FIREWALL BUYER'S GUIDE

THE DEFINITIVE GUIDE FOR EVALUATING ENTERPRISE NETWORK FIREWALLS

paloalto NETWORKS®

EVALUATING FIREWALLS FOR TODAY'S NEEDS

Even with more advanced features and supposed higher throughput than ever before, firewalls are not able to keep up with modern demands or advanced threats. Users are more distributed than ever, and so is data. Threats are changing rapidly, and traditional defense mechanisms that rely on layered security do not work.

INTRODUCTION

Your network is more complex than ever before. Your employees are accessing any application they want, using work or personal devices. Often times, 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 their new job. Adding another layer of complexity is the 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 always helpful to focus on the fundamentals. 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 simply states: Allow what you want; all else is implicitly denied.

Over time, these fundamental functions were nullified by the very traffic they were meant to control. Applications evolved at a faster pace than the firewall. As a result, these firewalls have trouble exerting the control needed to protect 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 using these applications to 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[1]. Employees may use some of them for work purposes; others will be for purely personal use. These applications are often highly extensible and include features that introduce unwarranted risk. These applications represent both business and security risks and your challenge will be how to strike an appropriate balance of blocking some and securely enabling others
- **Core business applications:** These are the applications that run your business; they house your most valued assets. They include databases, directories and ERP applications in your data center, and applications such as Salesforce® and Workday® in the cloud. This group of applications is heavily targeted by cyberattackers who use multifaceted attacks. Your challenge will be how best to isolate and protect them from stealthy attacks that use common evasion techniques to easily evade your firewall and IPS.
- **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 are using a wide range of non-standard ports. Many of the applications that use SSL never use port 443, while others hop ports.

To address these challenges, there has been an increased focus on the fundamentals of the firewall. Every network firewall vendor is 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 at which all the traffic flowing across your network is seen, 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.

REVOLUTION, NOT EVOLUTION

There is too much traffic, too many applications, and too little tolerance for negative performance impacts to keep adding devices and new software “modules” that purport to help analyze traffic.

NEXT-GENERATION FIREWALLS DEFINED

The next-generation firewall is well defined by Gartner® as something new and enterprise-focused, "incorporating full-stack inspection to support intrusion prevention, application-level inspection and granular policy control." 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, 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's response options be broader than just "allow" and "deny"?
- Will your network be protected from threats and cyber-attacks – both known and unknown?
- Can you systematically identify and manage unknown traffic?
- Can you implement the desired security policies without compromising performance?

Will the administrative efforts 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?

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.

NEXT-GENERATION FIREWALLS

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 inline deployment.

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 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 specifically look for and block threats or unwanted applications and to 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 increased administrative effort associated with multiple policies and duplicate log databases.

The remainder of this Buyer's Guide is broken down into three sections. The first section introduces *10 Things Your Next Firewall Must Do*, which 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. The remaining sections delve into how these 10 things should be used to select a vendor through the request for proposal (RFP) process and how you should physically evaluate the firewall solution.

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 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:

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. The concept of any application on any port is one of the fundamental changes in the application landscape that is driving the migration from port-based firewalls to next-generation firewalls. Any application on any port also underscores why a negative control model can't solve the problem. If an application can move to any port, a product based on negative control would require beforehand knowledge or have to run all signatures on all ports, all the time.

Requirements: This one is simple. 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. Traffic classification on all ports will be a recurring theme throughout the remaining items; otherwise, port-based controls will continue to be outwitted by the same techniques that have plagued them for years.

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 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 RDP 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. Cyberattackers know these applications are commonly used, and there are publicly documented cases in both the Verizon Data Breach Investigations Report (DBIR) and the Mandiant® report where these remote access tools were executed in one or more of the attack phases.

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 their 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.

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 such as PHProxy or CGIProxy. Remote access applications such as RDP, TeamViewer or 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 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 tactic. One more consideration: Applications that enable circumvention are regularly updated to make them harder to detect and control. So it is important to understand that your next firewall should identify these circumvention applications; it is also important 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: Currently, SSL accounts for about 14 percent of global application traffic bandwidth[2]. 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 easily configured by end users for non-work purposes in the same manner that a remote desktop tool is 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 and 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. Additional requirements to consider are the ability to identify and control the use of SSH. Specifically, SSH control should include the ability to determine if it is being used for port forwarding (local, remote, X11) or native use (SCP, SFTP and shell access). Knowledge of how SSH is being used can then be translated into appropriate security policies.

Your next firewall must provide application function control.

Business case: Application platform developers such as Google®, Facebook, Salesforce® or 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, allowing 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. Another example may be Google Mail (Gmail) and Google Talk (Gtalk). Once a user is signed into Gmail, which may be allowed by policy, they can easily switch context to Gtalk, 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: 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.

Requirements: Your next firewall must 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.

SAFE APPLICATION ENABLEMENT

To safely enable applications and technologies – and the business that rides atop them – network security teams need to put in place the appropriate policies governing use and the controls capable of enforcing them.

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) communicating to a malicious server (command and control), then the attacker has full freedom to use any port and protocol combination he chooses. In fact, in a recent three month analysis, 97 percent of all unknown malware delivered via FTP used completely non-standard ports.

Requirements: Part of safe enablement is allowing an application and scanning it for threats. These applications can communicate over a combination of protocols (e.g., SharePoint uses CIFS, HTTP and HTTPS, and requires a more sophisticated firewall policy than "block the application.") The first step is to identify the application (regardless of port or encryption), determine the functions you may want to allow or deny, and then scan the allowed components for any of the appropriate threats – exploits, viruses/malware, or spyware, or even confidential, regulated, or sensitive information.

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.

8

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.

9

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 – 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, etc.) 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

10

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 tear down of applications within a virtualized data center exacerbates the challenges of identifying and controlling applications using a port- and IP address-centric approach. In addition to delivering the features already described in *10 Things Your Next Firewall Must Do* 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.

ENABLING YOUR BUSINESS

In today's always-connected world, controlling applications is more than merely allowing or denying; it is about safely enabling applications to the betterment of the business.

FIREWALLS SHOULD SAFELY ENABLE APPLICATIONS AND BUSINESS

Your users continue to adopt new applications and technologies, often times 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 help your employees 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 also the controls capable of enforcing them.

10 Things Your Next Firewall Must Do describes the critical capabilities that will allow organizations to safely enable application usage and ultimately, the business. The next steps are to translate those requirements into actionable steps by selecting a vendor through an RFP process and formally evaluating solution offerings, ultimately resulting in the purchase and deployment of a next-generation firewall.

USING THE RFP PROCESS TO SELECT A NEXT-GENERATION FIREWALL

Typically, when selecting firewalls, IPS or other critical security infrastructure components, organizations will utilize an RFP as a means of ensuring that the specific needs are addressed. According to the Gartner's Magic Quadrant for Enterprise Firewalls 2016, "Enterprises with traditional firewalls seek to have firewalls that have application and user visibility, and to require enforcement options in their next refresh." As new deployment opportunities occur, organizations should expand their RFP selection criteria to include application visibility and control offered by next-generation alternatives. The previous section established the 10 key requirements your next firewall must do. This section will translate those requirements into tools you can use to identify and select a next-generation firewall.

Firewall Architecture and Control Model Considerations

There are many elements to consider when evaluating how effectively a vendor can deliver application visibility and control in the firewall. The firewall architecture, specifically its traffic classification engine, will dictate how effectively it can identify and control applications, instead of just ports and protocols. As mentioned earlier, the very first thing a new firewall of any type must do is accurately determine what the traffic is and then use that result as the basis for all security policy decisions.

In this model, the firewall policies are traditional positive control (block all, except that which you expressly allow). A positive model means you can control and enable applications, which is a critical requirement in the always-on, always connected-world that businesses are faced with today. Bolting on IPS-like elements that look for applications means that a negative control model is used (allow all, except that which is expressly denied by the IPS). A negative model means you can only block applications. The differences are analogous to turning the lights on in a room to see and control everything (positive) vs. using a flashlight in a room to see and control only what you are looking at (negative). Using this add-on to identify and block "bad" events is simply a patch and not the full solution because it is designed to look only at a partial set of traffic to avoid impeding performance, and cannot cover the breadth of cyberattacks and applications.

Application Visibility and Control

The RFP must determine the details around how the firewall architecture facilitates the identification and control of the entire spectrum of applications including business, personal or other, as well as protocols, no matter which port, SSL encryption or other evasive technique is in use. Consider the following questions and statements when issuing an RFP for next-generation firewalls.

- Many applications can evade detection using non-standard ports, port hopping, or by being configured to run on a different port.
- In traffic classification, is the first task that the firewall execute based on application identity or the network port?
  - Are the application identification mechanisms part of the core firewall traffic classification (i.e., enabled by default)?
  - Are the application identification mechanisms dependent on the application's standard port?
  - Can the signatures be applied to all ports and is the process automatic or manually configured?
- When traffic first hits the device, is it first classified based on port (this is port 80, therefore it is assumed to be HTTP) or application (this is Gmail)?
- Describe in detail how the firewall can accurately identify applications.
  - Which mechanisms, besides signatures, are used to classify traffic?
  - Describe the breadth of application and protocol decoder use.
  - How are SSL and SSH decryption and control implemented?
  - Are the traffic classification mechanisms applied equally across all ports?
- Which mechanisms are used to detect purposely evasive applications, such as UltraSurf or encrypted P2P?
- Is application identification actually performed in the firewall, or is it performed in a secondary process, after port-based classification?
  - What are the three key advantages of the supported architectural approach?
- Is application state tracked, and if so, how is it utilized to ensure consistent control of the application and associated secondary functions?
  - Give three examples of how application state is used in policy control.
- Is the identity of the application the basis of the firewall security policy, or is application control treated as a secondary policy element?
- How often is the application database updated and is it a dynamic update or a system reboot upgrade?
- In virtualized environments, describe how the traffic is classified throughout the virtual machine (east-west, north-south).
  - Describe the points of integration within the virtualized environment.
  - Describe the process of building security policies for newly created virtual machines.
  - Describe the features available to track virtual machine moves, adds and changes.
  - Describe the features available for integration with automation and orchestration systems.

Controlling Evasive Applications, SSL and SSH

A wide range of applications can be used to circumvent security controls. Some, such as external proxies, and non-VPN related encrypted tunnels are designed with circumvention as a goal. Others, such as remote server/desktop management tools have evolved to where non-IT or non-support staff employees use them to circumvent control mechanisms. As a means of security, SSL is becoming a standard configuration for many end-user applications, yet the problem arises when the use of SSL may be masking inbound threats or outbound data transfer. Today, SSL accounts for about 14 percent of global application traffic bandwidth[3] in some way. So it is important to determine the respective next-generation firewall vendors that address this category of applications. Consider the following questions and statements when issuing an RFP for next-generation firewalls.

- Describe the process by which SSL encrypted applications are identified on all ports, including non-standard ports.
- What policy controls are available to selectively decrypt, inspect, and control applications that are using SSL?
- Is bi-directional SSL identification, decryption, and inspection supported?
- Is SSL decryption a standard feature, or an extra cost? And is a dedicated device required?

- SSH is a commonly used tool for IT, support, and tech-savvy employees as a means of accessing remote devices.
  - Is SSH control supported and if so, describe the depth of control.
- What mechanisms are used to identify purposely evasive applications such as UltraSurf or Tor?
- Describe how the product can automatically identify a circumventer that is using a non-standard port.

Policy-Based Application Enablement

In today's always-connected world, controlling applications means more than merely allowing or denying; it is about safely enabling applications to the betterment of the business. Many "platforms" (Google, Facebook, Microsoft) make different applications available to the user after their initial login. It is imperative that you determine how the firewall vendor monitors the state of the application, detects changes in the application, and classifies the change in state. Consider the following questions and statements when issuing an RFP for next-generation firewalls.

- Is stateful-inspection traffic classification performed separately, prior to application identification? Once an application is identified, describe how the changes in application state are monitored, tracked and used within policy.
- Describe how the application database hierarchy (flat, multi-level, other) exposes functions within the parent application for more granular enablement policies.
- Describe the levels of control that can be exerted over individual applications and their respective functions:
  - Allow
  - Allow based on application, application function, category, subcategory, technology or risk factor
  - Allow based on schedule, user, group or port
  - Allow and scan for viruses, application exploits, spyware and drive-by downloads
  - Allow and shape/apply quality of service (QoS) controls
  - Deny
- Can port-based controls be implemented for all applications in the application database so that an administrator can enforce, by policy, the application and port relationship? For example:
  - Force Oracle® database developers over a specific port or range of ports?
  - Ensure the IT personnel are the only ones allowed to use SSH and RDP.

PERFORMANCE MATTERS

It is critical to determine the performance on the network when all security features are enabled and analyzing a real-world mix of traffic.

  - Detect and block malware within the application, even if it is on a non-standard port.
- List all the enterprise identity repositories supported for user-based controls.
- Is an API available for custom or non-standard identity-infrastructure integration?
- Describe how policy-based controls are implemented by users and groups for terminal services environments.
- Describe any differences in application enablement options for hardware and virtualized instances.

Systematically Managing Unknown Applications

Every network has some unknown application traffic; the typical source is an internal or custom application, but it may also be an unidentified commercial application or, worst case, some malicious code. The key elements to determine through the RFP and the evaluation process are a specific description of how the vendor enables you to systematically manage the unknown traffic, which represents a higher business and security risk. Consider the following questions and statements when issuing an RFP for next-generation firewalls.

- Provide a detailed description of how unknown traffic can be identified for analysis.
- Are the mechanisms used for analysis part of the standard feature set, or are they secondary, add-on products?
- What, if any, actions can be taken on unknown traffic (allow, deny, inspect, shape, etc.)?
- Describe the recommended best practices for managing unknown application traffic.
  - Can it be controlled by policy, in the same manner as an officially supported applications (e.g., allow, deny, inspect, shape, control by user, zone, etc.)?
  - Can the internal traffic be “renamed”?
  - Can a custom application signature be created?
- What is the process for submitting requests for new or updated application signatures?
- Once an application is submitted, what is the SLA turnaround time?
- What mechanisms are available to determine if the unknown traffic is malicious code?

Threat Prevention

Threats are increasingly tied to a variety of applications both as vectors for exploits and infection as well as ongoing command and control of infected devices. For this reason, analysts are consistently recommending that enterprises consolidate traditional IPS and threat prevention technologies as a component of the next-generation firewall. Consider the following questions and statements when issuing an RFP for next-generation firewalls.

- Describe all threat prevention mechanisms in use (IPS, antivirus, antispyware, URL filtering, data filtering, etc.).
- How are these threat prevention mechanisms licensed?
- Describe which threat prevention mechanisms are developed in house or obtained via a third party or service.
- How are threats prevented that are embedded within applications on non-standard ports?
- Is application identification information integrated or shared with the threat prevention technologies? If so, describe the level of integration.
- Describe which threat prevention disciplines (IPS, AV, etc.) are port-based as opposed to application-based.
- Can the threat prevention engine scan inside of compressed content such as ZIP or GZIP?
- Can the threat prevention engine scan within SSL encrypted content?
- Describe how the firewall can detect and defend against custom or polymorphic malware.
  - Which mechanisms are used to block the malware?
- Describe the threat prevention research and development process.
- Which process is used to discover unknown threats? Is the process scalable for thousands of users?
- If a threat is found that was previously unknown, how do all firewalls get protected from this threat? How soon do all firewalls get protected?
- Does the firewall protect our network from unknown threats that are discovered in other customer environments? How?

Securing Remote Users

Modern network users assume the ability to connect and work from many locations beyond the traditional perimeter of the network. These users must remain protected even in instances where they are beyond the network perimeter, using a PC, a smartphone or a tablet. The goal of this section is to determine what capabilities are available to secure these remote users and how this level of protection differs when the user is on or off of the physical network. Consider the following questions and statements when issuing an RFP for next-generation firewalls.

- Provide a detailed description, including all necessary components, of the available options for securing remote users.
- If a client component is included, how is it distributed?
- Describe the sizing requirements. How many users can be supported simultaneously?
- Is the remote user security feature set transparent to the client?
- Describe how policy control over remote users is implemented (e.g., in the firewall policy, in a separate policy/device, other).
- List all features and protections provided by the remote capabilities (SSL, application control, IPS, etc.)
- Can your firewall keep users connected in order to ensure consistent policy enforcement regardless of location?
- How do you address mobile device users? Will you be able to provide consistent policy enforcement when users are on external networks as well as internal wireless networks?
- Can the firewall address BYOD issues such as providing a way to safely enable both corporate and personally owned laptops, phones and tablets?

Management

Management is a critical element for implementing effective network security. In moving to your next firewall, a key goal must be to simplify security management wherever possible by adding application visibility and control. Consider the following questions and statements when issuing an RFP for next-generation firewalls.

- Does device management require a separate server or device?
- Describe all of the management options that are supported: Command line interface (CLI)? Browser? Software client? Centralized server?
    - For each of the management alternatives supported, describe how much effort is required to move from one management technique to another.
- Describe the centralized management architecture and deployment options.
    - What visibility tools, outside of the log viewer and reporting, are available to enable a clear picture of the applications, users and content traversing the network?
        - Are the visibility tools included as part of the base functionality, or are they extra cost/added licenses?
        - Are the visibility tools deployed on-box, or are they a separate device/appliance?
- Provide a detailed description of the effort and steps required to begin “seeing a comprehensive view of all application traffic” on the network.
- Can the application policy controls, firewall policy controls, and threat prevention features all be enabled in a single rule in the firewall policy editor?
- Describe the logging and reporting capabilities – are they on-box and if so, what is the performance degradation when logging is enabled?
    - Is full log analysis available on-box, or is it an extra cost, added license or separate device?
- Are fully customizable reporting tools available to understand how the network is being used and to highlight changes in network usage?
    - Are they an extra cost, added license, or separate device?
- Describe how management access is ensured when the device is under heavy traffic load.
- Describe the relationship between individual device and centralized management of multiple devices.
- Describe the differences in management between hardware and virtualized instances.

Performance

Real-world performance is a critical component of a security deployment. Application control requires a far deeper investigation of traffic than port-based firewalling and as such, is far more computationally intensive. Adding threat inspection and policy control to that same traffic only adds to the processing burden placed on the firewall. It is critical to determine the performance on the network when all security features are enabled and analyzing a real-world mix of traffic. Consider the following questions and statements when issuing an RFP for next-generation firewalls.

- Verify whether the product is software-based, an OEM server, or a purpose-built appliance.
- Investigate the hardware architecture to confirm appropriate processing power for continuous application-level traffic classification and inspection.
- Describe the traffic mix used to produce the published performance metrics for:
    - Firewall + logging
    - Firewall + application control
    - Firewall + application control + threat prevention
- What is the rated throughput for:
    - Firewall + logging?
    - Firewall + application control?
    - Firewall + application control + threat prevention?

Networking

Your firewall must provide a flexible networking architecture that includes support for dynamic routing, switching and VPN connectivity, and enables you to deploy the firewall into nearly any networking environment. Consider the following questions and statements when issuing an RFP for next-generation firewalls.

- What deployment modes are supported on the firewall?
- During evaluation, can the firewall passively monitor traffic flows across a network by way of a switch SPAN or mirror?
- When no switching or routing is needed, can the firewall be installed transparently on a network segment by binding two ports together (e.g. virtual wire)?
- Can the firewall provide switching between two networks (Layer 2)?
- Can the firewall route traffic between multiple ports (Layer 3)?
- Can different interfaces be configured in Layer 2, Layer 3 and virtual wire modes based on networking requirements (i.e. mixed mode)?
- Which routing protocols does the firewall support, e.g., RIP, OSPF, BGP?

Additional RFP Considerations

Every organization has varied requirements over and above the items listed within this document. Examples may include company viability, customer references and quality of customer support. The recommended best practices for an RFP are to be very systematic in driving the vendors towards proving that their offering delivers the claimed functionality.

EVALUATING NEXT-GENERATION FIREWALLS THROUGH FORMAL TESTING

Once the final vendor, or the "short-list" of vendors, has been selected via the RFP, the next step is to physically evaluate the firewall using traffic patterns, objects and policies that are accurate representations of the organization's business. This section provides some recommendations on how to physically evaluate a next-generation firewall. The evaluation will give you the ability to see, in a real-world environment, how well a firewall vendor will address the key requirements. Note that the tests suggested below represent a sample of the next-generation firewall functions required, and are meant as guidelines from which a more detailed, step-by-step test plan can be developed.

Application Visibility and Control

The goal of this section is threefold. First, verify that the first task the device under test (DUT) executes is traffic classification based on the application identity, not the network port. Second, verify that the DUT classifies applications regardless of evasive tactic, such as hopping ports, non-standard ports, or other evasive tactic, as a means of enhancing accessibility. Third, determine that the application identity becomes the basis of the firewall policy, as opposed to an element within a secondary policy.

Application Identification

- Confirm that the firewall can identify various applications. The ideal way to execute this test is to deploy the DUT in tap or transparent mode on the target network.
- Verify that the DUT correctly identifies the application traffic using both graphical, summary level tools and forensic investigative tools.
    - Determine the amount of administrative effort associated with this task.
- Evaluate the steps required to initially enable application identification. How quickly can a user set a policy and begin "seeing" application traffic? Are there extra steps required to gain visibility into applications that hop ports or use non-standard ports?

Identify Applications That Port Hop or Use Non-Standard Ports

- Verify that the firewall can identify and control applications running on ports other than the application's default port. For example, SSH on port 80 and Telnet on port 25.
- Confirm that the firewall can identify applications that hop ports using a known port-hopping application such as Microsoft Lync, Skype, or one of the many P2P applications.

Application Identity as a Basis of Firewall Security Policy

- Confirm that when creating a firewall policy, the application, not the port, is used as the primary policy element.
    - Does the application control policy require a port-focused rule first?
    - Is the application control element a completely separate policy editor?
- Create a policy to allow certain applications and block others, and verify that the applications are controlled as expected.
- Does an application-based policy support the deny-all-else premise that a firewall is based upon?

Identify and Control Circumventors

- Confirm that the DUT can identify and control a range of applications that are used to circumvent security controls. Applications that fall into this group include external proxies (PHproxy, Kproxy), remote-desktop access (RDP, LogMeIn!, TeamViewer, GoToMyPC) and non-VPN related encrypted tunnels (Tor, Hamachi, UltraSurf).
- Confirm that each of the circumventors is identified accurately during the test.
- Verify that all the circumventors can be blocked, even when they are enabled on a non-standard port.

Identify and Control Applications Using SSL or SSH

With more and more applications using SSL encryption and the use of SSH for alternative purposes, you need to evaluate the ability to identify and control application using SSL and SSH.

- Verify that the DUT can identify and decrypt applications that are known to use SSL encryption.
- Confirm that the DUT can identify, decrypt, and apply security policy to the decrypted application.
- Confirm that the DUT allows exempting some traffic from decryption based zone, user, user group or URL category.
- Validate that if the decrypted application is "allowed," it will be re-encrypted and sent on its way.
- Confirm the ability to perform inbound and outbound SSL decryption and inspection.
- Verify SSH is identified accurately, regardless of port.
- Validate that SSH control delineates between port forwarding (local, remote, X11) and native use (SCP, SFTP and shell access).

Identify and Control Applications Sharing the Same Connection

Determine if the application classification mechanisms continually monitor the state of the application, looking for changes in the application, and more importantly, if the change in state is classified correctly. Many "platforms" (Google, Facebook, Microsoft) enable different applications once the user initially logs in. Tracking that change in the application state is a critical component to a next-generation firewall.

- When using an application, such as WebEx or SharePoint, first confirm that the DUT identifies the initial application (as WebEx or SharePoint).
- Without logging out of the application, switch to a separate function (WebEx Desktop Sharing, SharePoint Admin, SharePoint Docs), and validate that the change in state is tracked and that the new application/function is indeed correctly identified.
- Validate policy control and inspection over the application function.

Application Function Control

Determine the ability for the DUT to identify and control specific functions within an application. Function-level control is critical to enabling the use of an application, yet exerting some level of control to address the associated business and security risks. File transfer is a common example, but other examples may include administrative functions, VoIP features, social media posting, and chat capabilities within the parent application.

- Confirm that the DUT provides visibility into the application hierarchy (both core application and additive functions).
- Verify file transfer function control by identifying and controlling an application that supports file transfers.
- Confirm the DUT's ability to block file upload/download by application and file type. For example, the ability to prevent a user from transferring a Word document using a web-based email application.

Systematically Manage Unknown Traffic

All networks have a small amount of unknown traffic, and you need to determine how quickly you can identify what the unknown traffic is and take an appropriate action.

- Validate that visibility into unknown traffic is available and includes, at a minimum:
    - The volume of traffic
    - User and/or IP addresses
    - The port in use
    - The associated content – file, threat, other
- What is the level of effort associated with investigating unknown traffic?
- Can you set a firewall policy (allow, block, inspect, etc.) for unknown traffic?
- Confirm the options available to more accurately identify and control the unknown application traffic.
    - Can the traffic be "renamed"?
    - Can the user create a custom identification mechanism?
    - Will the vendor provide a custom identification mechanism and if so, how quickly?

Threat Prevention

To protect your network, you will need to strictly control the exposure to threats and reliably prevent known and unknown threats present within allowed application traffic. You need to test the ability of the DUT to enforce security in a real-world environment, including previously unknown threats; threats carried by applications running on non-standard ports; and threats obscured by compression, all the while meeting enterprise performance requirements.

- Confirm the granularity of the threat prevention profiles. Are they global (only) or can they be set individually, based on the traffic, threat, user, etc.?
- Verify that threat prevention techniques (IPS, malware, content filtering) are consistently applied to applications (and threats) that can use non-standard ports. This means that not only should the DUT control applications on non-standard ports but the threat prevention should stop threats traveling over non-standard ports as well.
- Verify that the DUT detects malware and unapproved files even when compressed such as with ZIP or GZIP.
- Determine the process for identifying and blocking unknown malware.
- Verify the performance of the DUT with all threat prevention enabled to ensure the real-world applicability of threat prevention features.

Securing Remote Users

First, determine if the DUT can protect remote users by applying the same policy as used internally; and second, determine the management effort and deployment complexity.

- Verify that the DUT can protect remote users using more than an SSL VPN connection or a backhaul connection.
- Confirm ease-of-deployment and management by establishing a remote group of users and deploying a test policy.
- Can the DUT provide policies based on the type of device?
- Can the DUT protect against mobile malware as well as mobile OS vulnerabilities?
- Can the DUT provide application control for mobile applications?
- Close the test out by monitoring remote users via the log viewer.

ATTACK SURFACE REDUCTION

To protect your network, you will need to both strictly control the exposure to threats and reliably prevent threats present within allowed application traffic.

Management

You need to look at the complexity of managing the DUT in terms of separate devices, as well as the difficulty (number of steps, clarity of UI, etc.) of the task at hand.

- Confirm the management methodology of the DUT. Does individual device management require a separate device or server? Can the DUT be managed via a browser, or is a "fat client" required?
- Verify the availability of visualization tools that provide network intelligence via a summary view of the applications, threats, and URLs on the network.
  - Are logs stored centrally, or in separate function-level databases (e.g., firewall, application control, IPS)?
  - Measure the administrative effort associated with log analysis for both visibility and forensic, incident investigation purposes.
- Validate that application policy controls, firewall policy controls, and threat prevention features can be all enabled from the same policy editor.
  - Is a port-based firewall rule created and applied prior to application-level control?
  - If multiple policies are used (e.g., firewall, application control, IPS), are there any policy reconciliation tools available to find potential policy holes?

Performance With Services Enabled

Application control is far more computationally intensive than traditional port-based firewalling, therefore it is critical to validate that the target DUT can perform adequately when identifying and controlling applications.

- Verify whether the DUT is software-based, an OEM server, or a purpose-built appliance.
- If it is an appliance, investigate the hardware architecture to confirm that appropriate processing power is meeting your network performance requirements when all services are enabled.
- Test it! Evaluate the actual performance in a test environment using traffic patterns that are representative of the target network environment.

Hardware or Virtualized Form Factor Considerations

If target deployment location is a virtualized data center, then you should pick and choose the tests above to ensure that the firewall functionality in a virtualized form factor is adequately tested. For virtualized environments, additional considerations should include:

- What is the process for managing the policy to virtual-machine instance relationship? How many steps are involved?
- Can the same types of policies be created for both physical and virtual instances?
- Are the exact same features supported in both hardware and virtualized instances?
- Verify that DUT can secure all traffic between virtual machines on the same virtualized server.
- Verify that the DUT can deliver application, user and content policies on the same virtual instance.
- Verify that the DUT can continue to enforce policies even with guest virtual-machine migrations.
- Confirm and validate the interaction with the virtualization platform management system.
- Confirm and validate the interaction with automation and orchestration systems.

Networking Considerations

Verify that the DUT supports a flexible networking architecture that enables you to deploy the firewall into nearly any networking environment.

- Test how easy it is to evaluate the firewall in a monitor mode. Can you passively monitor traffic flows across a network by way of a switch SPAN or mirror port, or does the DUT require a network change for an evaluation?
- Test the firewall configuration in multiple modes, e.g., Layer 2, Layer 3, virtual wire (with no routing or switching).
- Test if the DUT supports configuring different interfaces for a different mode, e.g., eth1/1 and eth1/2 in Layer 2, and eth1/3 and eth1/4 in Layer 3 mode.
- Check what routing protocols are supported based on what you need, e.g., RIP, OSPF, BGP.

Additional Evaluation Considerations

The evaluation and testing process for network security products will vary from organization to organization, and in nearly all cases, will expand beyond the scope of this document. Examples may include checking the overall stability of the DUT and testing the responsiveness of customer support. The recommended best practices for a firewall evaluation is to build a specific set of evaluation criteria and put each device through the entire suite of tests, documenting in detail the results so that the final selection can be made in a systematic manner.

SECURE APPLICATION ENABLEMENT WITH NEXT-GENERATION FIREWALLS

At one time, the concept of allowing an employee to use an external or personal application for work-related purposes was unheard of. Today, employees are always online and are continually using the latest applications, oftentimes melding personal and work-related usage. Summarily blocking these applications is equivalent to blocking the business.

10 Things Your Next Firewall Must Do validates the fact that the best location to execute secure application enablement is at the firewall by using the application identity and traditional positive control model (firewall) policies that allow administrators to define, based on the business, which applications are enabled and which are denied. It should be clear after using the tools within this document that attempts to claim secure application enablement using a negative control model, IPS-like, bolt-on approach are unrealistic.

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 – with important threat information being shared across security functions system-wide – and architected to operate in modern networks with new technology initiatives like cloud and mobility, customers benefit from better security than legacy or point security products provide and realize better total cost of ownership.

Find out more at **www.paloaltonetworks.com**.

[1] Palo Alto Networks Application Usage and Threat Report (Research by Unit 42), 2015

[2] Palo Alto Networks Application Usage and Threat Report (Research by Unit 42), 2015

[3] Palo Alto Networks Application Usage and Threat Report (Research by Unit 42), 2015