


Palo Alto Networks Next-generation Firewall Overview

Fundamental shifts in application usage, user behavior, and network infrastructure create a threat landscape that exposes weaknesses in traditional port-based network security. Your users want access to an increasing number of applications operating across a wide range of device types often with little regard for the business or security risks. Meanwhile, datacenter expansion, network segmentation, virtualization, and mobility initiatives are forcing you to rethink how to enable access to applications and data, while protecting your network from a new, more sophisticated class of advanced threats that evade traditional security mechanisms.


Historically, you were left with two basic choices—either block everything in the interest of network security, or enable everything in the interest of your business. These choices left little room for compromise. The Palo Alto Networks® enterprise security platform provides you with a way to safely enable the applications your users need by allowing access while preventing cybersecurity threats.

Our next-generation firewall is the core of the enterprise security platform, designed from the ground up to address the most sophisticated threats. The next-generation firewall inspects all traffic - inclusive of applications, threats, and content—and ties it to the user, regardless of location or device type. The application, content, and user—the elements that run your business—become integral components of your enterprise security policy. The result is the ability to align security with your 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 access to relevant applications.
- Prevent threats by eliminating unwanted applications to reduce your threat footprint and apply targeted security policies to block known vulnerability exploits, viruses, spyware, botnets, and unknown malware (APTs).
- Protect your datacenters through the validation of applications, isolation of data, control over rogue applications, and high-speed threat prevention.
- Secure public and private cloud computing environments with increased visibility and control; deploy, enforce and maintain security policies at the same pace as your virtual machines.
- Embrace safe mobile computing by extending the enterprise security platform to users and devices no matter where they are located.



The enterprise security platform helps your organization address a spectrum of security requirements based upon a common principle. Using a balanced combination of network security with global threat intelligence and endpoint protection, your organization can support business initiatives while improving your overall security posture, and reducing security incident response time.



Applications, content, users, and devices—all under your control.

Using Security to Empower Your Business

Our enterprise security platform allows you to empower your business with policies that revolve around applications, users, and content. It uses a positive control model, a design unique to our platform, that permits you to enable specific applications or functions, and block all else (implicitly or explicitly). The next-generation firewall performs a full stack, single pass inspection of all traffic across all ports, thus providing complete context of the application, associated content, and user identity as the basis for your security policy decisions.

- Classify all traffic, across all ports, all the time. Today, applications and their associated content can easily bypass a port-based firewall using a variety of techniques. Our enterprise security platform natively applies multiple classification mechanisms to the traffic stream to identify applications, threats and 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 categorized for systematic management.
- Reduce the threat footprint; prevent cyber attacks. Once traffic is fully classified, you can reduce the network threat footprint by allowing specific applications and denying all others. Coordinated cyber attack prevention can then be applied to block known malware sites and prevent vulnerability exploits, viruses, spyware, and malicious DNS queries. Any custom or unknown malware is analyzed and identified by executing the files, and directly observing


their malicious behavior in a virtualized sandbox environment. When 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's critical to map application usage to user and device type—and be able to apply that context to your 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 the type of device they are using.

Establishing the context of the specific applications in use, the content or threat they may carry, and the associated user or device helps you streamline policy management, improve your security posture, and accelerate incident investigation.

Complete Context Means Tighter Security Policies

Security best practices dictate that the decisions you make regarding policies, your ability to report on network activity, and your forensics capacity depend 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 effort to protect your network. When you know exactly which applications are traversing your Internet gateway, operating within your datacenter or cloud environment, or being used by remote users, 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 empowers 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 visualization 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, the changes are 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 analyzed 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 datacenter, verify all applications under use, and ensure that they are only being used by authorized users. Added visibility into datacenter activity can confirm that there are no misconfigured applications or rogue use of SSH or RDP.
- In public and private cloud environments, enforce policy and protect applications with the enterprise security platform while keeping pace with the creation and movement of your virtual servers.
- Across all deployment scenarios, unknown applications – typically a small percentage on every network—can be categorized 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



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

business-centric approach that helps you strike a balance between the traditional deny-everything approach and the allow-all approach.

- 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 (EXE, DLL, ZIP files, PDF documents, Office Documents, Java, and Android APK) 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 customized 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 and threat prevention capabilities to remote users with GlobalProtect™.

In the datacenter, use context to confirm that your datacenter applications are running on their standard ports, find rogue applications, validate users, isolating data, and protect 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 datacenter.
- In your virtual datacenter, use dynamic objects to help automate security policy creation as SharePoint virtual machines are established or taken down or travel across your virtual environment.

Protecting Enabled Applications and Content

When you apply threat prevention and content scanning policies, the context of the application and the user become integral components of your security policy. Full context within your threat prevention policies neutralizes evasion tactics such as port-hopping and tunneling. Reduce the threat target surface area by enabling a select set of applications, and then apply threat prevention and content scanning policies to that traffic.

Threat protection and content scanning elements available within your policies include:

- **Prevent known threats using IPS and network antivirus/anti-spyware.** Protection from a range of known threats is accomplished with single pass inspection using a uniform signature format and a stream-based scanning engine. Intrusion Prevention System (IPS) features block network and application layer vulnerability exploits, buffer overflows, DoS attacks, and port scans. Antivirus/anti-spyware protection blocks millions of malware variants, including those hidden within compressed files or web traffic (compressed HTTP/HTTPS), as well as known PDF viruses. For traffic encrypted with SSL, you can selectively apply policy-based decryption and then inspect the traffic for threats, regardless of port.
- **Block unknown or targeted malware with WildFire.** Unknown or targeted malware (e.g., Advanced Persistent Threats) hidden within files can be identified and analyzed by WildFire, which directly executes and observes unknown files virtualized sandbox environment in the cloud or on the WF-500 appliance. WildFire monitors more than 100 malicious behaviors and if malware is found, a signature is automatically developed and delivered to you in as little as 15 minutes. All major file types are supported by WildFire including: PE files; Microsoft Office .doc, .xls, and .ppt; Portable Document Format (PDF); Java Applet (jar and class); and Android Application Package (APK). In addition, WildFire analyzes links in email to stop spearphishing attacks.



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.

- Identify bot-infected hosts and disrupt network activity from malware.** Complete, contextual classification of all applications, across all ports, including any unknown traffic, can often expose anomalies or threats in your network. Use command & control App-ID™s, behavioral botnet report, DNS sinkholing, and passive DNS to quickly correlate unknown traffic, suspicious DNS, and URL queries with infected hosts. Apply global intelligence to intercept and sinkhole DNS queries for malicious domains.
- Limit unauthorized file and data transfers.** Data filtering features enable your administrators to implement policies that will reduce the risks associated with unauthorized 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, customizable 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.

- Device-based policy for application access. Using GlobalProtect, an organization can set specific policies to control which devices can access particular applications and network resources. For example, ensure that laptops are compliant with the corporate image before allowing access to the datacenter. Check if the mobile device is up to date, corporate owned, and fully patched before accessing sensitive data.

Centralized Management

The 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 deliver visibility, policy editing, reporting, and logging features for all of your hardware and virtual appliance firewalls.

Panorama provides 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 centralized 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 synchronization 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 analyzing security incidents or generating day-to-day reports.

- **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 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.

In addition to the reporting and logging capabilities provided by the Palo Alto Networks enterprise security platform, integration is also available with 3rd party SIEM tools, such as Splunk for Palo Alto Networks. These tools provide additional reporting and data visualization capabilities, and enable you to correlate security events across multiple systems in your enterprise.

Purpose-built Hardware or Virtualized Platforms

Our next-generation firewall is available in either a purpose-built hardware platform that scales from an enterprise branch office to a high-speed datacenter, or in a virtualized form factor to support your cloud-based computing initiatives. Palo Alto Networks supports the broadest range of virtual platforms to cover your diverse virtualized data center and public & private cloud requirements. The VM-Series firewall platform is available for VMware ESXi, NSX, Citrix SDX, Amazon AWS, and KVM hypervisors. When you deploy our platforms in either hardware or



Panorama can be deployed on a dedicated appliance or in a distributed manner to maximize scalability.

virtual form factors, you can use Panorama for centralized management.

Purpose-Built Hardware or Virtualized Platforms

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