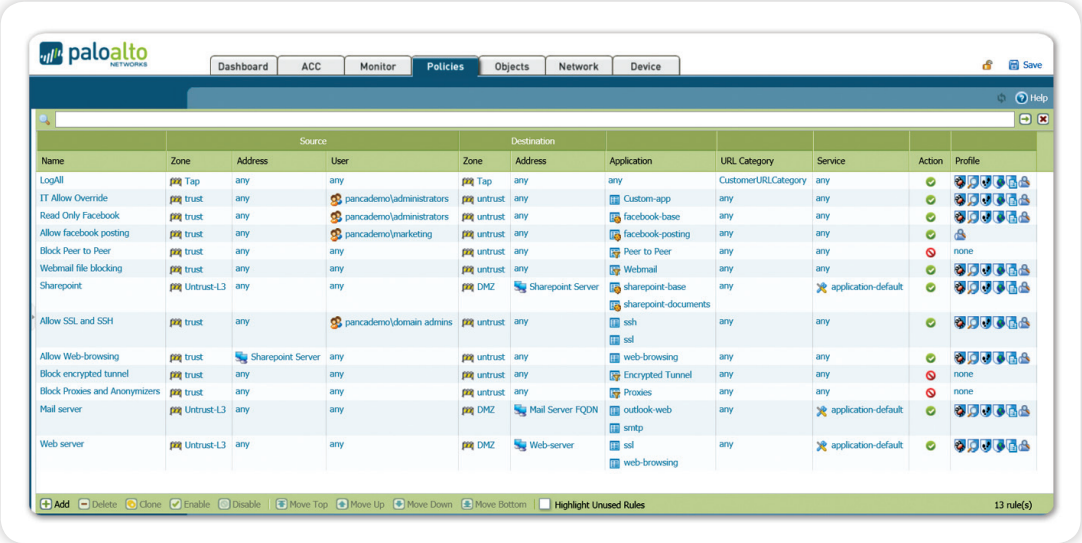

Unified Policy Editor
A familiar look and feel enables the rapid creation and deployment of policies that control applications, users and content.

- Implement QoS policies to allow the use of both bandwidth-intensive media applications and websites but limit their impact on VoIP applications.
- Decrypt SSL traffic to social networking and webmail sites and scan for malware and exploits.
- Allow downloads of executable files from uncategorized websites only after user acknowledgement to prevent drive-by-downloads via zero-day exploits.
- Deny all traffic from specific countries or block unwanted applications such as P2P file sharing, circumventors, and external proxies.

The tight integration of application control, based on users and groups, and the ability to scan the allowed traffic for a wide range of threats, allows organizations to dramatically reduce the number of policies they are deploying along with the number of employee adds, moves, and changes that may occur on a day-to-day basis.

Policy Editor: Protecting Enabled Applications

Securely enabling applications means allowing access to the applications, then applying specific threat prevention and file, data, or URL filtering policies. Each of the elements included in Content-ID can be configured on a per-application basis.

- **Intrusion Prevention System (IPS):** Vulnerability protection integrates a rich set of intrusion prevention system (IPS) features to block network and application-layer vulnerability exploits, buffer overflows, DoS attacks, and port scans.
- **Network Antivirus:** Stream-based antivirus protection blocks millions of malware variants, including PDF viruses and malware hidden within compressed files or web traffic (compressed HTTP/HTTPS). Policy-based SSL decryption enables organizations to protect against malware moving across SSL encrypted applications.
- **URL Filtering:** A fully-integrated, customizable URL filtering database allows administrators to apply granular web-browsing policies, complementing application visibility and control policies and safeguarding the enterprise from a full spectrum of legal, regulatory, and productivity risks.
- **File and Data Filtering:** Data filtering features enable administrators to implement policies that will reduce the risks associated with file and data transfers. File transfers and downloads can be controlled by looking inside the file (as opposed to looking only at the file extension), to determine if it should be allowed or not. Executable files, typically found in drive-by downloads, can be blocked, thereby protecting the network from unseen malware propagation. Finally, data filtering features can detect, and control the flow of confidential data patterns (credit card and social security numbers).

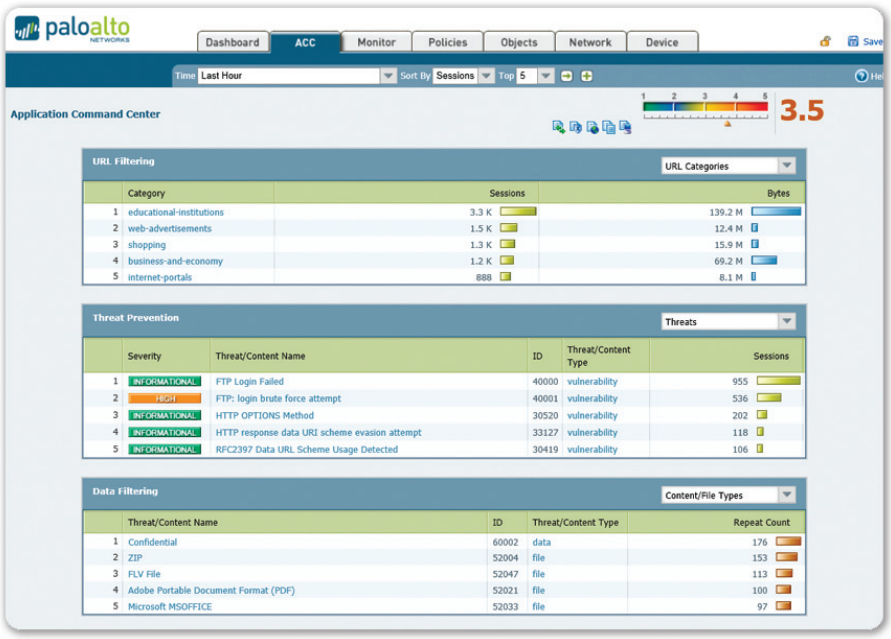

Content and Threat Visibility
View URL, threat and file/data transfer activity in a clear, easy-to-read format. Add and remove filters to learn more about individual elements.

Modern Malware Detection and Prevention

Malware has evolved to become an extensible networked application that provides attackers with unprecedented access and control inside of the targeted network. As the power of modern malware increases, it is critical that enterprises be able to detect these threats immediately, even before the threat has a defined signature. Palo Alto Networks next-generation firewalls provide organizations with a multi-faceted approach based on the direct analysis of both executable files and network traffic to protect their networks even before signatures are available.

- **WildFire™:** Using a cloud-based approach, WildFire exposes previously unseen malicious executable files by directly observing their behavior in a secure virtualized environment. WildFire looks for malicious actions within Microsoft Windows executable files such as changing registry values or operating system files, disabling security mechanisms, or injecting code into running processes. This direct analysis quickly and accurately identifies malware even when no protection mechanism is available. The results are immediately delivered to the administrator for an appropriate response and a signature is automatically developed and delivered to all customers in the next available content update.
- **Behavioral Botnet Detection:** App-ID classifies all traffic at the application level, thereby exposing any unknown traffic on the network, which is often an indication of malware or other threat activity. The behavioral botnet report analyzes network behavior that is indicative of a botnet infection such as repeatedly visiting malware sites, using dynamic DNS, IRC, and other potentially suspicious behaviors. The results are displayed in the form of a list of potentially infected hosts that can be investigated as possible members of a botnet.

Traffic Monitoring: Analysis, Reporting and Forensics

Security best practices dictate that administrators strike a balance between being proactive, continually learning and adapting to protect the corporate assets, and being reactive, investigating, analyzing, and reporting on security incidents. ACC and the policy editor can be used to proactively apply application enablement policies, while a rich set of monitoring and reporting tools provide organizations with the necessary means to analyze and report on the application, users and content flowing through the Palo Alto Networks next-generation firewall.

- **App-Scope:** Complementing the real-time view of applications and content provided by ACC, App-scope provides a dynamic, user-customizable view of application, traffic, and threat activity over time.
- **Reporting:** Predefined reports can be used as-is, customized, or grouped together as one report in order to suit the specific requirements. All reports can be exported to CSV or PDF format and can be executed and emailed on a scheduled basis.
- **Logging:** Real-time log filtering facilitates rapid forensic investigation into every session traversing the network. Log filter results can be exported to a CSV file or sent to a syslog server for offline archival or additional analysis.
- **Trace Session Tool:** Accelerate forensics or incident investigation with a centralized correlated view across all of the logs for traffic, threats, URLs, and applications related to an individual session.

GlobalProtect: Consistent Security Everywhere

Applications are not the only forces of change in the enterprise. Increasingly, end-users simply expect to connect and work from any location with any device of their choosing. As a result, IT teams are struggling to extend security to these devices and locations that may be well beyond the traditional perimeter of the enterprise. GlobalProtect meets this challenge by extending consistent security policies to all users regardless of their location and device.

First, GlobalProtect ensures secure connectivity for all users with a transparent VPN that supports a range of devices including Microsoft Windows, Apple Mac OS X, and Apple iOS. Once connected, all traffic is classified by the firewall, enablement policies are applied, traffic is scanned for threats, and the network and user are protected.

Additionally, GlobalProtect can apply additional controls based on the state of the end-user's device. For example, a user could be denied access to certain applications or sensitive areas of the network if the device has an out of date antivirus or does not have disk encryption enabled. This allows IT to safely enable application usage across a range of end-user device types, while retaining a consistent next-generation approach to security.

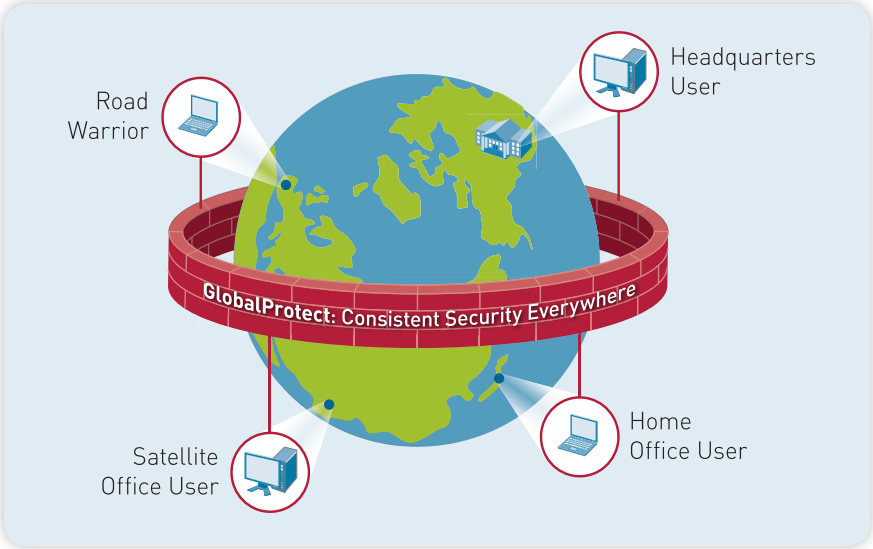

GlobalProtect
Enforce consistent secure application enablement policies for all users, no matter where they are located.

paloalto NETWORKS
the network security company™

3300 Olcott Street
Santa Clara, CA 95054

Main: +1.408.573.4000
Sales: +1.866.320.4788
Support: +1.866.898.9087

www.paloaltonetworks.com