

Palo Alto Networks Next Generation Firewalls

Fundamental shifts in application usage, user behavior, and network infrastructure have resulted in an evolved threat landscape that has exposed weaknesses in traditional port-based firewall protection.

Users are accessing an increasing number of applications with a wide range of device types, often times to get their job done, yet with little regard to the business or security risks. Meanwhile, datacentre expansion, network segmentation, virtualisation and mobility initiatives are forcing you to re-think how to enable access to applications and data, yet protect your network from a new, more sophisticated class of advanced threats that are adept at evading traditional security mechanisms.

Next Generation Firewalls

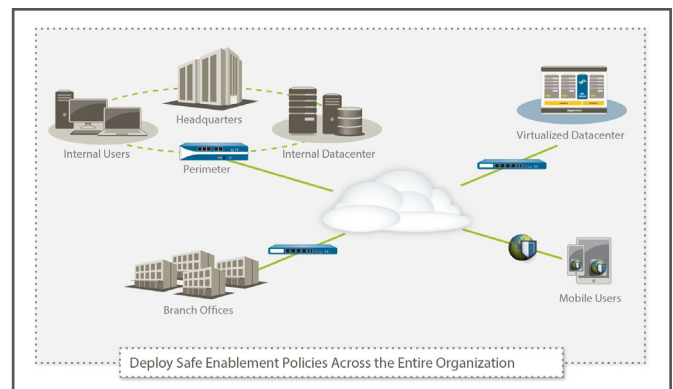
Historically you were left with two basic choices, either block everything in the interest of network security or enable everything in the interest of business. These choices left little room for compromise. Palo Alto Networks pioneered the next-generation firewall to enable you to accomplish both objectives—safely enable applications while protecting against both known and unknown threats. The Palo Alto Networks next-generation firewall acts as the basis of an enterprise security platform that is designed from the ground up to address the most sophisticated threats. Unique to the platform is a traffic classification that natively inspects all applications, threats and content, then ties that traffic to the user, regardless of location or device type. The application, content, and user—the elements that run your business—then become integral components of your enterprise security policy. The result is the ability to align security with key business initiatives.

- Safely enable applications, users and content by classifying all traffic, determining the business use case and assigning policies to allow and protect the relevant applications.
- Prevent threats by eliminating unwanted applications to reduce the threat footprint and applying targeted security policies to block known vulnerability exploits, viruses, spyware, botnets and unknown malware (APTs).
- Embrace mobile computing by ensuring devices are properly configured and that they are protected from threats.
- Protect your datacentres through validation of applications, isolation of data, control over rogue applications and high speed threat prevention.
- Secure cloud-computing environments with increased visibility and control; deploy and maintain security policies at the same pace as your virtual machines.

Enterprise Security

Your business groups are either already using, or demanding to use, the latest and greatest applications for both personal and professional purposes.

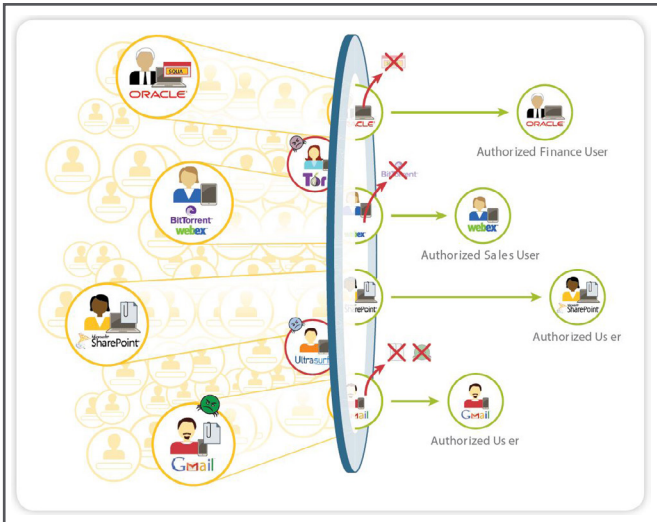
The Palo Alto Networks enterprise security platform allows you to support business initiatives while improving your overall security posture and reducing security incident response time.



Empower Your Business

The Palo Alto Networks enterprise security platform allows you to empower your business with policies that revolve around applications, content and users. Unique to the Palo Alto Networks platform is the use of a positive control model that allows you to enable specific applications or functions and block all else (implicitly or explicitly).

In order to achieve this level of control, all traffic—including of the application, the associated content or threat and the user—must be proactively inspected and classified at the firewall (not after the fact). The complete context of the application, associated content, and user identity is then used for all security policy decisions.



Applications, content, users, and devices - all under your control

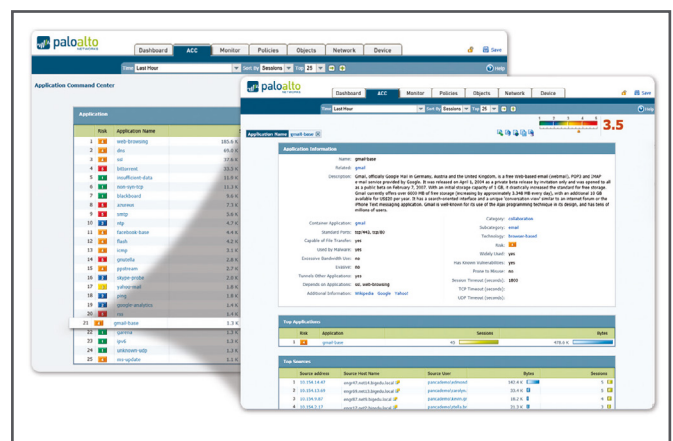
- Classify all traffic, across all ports, all the time.** Today, applications and the associated content can easily bypass a port-based firewall using a variety of techniques. The Palo Alto Networks enterprise security platform addresses the traffic classification visibility limitations that plague port-based security by natively applying multiple classification mechanisms to the traffic stream as soon as the platform sees it, to determine the identity of the applications traversing your network and if they are carrying any threats or malware. All traffic is classified regardless of port, encryption (SSL or SSH) or evasive techniques employed. Unidentified applications, typically a small percentage of traffic, yet high in potential risk, are automatically categorised for systematic management.
- Reduce the threat footprint, prevent cyber attacks.** Once the traffic is fully classified, you can protect your network from a range of cyber attacks by allowing specific applications and denying all others to reduce the network threat footprint. Coordinated cyber attack protection can then be applied to the allowed traffic, blocking known malware sites, preventing vulnerability exploits, viruses, spyware and malicious DNS queries. Custom or otherwise unknown malware found in the applications on your network is analysed and identified by executing the files and directly observing their malicious behavior in a virtualised sandbox environment. If new malware is discovered, a signature for the infecting file and related malware traffic is automatically generated and delivered to you.

- Map application traffic and associated threats to users and devices.** To improve your security posture and reduce incident response times it is critical that you be able to determine application usage, mapped to user and device type, and be able to apply that context to security policies. Integration with a wide range of enterprise user repositories provides the identity of the Microsoft Windows, Mac OS X, Linux, Android, or iOS user, and device accessing the application. The combined visibility and control over both users and devices means you can safely enable the use of any application traversing your network, no matter where the user is or what device type they are using.

By establishing complete context of which applications are in use, what content or threat they may carry and the associated user or device, you gain a more comprehensive view into network activity that can help to streamline policy management, improve your security posture and accelerate incident investigation.

Complete Context

Security best practices dictate that the decisions you make regarding policies, your ability to report on network activity and your forensics capacity are dependent on context. The context of the application in use, the website visited, the associated payload and the user are all valuable data points in your quest to protect your network. Knowing exactly which applications are traversing your Internet gateway, operating within your datacentre or being used by remote users as opposed to the broader set of traffic that is port-based, means you can apply specific policies to those applications, complete with coordinated threat protection. The knowledge of who the user is, not just their IP address, adds another contextual element that allows you to be more granular in your policy assignment.



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.

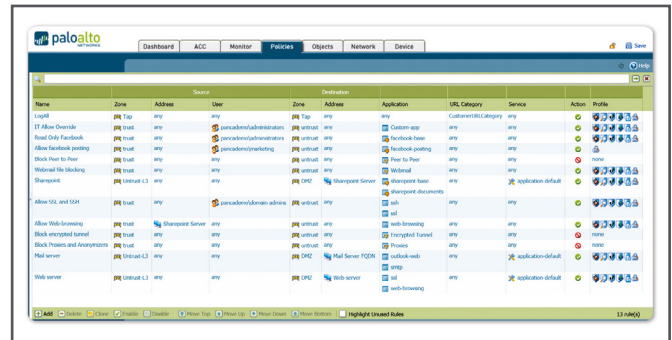
A rich set of graphical visualisation and log filtering tools provides you with the context of the application activity, the associated content or threat, who the user is, and on what type of device. Each of these data points by themselves paints a partial picture of your network, yet when taken in complete context provide a full view of the potential security risk, allowing you to make more informed policy decisions. All traffic is continuously classified and as their state changes, it is logged for analysis and the graphical summaries are dynamically updated, displaying the information in an easy-to-use, web-based interface.

- At the Internet gateway, you can investigate new or unfamiliar applications to quickly see a description of the application, its behavioral characteristics, and who is using it. Additional visibility into URL categories, threats, and data patterns provides a more well-rounded picture of network traffic traversing the gateway.
- All files analysed for unknown malware by WildFire™ are logged on-box with full access to details including the application used, the user, the file type, target OS, and malicious behaviors observed.
- Within the datacentre, you can verify all applications under use, and ensure that they are being used by only those who have been authorised. Added visibility into datacentre activity can confirm that there are no misconfigured applications or rogue use of SSH or RDP.
- Across all deployment scenarios, unknown applications, typically a small percentage on every network, is categorised for analysis and systematic management.

In many cases, you may not be fully aware of which applications are in use, how heavily they are used or by whom. Complete visibility into the business relevant aspects of your network traffic—the application, the content, and the user—is the first step towards more informed policy control.

Reducing Risk by Enabling Applications

Traditionally, the process of reducing risk meant that you had to limit access to network services and possibly hinder your business. Today, risk reduction means safely enabling applications using a business-centric approach that helps you strike a balance between the traditional deny-everything approach and the allow-all approach.



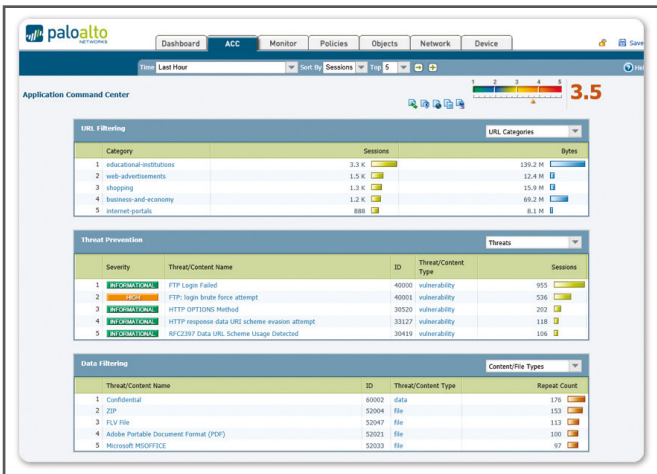
Unified Policy Editor: A familiar look and feel enables the rapid creation and deployment of policies that control applications, users and content.

- Use application groups and SSL decryption to limit webmail and instant messaging to a few specific application variants, inspect them for all threats, and upload unknown suspect files (PDFs, Office Documents, EXEs) to WildFire for analysis and signature development.
- Control web-surfing for all users by allowing and scanning traffic to business-related websites and blocking access to obvious non-work related websites; “coach” access to questionable sites through customised block pages.
- Explicitly block all peer-to-peer file transfer applications for all users using dynamic application filters.
- Embrace mobile devices by extending your Internet gateway policies to remote users with GlobalProtect™.

In the datacentre your policies will utilise context as a means of confirming that your datacentre applications are running on their standard ports, finding rogue applications, validating users, isolating data, and protecting business critical data from threats. Examples may include:

- Using security zones, isolate the Oracle-based credit card number repository, forcing the Oracle traffic across its standard ports while inspecting the traffic for inbound threats and limiting access only to the finance group.
- Create a remote management application group (e.g., SSH, RDP, Telnet) for only the IT department to use within the datacentre.

In your virtual datacentre, use dynamic objects to help automate security policy creation as SharePoint virtual machines are established or taken down or travel across your virtual environment.



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.

- Limit unauthorised file and data transfers.** Data filtering features enable your administrators to implement policies that will reduce the risks associated with unauthorised file and data transfers. File transfers can be controlled by looking inside the file (as opposed to looking only at the file extension), to determine if the transfer action should be allowed or not. Executable files, typically found in drive-by downloads, can be blocked, thereby protecting your network from unseen malware propagation. Data filtering features can detect, and control the flow of confidential data patterns (credit card or social security numbers as well as custom patterns).
- Control web surfing.** A fully-integrated, customisable URL filtering engine allows your 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. Your URL filtering policies are applied to all web traffic including cached results, and translation engine results, both of which are common URL filtering policy evasion tactics. If you are leveraging the browser-based safe-search settings, block policies for all search results will be enforced when the "strict" setting is enabled. In addition, the URL categories can be leveraged into the policies to provide further granularity of control for SSL decryption, QoS, or other rule bases. In many cases, the complementary nature of application control and URL filtering can help you reduce administrative efforts and your overall cost of ownership.

Centralised Management

Your global deployment of the Palo Alto Networks enterprise security platforms can be managed individually via a command line interface (CLI) or through a full-featured browser-based interface.

For large-scale deployments, you can use Panorama to globally control all device aspects using device groups and templates. All of the visibility, policy editing, reporting and logging features are supported by Panorama, giving you the same level of contextual control over your global deployment as you have over a single appliance.

Role based administration combined with pre-and post rules allows you to balance centralised control with the need for local policy editing and device configuration flexibility. Whether using the device's web interface or Panorama's, the interface look and feel is identical, ensuring that there is no learning curve when moving from one to another.

Your administrators can use any of the provided interfaces to make changes at any time without needing to worry about synchronisation issues. Additional support for standards-based tools such as SNMP, and REST-based APIs allow you to integrate with third-party management tools.

Reporting and Logging

Security best practices means striking a balance between ongoing management efforts and being reactive, which may involve investigating and analysing security incidents or generating day-to-day reports.

- Reporting:** Predefined reports can be used as-is, customised, 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 your network. Complete context of the application, the content, including malware detected by WildFire, and the user can be used as a filter criteria and the results can be exported to a CSV file or sent to a syslog server for offline archival or additional analysis. Logs that have been aggregated by Panorama can also be sent to a syslog server for added analysis or archival purposes.

Purpose-Built Hardware or Virtualised Platforms

The Palo Alto Networks enterprise security platform is available in either a purpose-built hardware platform that scales from an enterprise branch office to a high-speed datacentre or as a virtualised form factor to support your cloud-based computing initiatives. When you deploy the platforms in either hardware or virtual form factors, you can use Panorama, an optional centralised management offering to gain visibility into traffic patterns, deploy policies, generate reports and deliver content updates from a central location.