10 THINGS YOUR NEXT FIREWALL MUST DO

paloalto NETWORKS®

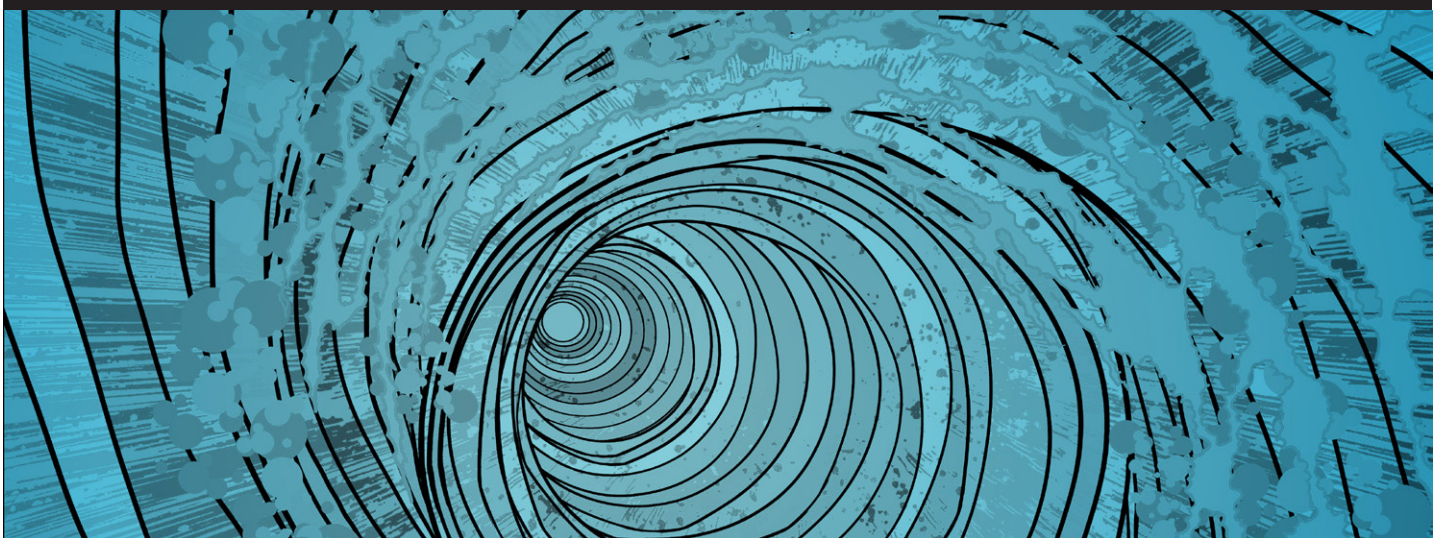

Your network is more complex than ever before. Your employees are accessing any application they want, using work or personal devices. Oftentimes, these applications span both personal and work-related usage, but the business and security risks are often ignored. Prospective employees are asking about application usage policies before accepting a job. Adding another layer of complexity is the underlying concern about the effectiveness of your cybersecurity posture. Is your business a target for a cyberattack? Is it a question of when, as opposed to if? And are you as prepared as you could be? The complexity of your network and your security infrastructure may limit or slow your ability to respond to these and other cybersecurity challenges.

When increasing complexity limits or slows the decision-making process, it's helpful to focus on the fundamentals as a means of addressing the situation at hand in a more effective manner. Remember the three fundamental functions that your firewall was designed to execute:

1. Operate as the core of your network security infrastructure.
2. Act as the access control point for all traffic – allowing or denying traffic into the network based on policy.
3. Eliminate the risk of the unknown by using a positive control model, which allows what you want; all else is implicitly denied.

Over time, these fundamental functions have been nullified by the very traffic they were meant to control. Applications evolved at a faster pace than the firewall. As a result, the firewall has trouble exerting the levels of control you need to protect your digital assets.

Port hopping, the use of non-standard ports, and the use of encryption are a few of the ways in which applications have become more accessible. These same techniques are also used by cyberattackers, both directly in the cyberthreats they create and indirectly by hiding the threats within the application traffic itself. Further complicating these challenges is the fact that your employees are probably using these applications to help get their jobs done. Some examples of the applications and threats found on your network include:

- **Common end-user applications:** These applications include social media, filesharing, video, instant messaging and email. Collectively they represent more than 35 percent of the applications on your network. Employees may use some of them for work purposes, while others will be purely for personal use. These applications are often highly extensible and include features that may introduce unwarranted risk. These applications represent both business and security risks, and your challenge will be how to strike an appropriate balance between blocking some and securely enabling others.
- **Core-business applications:** These are the applications that run your business; they house your most valued assets, and include databases, directories and ERP applications in your data center and such applications as Salesforce and Workday in the cloud. This group of applications is heavily targeted by cyberattackers using multifaceted attacks. Your challenge is going to be how best to isolate and protect these applications from stealthy attacks that easily evade your firewall and IPS through common evasion techniques.
- **Infrastructure and custom applications:** This group of applications represents core infrastructure applications, such as SSL, SSH and DNS, as well as internally developed, custom or unknown applications. These applications are commonly used to mask command-and-control traffic generated by bots and other types of malware. Interestingly, many of these applications use a wide range of non-standard ports. Many of the applications that use SSL never use port 443, and others hop ports.

To try and address these challenges, there has been an increased focus on the fundamentals of the firewall: All network firewall vendors are rethinking how they identify and control traffic based on the application itself, instead of just the port and protocol.

Collectively, firewalls that are capable of exerting an application-centric approach to firewall control are now described as "next-generation," and every firewall vendor acknowledges that application control is an increasingly critical part of network security.

There are two obvious reasons for this renewed focus on the fundamentals. First, applications and the associated threats can easily slip by port-based firewalls as well as the additive threat prevention elements. Second, the firewall is the only place that sees all the traffic flowing across your network and it is still the most logical location to enforce access control policies. The value of this renewed focus is obvious: Your security posture should improve, while the administrative effort associated with firewall management and incident response should shrink or, at a minimum, remain constant.

Next-Generation Firewalls Defined

Most network security vendors are now offering application visibility and control by either adding application signatures to their IPS engine or offering you an add-on license for an application control module. In either case, these options are additive to a port-based firewall and do little to help you focus on the fundamental tasks your firewall is designed to execute.

How effectively your business operates is heavily dependent upon the applications your employees use and the content that the applications themselves carry. Merely allowing some and then blocking others, may inhibit your business. If your security team is looking at next-generation firewall features and capabilities, the most important consideration is whether the next-generation firewall will empower your security team to safely enable applications to the benefit of the organization. Consider the following:

- Will the next-generation firewall increase visibility and understanding of the application traffic, including that destined to SaaS applications?
- Will the traffic-control-policy response options be broader than just "allow" or "deny"?
- Will your network be protected from threats and cyberattacks – both known and unknown?
- Can you systematically identify and manage unknown traffic?
- Can you implement the desired security policies without compromising on performance?
- Will the administrative effort your team devotes to firewall management be reduced?
- Will your job of managing risk be easier and more effective?
- Can the policies you enable help contribute to the business bottom line?

Next-Generation Firewall Requirements

1. Identify applications regardless of port, protocol, evasive tactic or decryption.
2. Identify users regardless of device or IP address.
3. Decrypt encrypted traffic.
4. Protect in real time against known and unknown threats embedded across applications.
5. Deliver predictable, multi-gigabit in-line deployment.

If the answers to the above questions are "yes," then your decision to transition from legacy firewalls to next-generation firewalls

is easy to justify. The next step is to consider the alternative solutions that firewall vendors are providing. When evaluating the available alternatives, it is important to consider the architectural differences between the next-generation firewall offerings and the associated impacts in terms of real-world functions/features, operations and performance.

Architectural Considerations for Firewall Traffic Classification

In building next-generation firewalls, security vendors have taken one of two architectural approaches:

1. Build application identification into the firewall as the primary classification engine.
2. Add an application signature pattern-matching engine to a port-based firewall.

Both approaches can recognize applications but with varying degrees of success, usability and relevance. Most importantly, these architectural approaches dictate a specific security model for application policies – either positive (define what is allowed; deny all else), or negative (define what to block; allow all else).

- A positive security model (firewall or otherwise) gives you the ability to write policies that allow specific applications or functions (e.g., WebEx®, SharePoint®, Gmail™) and then everything else is implicitly denied. In order to achieve this level of control, all traffic must be proactively classified at the firewall (not after the fact) to ensure the appropriate traffic is allowed and the rest denied. By establishing full visibility into all traffic, businesses are able to reduce the administrative effort associated with gaining visibility into network activity, policy management and incident investigation. Security implications may include better protection against known and unknown cyberattacks, even though you may be allowing a wider range of applications on your network and improved control over unknown applications through the deny-all-else premise a firewall provides.
- A negative security model (IPS, AV, etc.) gives you the ability to look for and block threats or unwanted applications specifically and allow everything else. This means that all traffic is not necessarily classified – only enough to fulfill the targeted block list. This technique may be sufficient in selectively finding and blocking threats or unwanted applications, but a negative security model is ill-suited to act as the primary means of controlling all traffic on your network, relegating this technique to be a port-based firewall helper. The business ramifications of a negative security model include the increased administrative effort associated with multiple policies and duplicate log databases.

10 Things Your Next Firewall Must Do should be viewed as proof points that the architecture and control model outlined above are critical to delivering on the promise of identifying and safely enabling applications at the firewall. Use this as a mechanism to assemble your own next-generation firewall criteria.

10 Things Your Next Firewall Must Do

Firewall selection criteria will typically fall into three areas: security functions, operations and performance. The security functions element corresponds to the efficacy of the security controls and the ability of your team to manage the risk associated with the applications that are traversing your network. From an operations perspective, the big question is, "Where does application policy live, and how hard or complex is it for your team to manage?" The performance difference is simple: Can the firewall do what it's supposed to do at the required throughput your business needs? While each organization will have varied requirements and priorities within the three selection criteria, the 10 things your next firewall must do are:

1 Your next firewall must identify and control applications and application functions on all ports, all the time

Business case: Application developers no longer adhere to standard port/protocol/application development methodology. More and more applications are capable of operating on non-standard ports or can hop ports (e.g., instant messaging applications, peer-to-peer file sharing, or VoIP). Additionally, users are increasingly savvy enough to force applications to run over non-standard ports (e.g., RDP, SSH). In order to enforce application-specific firewall policies where ports are increasingly irrelevant, your next firewall must assume that any application can run on any port. In addition, application platform developers, such as Google®, Facebook®, Salesforce® and Microsoft®, provide users with a rich set of features and functions that help to ensure user loyalty but may represent very different risk profiles. For example, WebEx® is a valuable business tool, but using WebEx desktop sharing to take over your employees' desktop from an external source may be an internal or regulatory compliance violation. Other examples may be Gmail and Google Hangouts. Once a user is signed into Gmail, which may be allowed by policy, they can easily switch context to Google Hangouts audio and video, which may not be allowed. Your next firewall must be able to recognize and delineate individual features and functions so that an appropriate policy response can be implemented.

Requirements: You must assume that any application can run on any port and your next firewall must classify traffic by application on all ports all the time, by default. The firewall must also continually classify each application, monitoring for changes that may indicate when a different function is being used. The concept of "once and done" traffic classification is not an option, as it ignores the fact that these commonly used applications share sessions and support multiple functions. If a different function or feature is introduced in the session, the firewall must note it within the state tables and perform a policy check. Continual state tracking to understand the different functions that each application may support, and the different associated risks, is a critical requirement for your next firewall.

2 Your next firewall must identify and control users reliably

Business case: Employees, customers and partners connect to different repositories of information within your network, as well as to the internet, to perform various aspects of their jobs. These people and their many devices represent your network's users. It's important to your organization's risk posture that you're able to identify who they are beyond IP address, and the inherent risks they bring with them based on the particular device they're using, especially when security policies have been circumvented or new threats have been introduced to the organization. In addition, users are constantly moving to different physical locations and using multiple devices, operating systems and application versions to access the data they need. Subnet IP is mapped only to the physical location, not to individual users. This means that

if users move around – as they tend to do now, even within the office – policy doesn't follow them.

Requirements: User and group information must be directly integrated into the technology platforms that secure modern organizations. Your next firewall must be able to get the user identity from multiple sources, including VPN, WLAN access controllers, directory servers, email servers and captive portal. Knowing who is using each of the applications on your network, and who may have transmitted a threat or is transferring files, strengthens security policies and reduces incident response times. The firewall must allow policies to safely enable applications based on users or groups of users, in either outbound or inbound directions (for example, allow only the IT department to use tools such as SSH, telnet and FTP). With user-based policies, policy follows the users no matter where they go – headquarters, branch offices or at home – and whatever device they may use.

Your next firewall must identify and control security evasion tools

Business case: A small number of the applications on your network may be used to purposely evade the very security policies you have in place to protect your organization's digital assets. Two classes of applications fall into the category of security evasion tools – those that are expressly designed to evade security (e.g., external proxies, non-VPN-related encrypted tunnels) and those that can be adapted to easily achieve the same goal (e.g., remote server/desktop management tools).

External proxies and non-VPN-related encrypted tunnel applications are specifically used to circumvent the in-place security controls using a range of evasion techniques. These applications have no business value to your network as they are designed to evade security, introducing unseen business and security risks.

Remote server/desktop management tools, such as Microsoft Remote Desktop Services and Teamviewer, are typically used by support and IT professionals to work more efficiently. They also are frequently used by employees to bypass the firewall, establishing connections to their home or other computer outside of the network.

To be clear, not all of these applications carry the same risks – remote-access applications have legitimate uses, as do many encrypted tunnel applications. However, these same tools are increasingly being adopted by attackers as part of ongoing, persistent attacks. Without the ability to control these security evasion tools, organizations cannot enforce their security policies, exposing themselves to the very risks they thought their controls mitigated. The Verizon Data Breach Investigation Report mentions that proxy servers masking true IP addresses are a hurdle to finding evidence for a cyberattack prosecution.

Requirements: There are different types of circumvention applications – each using slightly different techniques. There are both public and private external proxies (see proxy.org for a large database of public proxies) that can use both HTTP and HTTPS. Private proxies are often set up on unclassified IP addresses (e.g., home computers) with applications like PHProxy or CGIProxy.

Remote access applications, such as Microsoft Remote Desktop Services, Teamviewer or Citrix® GoToMyPC® have legitimate uses, but due to the associated risk, should be managed more closely. Most other circumventors (e.g., Ultrasurf, Tor, Hamachi) have no business-use case for being on your network.

Regardless of your security policy stance, your next firewall needs to have specific techniques to identify and control all of these applications, regardless of port, protocol, encryption or other evasive tactics. One more consideration: Applications that enable circumvention are regularly updated to make them harder to detect and control. So it is important to understand not only that your next firewall can identify these circumvention applications but also to know how often that firewall's application intelligence is updated and maintained.

Your next firewall must decrypt and inspect SSL and control SSH

Business case: 25 to 35 percent of enterprise traffic is SSL and, depending on the industry vertical, the percentage of SSL traffic can reach as high as 70, according to NSS research. Given the increasing adoption of HTTPS for many high-risk, high-reward applications that end users employ (e.g., Gmail, Facebook), and users' ability to force SSL on many websites, your security team has a large and growing blind spot without the ability to decrypt, classify, control and scan SSL-encrypted traffic. Certainly, a next-generation firewall must be flexible enough that certain types of SSL-encrypted traffic can be left alone (e.g., web traffic from financial services or healthcare organizations) while other types (e.g., SSL on non-standard ports, HTTPS from unclassified websites in Eastern Europe) can be decrypted via policy. SSH is used nearly universally and can be configured easily by end users for non-work purposes in the same manner that a remote desktop tool can be used. The fact that SSH is encrypted also makes it a useful tool to hide non-work-related activity.

Requirements: The ability to decrypt SSL is a foundational element – not just because it's an increasingly significant percentage of enterprise traffic but also because it enables a few other key features that would end up incomplete or ineffective without the ability to decrypt SSL. Key elements to look for include recognition and decryption of SSL on any port, inbound or outbound; policy control over decryption; and the necessary hardware and software elements to perform SSL decryption across tens of thousands of simultaneous SSL connections with predictable performance – an additional requirement to consider.

Your next firewall must systematically manage unknown traffic

Business case: Unknown traffic exists in small amounts on every network, yet to you and your organization, it represents significant risks. There are several important elements to consider with unknown traffic. Is it categorized? Can you minimize it through policy control? Can your firewall easily characterize custom applications so they are "known" within your security policy? Does your firewall help you determine if the unknown traffic is a threat?

Unknown traffic is also strongly tied to threats in the network. Attackers are often forced to modify a protocol in order to exploit a target application. For example, to attack a web server, an attacker may need to modify the HTTP header so much that the resulting traffic is no longer identified as web traffic. Such an anomaly can be an early indication of an attack. Similarly, malware will often use customized protocols as part of its command and control model, enabling security teams to root out any unknown malware infections.

Requirements: By default, your next firewall must classify all traffic on all ports – this is one area where the earlier explanation about architecture and the security control model becomes very

important. Positive (default deny) models classify everything; negative (default allow) models classify only what they're told to classify. Classifying everything is only a small part of the challenge that unknown traffic introduces. Your next firewall must give you the ability to see all unknown traffic, on all ports, in one management location and quickly analyze the traffic to determine if it is (1) an internal or custom application, (2) a commercial application without a signature, or (3) a threat. Additionally, your next firewall must provide you with the necessary tools to not only see the unknown traffic but to systematically manage it by controlling it via policy, creating a custom signature, submitting a commercial application PCAP for further analysis, or performing a forensic investigation to determine if it is a threat.

Your next firewall must protect your network from known and unknown threats in all applications and on all ports

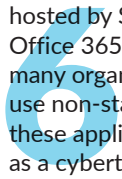

Business case: Organizations continue to adopt a wide range of applications to enable the business – they may be hosted internally or outside of your physical location. Whether it's hosted by SharePoint®, Box.com, Google Docs™, Microsoft Office 365™, or an extranet application hosted by a partner, many organizations require the use of an application that may use non-standard ports, SSL or can share files. In other words, these applications may enable the business, but they can also act as a cyberthreat vector. Furthermore, some of these applications (e.g., SharePoint) rely on supporting technologies that are regular targets for exploits (e.g., IIS, SQL Server). Blocking the application isn't appropriate, but neither is blindly allowing the applications and the (potential) associated business and cybersecurity risks.

This tendency to use non-standard ports is highly accentuated in the world of malware. Since malware resides in the network, and most communication involves a malicious client (the malware) that communicates with a malicious server (command and control), the attacker has full freedom to use any port and protocol combination.

Requirements: A critical part of safe enablement is allowing an application and scanning it for threats. First, your next firewall must identify the application (regardless of port or encryption), determine the functions you may want to allow or deny, and protect the organization from known and unknown threats – exploits, viruses/malware or spyware. Second, it must do so in a way that is scalable. Looking for threats can be performance-intensive, especially at scale when protecting thousands of end users. Offloading some of these detection capabilities to the cloud is one way to scale. Third, the firewall must automatically update itself in near-real time to protect from the newest threats being discovered globally.

Your next firewall must deliver consistent controls to all users, regardless of location or device type

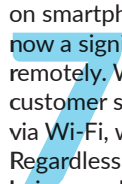

Business case: Your users are increasingly outside the four walls of the organization, oftentimes accessing the corporate network on smartphones or tablets. Once the domain of road warriors, now a significant portion of your workforce is capable of working remotely. Whether working from a coffee shop, home or a customer site, your users expect to connect to their applications via Wi-Fi, wireless broadband, or by any means necessary. Regardless of where the user is, or even where the application being employed might be, the same standard of firewall control should apply. If your next firewall enables application visibility and control over traffic inside the four walls of the organization, but not outside them, it misses the mark on some of the riskiest traffic.

Requirements: Conceptually, this is simple – your next firewall must have consistent visibility and control over traffic, regardless of where the user is. This is not to say that your organization will have the same policy for both; for example, some organizations might want employees to use Skype™ when on the road, but not inside headquarters, where others might have a policy that states users may not download Salesforce.com attachments unless they have hard-disk encryption turned on. This should be achievable on your next firewall without introducing significant latency for the end user, undue operational hassle for the administrator, or significant cost for the organization.

Your next firewall must simplify network security with the addition of application control

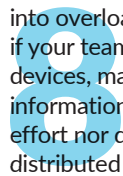

Business case: Many organizations struggle with incorporating more information feeds, more policies, and more management into overloaded security processes and people. In other words, if your team can't manage what it's already got, adding more devices, managing interfaces along with associated policies and information doesn't help you reduce your team's administrative effort nor does it help reduce incident response time. The more distributed the policy is (e.g., a port-based firewall allows port 80 traffic, IPS looks for and blocks threats and applications, and a secure web gateway enforces URL filtering), the harder it is to manage that policy. Which policy does your security team use to enable WebEx? How do they determine and resolve policy conflicts across these different devices? Given that typical port-based firewall installations have thousands of rules, adding thousands of application signatures across tens of thousands of ports is going to increase complexity by several orders of magnitude.

Requirements: Your business is based on applications, users and content, and your next firewall must allow you to build policies that directly support your business initiatives. Shared context across the application, user, and content in all aspects – visibility, policy control, logging and reporting – will help you simplify your security infrastructure significantly. A firewall policy based on port and IP address, followed by separate policies for application control, IPS and anti-malware will only complicate your policy management process and may end up inhibiting the business.

Your next firewall must deliver the same throughput and performance with application control fully activated

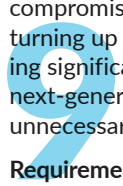

Business case: Many organizations struggle with the forced compromise between performance and security. All too often, turning up security features on your firewall means accepting significantly lower throughput and performance. If your next-generation firewall is built the right way, this compromise is unnecessary.

Requirements: The importance of architecture is obvious here too – but in a different way. Cobbling together a port-based firewall and other security functions from different technology origins usually means there are redundant networking layers, scanning engines and policies – which translates to poor performance.

From a software perspective, the firewall must be designed to do this from the beginning. Furthermore, given the requirement for

computationally intensive tasks (e.g., application identification, threat prevention on all ports and SSL decryption) performed on high-traffic volumes and with the low tolerance for latency associated with critical infrastructure, your next firewall must have hardware designed for the task as well – meaning dedicated, specific processing for networking, security and content scanning.

Your next firewall must deliver the same firewall functions in both a hardware and virtualized form factor

Business case: The explosive growth of virtualization and cloud computing introduces new security challenges that are difficult or impossible for legacy firewalls to manage effectively due to inconsistent functionality, disparate management, and a lack of integration points with the virtualization environment. In order to protect traffic flowing in and out of the data center within your virtualized environments and in the public cloud, your next firewall must support the same functionality in both a hardware and virtualized form factor.

Requirements: The dynamic setup and tearing down of applications within a virtualized data center and the public cloud exacerbates the challenges of identifying and controlling applications using an approach based on port and IP addresses. In addition to delivering the features already described here to manage effectively in both hardware and virtualized form factors, it is imperative that your next firewall provide in-depth integration with the virtualization environment to streamline the creation of application-centric policies as new virtual machines and applications are established and taken down. This is the only way to ensure you can support evolving data center architectures with operational flexibility while addressing risk and compliance requirements.

Firewalls Should Safely Enable Applications – and Business

Your users continue to adopt new applications and technologies, oftentimes to get their jobs done but with little regard to the associated business and security risks. In some cases, if your security team blocks these applications, it may hinder your business.

Applications enable your employees to get their jobs done and maintain productivity in the face of competing personal and professional priorities. Because of this, safe application enablement is increasingly the correct policy stance. To safely enable applications and technologies on your network and the business that rides atop them, your network security teams need to put in place the appropriate policies governing use and the controls capable of enforcing them.

About Palo Alto Networks

Palo Alto Networks® is the next-generation security company, leading a new era in cybersecurity by safely enabling applications and preventing cyber breaches for tens of thousands of organizations worldwide. Because our platform was built from the ground up with breach prevention in mind and important threat information being shared across security functions system-wide, and architected to operate in modern networks using new technology initiatives like cloud and mobility, customers get better security than they would from legacy or point security products and they realize a better total cost of ownership.
Find out more at **www.paloaltonetworks.com**.

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