



2011-2012

**FIREWALL
BUYERS
GUIDE**

The definitive guide for evaluating enterprise network firewalls.



the network security company™

WITH FOREWORD BY IANS

Foreword

Palo Alto Networks Firewall Buyers Guide Report

The threats that enterprise network and security teams face are evolving rapidly, and the security products that they implement play a major role in the success or failure of the organization's overall security strategy. The firewall has been a mainstay of network security for many years, but the needs of organizations are changing rapidly, and the firewall must evolve to meet the challenges of today's dynamic environments. Palo Alto Networks has asked IANS Research, an independent IT security, risk, and compliance research organization, to set the context for this Firewall Buyers Guide by evaluating how aspects of today's network security operations are changing, how user behaviors are shifting and the complexity of threats is increasing, and what developments in next-generation firewall technology are how a renewed focus on the role of the firewall is needed in order to keep up.



Palo Alto Networks: 2011-2012 Firewall Buyers Guide

The definitive guide for evaluating enterprise network firewalls.

The firewall is considered one of the most fundamental of network security controls. Over the last two decades, the firewall has evolved from a simple packet filtering device to a complex system capable of monitoring the state of numerous network traffic sessions simultaneously. Even with more advanced features and higher throughput than ever before, firewalls are having an identity crisis - the threats are changing rapidly, and traditional filtering for ports and IP addresses is no longer adequate to stop them. Enterprise networks are growing more and more complex, with myriad applications and services that network and security teams need to accurately identify and control. This paper will explore the current state of firewall technology, as well as the emerging requirements in today's network security environment that are driving the need for new firewall capabilities.

Threat Sophistication Demands Visibility.

The application landscape has changed dramatically today. Fewer applications are exhibiting traditional network traffic patterns, with standard protocols like FTP, SMTP, and SMB, while more and more applications are starting to use HTTP as a primary communication method, ranging from simple Web sites to complex Web services infrastructure using XML, SOAP, and AJAX. Users are accessing personal and work related email over HTTP, file sharing applications use HTTP, or they hop from port to port, we now have voice traffic using HTTP in addition to SIP and traditional VoIP protocols. While this introduces a revolution in terms of collaboration and development efforts, there is a significant downside to security. The threats organizations are facing today are significantly more widespread and advanced than ever before. Sensitive data breaches are commonplace, with companies like Citi losing more than 360,000 credit card account holders' data and Sony experiencing multiple incidents that total over 100 million compromised records. In many of these breaches, the attackers used sophisticated methods of social engineering and client-side software manipulation that resulted in very stealthy data exfiltration that was difficult, if not impossible, to detect using traditional network security tools. Much of the attackers' outbound traffic used standard ports like 80, 443, and others, which most organizations cannot inspect deeply due to the volume of Web data internal users are accessing at any given time. Indeed, this traffic blends in with the multitude of other Web-enabled applications so well that most organizations have no idea it's there at all.

The ability to match patterns of usage and behavior to detect malicious activity is critical in an age of sophisticated threats, many of which provide little signature other than traffic volumes and time of transfer. The ability to decipher what traffic (at the packet layer) is entering and, perhaps more importantly, leaving the organization is a cornerstone of advanced threat management. Many of the newest and most malicious bot variants affecting organizations today are emulating the most common traffic types to avoid detection of Command and Control (C&C) channels back to botnet controllers. Team Cymru, a network security think-tank and consulting firm chronicled one relatively innocuous example of this type of activity as far back as 2008. They witnessed infected bot clients communicating with a controller server using Base64-encoded strings such as the following¹:

GET

/cgi-bin/get.cgi?data=dmVyPTUmdWlkPTE4MDczMzM2NSZjb25uPSZvcz1YUCZzb2Nrcz0xNTczJmlwPTE5Mi4xNjguMTk3

This encoded string included data about the infected client's IP address, host operating system, and user ID. Newer bots such as Zeus and the TDL family of botnets also use similar encoding and encryption techniques. The TDL family of botnets uses SSL encryption, and also encodes command and control traffic with Base64. In addition, specific URL strings are known to be affiliated with this botnet, including these²:

/data/www/dm_engine/library/classes/DBase.php

/data/www/dm_engine/library/models/mSystems.php

/data/www/dm_engine/public/enginestatusn.php

/data/www/dm_engine/public/index.php



Change Breeds Innovation

Even with more advanced features and higher throughput than ever before, firewalls are having an identity crisis - the threats are changing rapidly, and traditional filtering for ports and IP addresses is no longer adequate to stop them.

As web-based applications, file sharing and social tools usage for both personal and business use explodes inside of organizations, so does the threat of information misuse and misappropriation. Detecting the details of information being shared inside and outside the organization is the only way to effectively control usage and enforce policy for internal users, but this task is often left to content filtering and proxy tools that are easily bypassed or circumvented by connecting to external proxies, using encryption, or tunneling over applications like Tor and others. In addition, traditional content inspection and filtering tools may not be effective at detecting simple obfuscation or encryption techniques that would ideally be found through behavioral analysis of user traffic patterns.

Complexity and Organizational Demands

Traditional firewall technology largely focuses on TCP and UDP ports as the primary identification factor for filtering network traffic. With the proliferation of Web-based applications and services both internally and externally, these "static" traffic identification techniques may prove to be increasingly inflexible, leaving many organizations blind to application-layer traffic that could pose a potential threat. In addition, many administrative consoles for applications and devices use Web-based functionality, on standard ports and others. For obvious reasons, these applications need to be supported, as do many others. This is one of the most critical elements of today's network security situation - business enablement is vitally important, yet network teams may lack the level of visibility and control needed to enable business activity while blocking malicious traffic or functionality that does not adhere to usage policies.

As the list of applications grows and IT teams shrink, the overhead for managing the myriad rules, syntax for these rules across network infrastructure platforms and finding the necessary expertise to coordinate all elements of defense is increasingly difficult. Many organizations accrue firewall rules and policies over many years, with very little in the way of documentation or proper justification for ongoing maintenance of particular rules. When firewall rule complexity and volume grows, several consequences may arise. First, rule overlap and conflict may lead to availability issues or inadvertent blocking of traffic that is necessary for business. This is obviously unacceptable, and as this scenario often brings criticism of the network and IT teams, they are loath to make any significant changes to the devices and their policies. The second major issue that tends to surface from this situation is a reduction in overall security, as "quick fixes" and "ANY:ANY" rules start to become more prevalent over time.

Network security teams are doing more with less. Unfortunately, the complexity of the application and threat landscape is leading to two ongoing trends in network security. First, the changes in the application landscape and competitive pressures have forced many security device manufacturers to add "application inspection" and advanced "anti-malware" capabilities to their products. Although some of this functionality is well designed and implemented, in many cases new problems arise. The features may introduce additional complexity in managing and maintaining the systems themselves, adding additional overhead and troubleshooting time that most operations teams cannot afford. Additionally, many of these features do not really work as advertised - in fact, in many cases, this functionality is nothing more than traditional signature-based intrusion prevention functionality. Although this may be a welcome addition, the performance impact on the device can be significant, which often leads to administrators disabling it.

The second trend is a proliferation of separate technology that performs a narrow function with regard to security. These "point solutions" may work well, but the growth in distinct systems adds managerial overhead, capital costs, and infrastructure complexity that may be too much for today's lean organizations to bear. Although separate controls seems to indicate a strong "defense in depth" posture, many organizations will be unable to manage all these products, and some will likely even become "shelfware" (unused), leading to a gap in network protection.

Complex Security Challenges Demand a New Approach

Paired with the complexity of the threats organizations have to deal with, many firms contend with a mix of traditional and cutting edge tools to secure the perimeter. With each additional piece of infrastructure and the benefits it promises comes a need for vendor-specific expertise, man-hours for management and the threat of adding another link to the security chain, thus weakening it. Converged infrastructure provides a single point at which to audit, apply, and update information security policies. Side benefits include lower operational costs in the short term, as well as lower capital expenditure and maintenance costs in the long term. This convergence could lead to significant cost savings in overall time spent developing policies, applying them, and maintaining changes over a longer period of time.

As threats become more sophisticated, actors more capable and policies more complex and greater in number, the role of the firewall as the centralized control point for information retention and security policy enforcement has never been greater. However, much of today's firewall technology is not able to provide all the necessary capabilities to prevent and detect data leakage (both malicious and inadvertent), malware infection, and cutting edge attacks from both inside and out. For network security to truly evolve, and firewalls to keep pace with the growing complexity of network environments with numerous applications, new functionality is needed.

New firewall capabilities must evolve in three categories:

- **Application identification and control:** This encompasses policy development and design, as well as the necessary parsing and dynamic interpretation of traffic to evaluate rules and apply them consistently even as traffic fingerprints change over time.
- **Enhanced visibility for both internal and remotely originating traffic:** Attackers are intelligent, and are building malware and attack toolkits that use encrypted channels like SSL to carry sensitive data and bot commands. In addition, many of the most critical business applications will require subtle variations of policy for different types of functionality, and firewalls need to be able to "understand" the subtleties of these applications to make proper policy decisions.
- **Improved operations:** First, firewalls must make security and network operations simpler, by consolidating functions currently performed by multiple network security controls and easing management and policy maintenance overhead. Second, the firewall itself must perform consistently at high speeds while still performing deep application inspection tasks.

Most of today's network security infrastructure can do some of this. However, many of the tools organizations are currently using are attempting to add functionality "organically" by building new modules for the existing devices. Neither these modules nor the hardware they run on were designed explicitly to analyze and categorize application traffic at high speeds, though, and enabling them can severely impact the performance of the devices, potentially leading to network availability issues.

Key Requirement: Application Identification and Control

Business application traffic, both Web-based and otherwise, comprises the majority of data flow within enterprise networks today. With the widespread proliferation of Web applications and other complex applications and protocols, what new capabilities do firewalls need in order to monitor and control traffic? First, the firewall must be able to identify key attributes of the applications in use beyond TCP and UDP ports. For example, the Secure Shell (SSH) remote access application commonly uses TCP port 22, but this can be changed with a trivial amount of effort to bypass port-based access controls and establish an encrypted tunnel. More and more malware variants are using HTTP and HTTPS for command and control channels, as mentioned earlier in the paper. In a SANS whitepaper entitled "Analysis of a Simple HTTP Bot," Daryl Ashley describes the behavior of an HTTP-based bot that uses simple encoding of commands. The bot sets the User-Agent header value to "inter easy" and also receives a scrambled Base64 encoded command which means "sleep": `<!-- 2upczxAX`.³

Most network security controls would pass this bot's traffic with no complaints, as it appears to resemble common Web application traffic. If a firewall was capable of analyzing all HTTP and HTTPS traffic and determine that the traffic was anomalous in some way, either based on behavior patterns of browsing, or the unusual request and response strings or patterns, then this could potentially be blocked. However, even most intrusion detection and prevention devices today would rely on a standard signature-based method to detect this, and most likely wouldn't.

Another key requirement for application detection and analysis is the ability to identify anonymous proxies and other “control bypass” techniques. Aside from simply blocking known domains and IP addresses hosting proxies on the Internet, a firewall should be able to identify common remote access tools like Remote Desktop, proxies like PHPProxy and others, and hosted remote access applications such as GoToMyPC. All of these tools have signatures that could be identified with a deeper level of application inspection. Additionally, many legitimate business applications have specific functions that should be disallowed for security reasons. For example, an organization may be using WebEx for online meetings and conference collaboration with employees, partners, and vendors. While the standard WebEx functionality is perfectly legitimate and allowed per policies, WebEx Desktop Sharing could inadvertently allow conference participants to view sensitive data on a company system, or worse, this feature could be surreptitiously enabled by malware or other threats. Network security controls should ideally be able to differentiate between these WebEx functions and block as needed.

Finally, network-based malware detection within approved application functions is a feature that newer firewalls should have, as well. For example, Microsoft SharePoint allows for simple sharing of documents using the HTTPS and CIFS protocols, and these may be infected with malware. Ideally, transfer of these documents should be detected and blocked by network firewalls instead of having to employ additional network-based security technologies, which add additional overhead and complexity to the environment.

Key Requirement: Visibility into Encrypted and Unknown Traffic

One of the most critical needs in network security at the moment is inspection of encrypted traffic without impacting the overall network performance. Although many organizations are using SSL-enabled traffic control systems currently, ranging from proxies and load balancers to specialized SSL Offload devices for firewalls and other systems, these SSL termination tools are usually not applying network security policies and deeply inspecting and analyzing the application traffic afterward. SSL traffic may comprise anywhere from 15% to over 50% of an organization’s traffic. For business use, SSL enables application traffic to communicate securely, and data can be safely sent into and out of the organization’s network with less concern over eavesdropping and data loss.

However, the use of SSL has traditionally left security teams with a glaring weakness - the inability to inspect and analyze that network traffic for threats. Savvy attackers know this, which explains the huge increase in botnet command and control channels using HTTPS. It is widely believed that data stolen in some of the largest breaches to date, including Heartland Payments and TJX, was exfiltrated over an encrypted channel to avoid detection as well. Firewalls should be able to natively decrypt and inspect SSL traffic regardless of the ports in use (although traditionally TCP port 443, this can be changed easily). Once decrypted, this traffic should be subject to policy analysis just like any other traffic traversing the environment, with a particular emphasis on data leakage and malware communication channels. With the growth in complex application traffic, firewalls have slowly shifted away from the classic “Default Deny” stance espoused by security professionals for several decades. This model takes a “whitelisting” approach to traffic control, classifying approved and allowed traffic and disallowing everything else. With this model, all unknown threats are hypothetically blocked from getting in or out. Unfortunately, this no longer works well using OSI layers 3 and 4, where IP addresses and ranges and TCP and UDP ports define what is acceptable and what isn’t. Firewalls will need to have a much more thorough inspection capability to analyze and “identify” traffic in order to place it on the whitelist, and nowhere is this more true than with custom-developed applications. Most of these, along with commercial applications and well-known services and protocols like DNS and FTP, have some identifiable signatures that can be used to classify and allow the legitimate traffic.

The growth in “mashup” and multi-function Web applications requires firewalls to evaluate multiple traffic types and use cases during the same sessions initiated by the same users. For example, a user signed into Gmail could initiate a Google Talk VoIP call, and also be using the Google Chat features within the context of the same online session in the browser. To a traditional firewall, this all looks like HTTP or HTTPS traffic to Google’s servers. However, these are very different types of communications that may need different security policy analysis, and a network firewall should be able to discern between the various application types and apply policies as needed.

Revolution, Not Evolution

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

Finally, firewalls should be able to apply deep application analysis techniques to all users, whether inside the primary business location or connecting via VPN or other remote access techniques. As more of the workforce is becoming increasingly mobile, workers are now connecting to internal systems from their homes, airports, coffee shops, and other far-flung locations. Corporate policies may permit the use of some applications while outside that are not allowed within the office environment, and network firewalls should be able to analyze remote users’ traffic and allow corporate application traffic while disallowing other traffic that should not be permitted internally.

Key Requirement: Operational Efficiency and Performance

Security and IT teams are increasingly taxed, and may be handling more duties and responsibilities due to downsizing or budget cuts. Adding additional security and network tools and controls to an already complex infrastructure may not be practical in terms of operational overhead. Often, new tools get installed with a base configuration established with the help of vendors, and operational teams quickly fall behind in updating and tuning these systems. Additionally, many of these tools are incredibly complex to configure and manage, and it is easy to make configuration mistakes that either impact network availability or permit too much traffic in the environment.

Unfortunately, as the threat landscape rapidly evolves, network and security teams need to be more vigilant than ever before, and need tools that can help them accurately identify and control ALL the traffic in their environments. Classifying applications and application traffic is not likely to be a top priority without significant help, which means that network filtering devices will need to simplify the process of creating policies and applying them flexibly.

At the same time, the addition of application inspection capabilities cannot impede the flow of traffic into and out of the network. As more organizations upgrade networks to 10GB speeds, inline filtering products will need to keep pace with the traffic while performing increasingly complex and sophisticated inspection tasks and applying policies on several interfaces at once. Without customized hardware and purpose-built software that can maximize the hardware’s efficiency, this is a difficult goal at best. The next generation of firewalls will need to withstand intense scrutiny for “speeds and feeds”, all while performing complex security filtering simultaneously.

Conclusion

The firewalls of yesterday are rapidly losing the ability to defend against modern threats. With the proliferation of Web-based attacks, sophisticated malware that uses high-volume protocols for command and control, and the increasing likelihood of sensitive data breaches, network filtering solutions will need to work harder than ever. Without the ability to scrutinize and identify numerous types of application traffic, it’s a lost cause. There is too much traffic, too many applications, and too little tolerance for negative performance impacts to keep adding in devices and new software “modules” that will help to analyze traffic. On top of this, security and network teams are barely keeping their heads above water as it is, so the addition of more tools only serves to make life more difficult in many cases, and operational tasks like policy creation, rule updates, and tuning filters and configurations are less likely to happen regularly. The next generation of network filtering tools will need to address all of these issues and more, paving the way for a more robust network security architecture that is more effective than current technologies while simpler to manage, as well.

About IANS

IANS is the leading provider of in-depth security insights delivered through its research, community, and consulting offerings. Fueled by interactions among IANS Faculty and end users, IANS provides actionable advice to information security, risk management, and compliance executives. IANS powers better, faster, technical and managerial decisions through experience-driven advice.

¹ <http://www.team-cymru.com/ReadingRoom/Whitepapers/2008/http-botnets.pdf>

² <http://www.securelist.com/en/analysis/204792131/TDSS#9>

³ http://www.sans.org/reading_room/whitepapers/malicious/analysis-simple-http-bot_33573

Contents

Introduction	9
Architecture and Security Model Considerations: Traffic is Best Classified in the Firewall	10
10 Things Your Next Firewall Must Do	11
Firewalls Should Safely Enable Applications – and Business	14
Using the RFP Process to Select a Next-Generation Firewall	15
Firewall Architecture and Control Model Considerations	15
Threat Prevention	17
Securing Remote Users	17
Management	18
Performance	18
Additional RFP Considerations	18
Evaluating Next-Generation Firewalls Through Formal Testing	19
Application Visibility and Control	19
Threat Prevention	20
Securing Remote Users	20
Management	21
Performance with Services Enabled	21
Additional Evaluation Considerations	21
Secure Application Enablement With Next-Generation Firewalls	21

Introduction

Much has been made about bringing application visibility and control into network security. The reason is obvious: applications can easily slip by traditional port-based firewalls. And the value is obvious: employees will use any application needed to get their jobs done—often indifferent to the risk that the use of these applications pose to the business. Nearly every network security vendor acknowledges that application control is an increasingly critical part of network security. While the next-generation firewall is well defined by Gartner as something new, enterprise-focused, and distinct, many network security vendors are claiming next-generation firewalls are a subset of other functions (e.g., UTM or IPS). Most traditional network security vendors are attempting to provide application visibility and control by using a limited number of application signatures supported in their IPS or other external database. But underneath, these capabilities are poorly integrated and their products are still based on legacy port-blocking technology, lacking next-generation firewall technology. Perhaps most importantly, many firewall vendors are missing the point – it’s not about blocking applications, it’s about safely enabling them. Unfortunately, the products currently offered by traditional network security vendors ignore much of what enterprises do with applications today – they use them to enable their business – and as such, need to make sure that those applications run securely. It is obvious that a next-generation firewall is a different and revolutionary class of product, but the interest from enterprise customers is so strong that vendors of traditional products are trying to subvert the interest of enterprise network security team by attempting to look like a next-generation firewall.

Definition: Next-Generation Firewall

Five Key Requirements

1. Identify applications regardless of port, protocol, evasive tactic or SSL
2. Identify users regardless of IP address
3. Protect in real-time against threats embedded across applications
4. Fine-grained visibility and policy control over application access/functionality
5. Multi-gigabit, in-line deployment with no performance degradation

For enterprises looking at next-generation firewalls, the most important consideration is: Will this new technology empower security teams to securely enable applications to the benefit of the organization? Consider the following:

- Will it increase visibility and understanding of application traffic?
- Will it expand traffic control options beyond blunt allow/deny?
- Will it help prevent threats?
- Will it eliminate the need to compromise between performance and security?
- Will it reduce costs for my organization?
- Will it make the job of risk management easier or simpler?

If the answers to the above questions are “yes,” then this transition is easy to justify.

There are substantial differences between next-generation firewalls and UTM-style devices – in terms of the organizations each solution is used in and in terms of architecture and security model. These differences have dramatic impacts on real-world functions/features, operations, and performance – as we’ve attempted to capture in the *10 Things Your Next Firewall Must Do* section that follows.

Architecture and Security Model Considerations: Traffic is Best Classified in the Firewall

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 application signatures to an IPS or IPS-like pattern matching engine which is then added to a port-based firewall.

Both 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 (default deny), or negative (default allow).

Firewalls use a positive security model. Another term is “default deny.” This means administrators write policies to allow traffic (e.g., allow WebEx)... and then everything else is denied or blocked. Negative policies (e.g., block Limewire) can be used in this model, but the most important fact is that the end of the policy in a positive security model says, “all else deny.” A key implication of this approach is all traffic must be classified in order to allow the appropriate traffic. This means that application visibility is complete and that policies can be used to enable applications. Another key result of this approach is that any unknown traffic is by default denied. In other words, the best next-generation firewall is a firewall.

Intrusion prevention systems (IPS) typically employ a negative security model, or default allow. This means that IPS identifies and blocks specific traffic (traditionally threats) and everything else is passed through. Traditional network security vendors are adding application signatures to an IPS-style engine and bolting it onto a traditional port-based firewall. The result is an “application prevention system.” The application control is in a negative security model. In other words, application control is not in a firewall. Implication: one only sees what is expressly looked for, and unknown traffic is, by default, allowed.

The remainder of this paper is broken down into three distinct sections. The first section introduces the 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 securely enabling applications at the firewall. The remaining sections will follow this same format. Sections two and three will 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 functional elements correspond to the efficacy of the security controls, and the ability for enterprises to manage risk associated with the applications traversing the network. From an operations perspective, the big question is, “where does application policy live, and how hard or complex is it to manage?” The performance difference is simple: can the firewall do what it’s supposed to do at the throughput it’s supposed to do it? While each organization will have varied requirements within each of these three areas, and differing levels or prioritization, the 10 things your next firewall must do are:

1. Identify and control applications on any port
2. Identify and control circumventors
3. Decrypt outbound SSL
4. Identify and control applications sharing the same connection
5. Provide application function control
6. Deal with unknown traffic by policy
7. Scan for viruses and malware in allowed collaborative applications
8. Enable the same application visibility and control for remote users
9. Make network security simpler, not more complex with the addition of application control
10. Deliver the same throughput and performance with application control active

1.

Your next firewall must identify and control applications on any port, not just standard ports (including applications using HTTP or other protocols).

Business case: Application developers no longer adhere to standard port/protocol/application mapping. 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., MS RDP, SSH). In order to enforce application-specific policies where ports are increasingly irrelevant, your next firewall must assume that any application can run on any port. This is one of the fundamental changes in technology that made the next-generation firewall an absolute necessity. It was this change to applications that made the positive control of traditional port-based firewalls obsolete. It 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 have to run all signatures on tens of thousands of ports.

Requirements: This one is simple, you should assume that any application can run on any port and your next firewall must classify traffic, by application, on all ports – all the time. Classification on all ports, all the time 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.

2.

Your next firewall must identify and control circumventors: proxies, remote access, and non-VPN-related encrypted tunnel applications.

Business case: Most organizations have security policies – and controls designed to enforce those policies. Proxies, remote access, and encrypted tunnel applications are specifically used to circumvent security controls like firewalls, URL filtering, IPS and secure web gateways. Without the ability to control these circumventors, organizations cannot enforce their security policies, and expose themselves to the very risks they thought their controls mitigated. To be clear, not all of these types of applications are the same – remote access applications have legitimate uses, as do some encrypted tunnel applications.

However, external anonymous proxies that communicate over SSL on random ports, or applications like Ultrasurf and Tor have only one real purpose - to circumvent security controls.

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 PHPProxy or CGIProxy. Remote access applications like MS RDP or GoToMyPC can have legitimate use – but due to the associated risk, should be managed. Most other circumventors, (e.g., Ultrasurf, Tor, Hamachi) don't have business uses. Regardless of the policy stance, your next firewall needs to have specific techniques to deal with 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 not only that your next firewall can identify these circumvention applications, but it is also important to know how often that firewall's application intelligence is updated and maintained.

3.

Your next firewall must decrypt outbound SSL.

Business case: Today, applications using SSL in some way, shape or form represent 25% of the applications and 23% of the overall bandwidth on corporate networks Application Usage and Risk Report (May 2011) In some industries (e.g., financial services), it's more than 50%. 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, network security teams have a large and growing blind spot without decrypting, classifying, controlling, and scanning 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 health care organizations) while other types (e.g., SSL on non-standard ports, HTTPS from unclassified websites in country X) can be decrypted via policy.

Requirements: The ability to decrypt outbound 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, policy control over decryption, and the necessary hardware and software elements to perform SSL decryption across tens of thousands of simultaneous SSL connections with good performance and high throughput. Additional requirements to consider should be the ability to decrypt and inspect inbound SSL traffic as well as the ability to identify and control the use of SSH. Specifically, SSH control should include the ability to determine (and control), if it is being used for port forwarding (local, remote, x11) or native use (SCP, SFTP and shell access).

4.

Your next firewall must identify and control applications sharing the same connection.

Business case: Applications share sessions. To ensure users are continuously using an application "platform," whether it's Google, Facebook, Microsoft, Salesforce.com, LinkedIn, or Yahoo, application developers integrate many different applications – which often have very different risk profiles and business value. For example, when using Gmail you can branch off to use Google Talk, which is a fundamentally different application with a different business and security risk profile, and your next firewall must be able to recognize that switch and enable the appropriate policy response for each.

Requirements: Simple classification of the application platform or website doesn't work. In other words, "fast path" is not an option – "once and done" classification ignores the fact that these commonly used applications share sessions. Traffic must be continuously evaluated to understand the application, its changes (see #5), when the user changes to a completely different application using the same session, and enforce the appropriate policy controls. Looking briefly at the technical requirements using our Gmail/Google Talk example: Gmail is by default HTTPS (see #3) so the first step is to decrypt – but it has to be continuous, as does the application classification, because at any time, the user can start a switch functions from email to chat... which may have a completely different policy associated with it.

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, but also controls capable of enforcing them.

5.

Your next firewall must provide application function control (e.g., SharePoint Admin vs. SharePoint Docs).

Business case: Many applications have significantly different functions, presenting different risk profiles and value to both the user, and the organization. Good examples of this include WebEx vs. WebEx Desktop Sharing, Yahoo Instant Messaging vs. the file transfer feature, and regular Gmail vs. sending attachments. In regulated environments or in organizations heavily dependent on intellectual property this is a significant issue.

Requirements: Continuous classification and fine-grained understanding of each application. Your next firewall has to continually evaluate the traffic and watch for changes – if a different function or feature is introduced in the session, the firewall must note it and perform a policy check. Understanding the different functions of each application and the different associated risks is equally important. Unfortunately, many firewalls classify a traffic flow once, and then “fast path” it (read: never look at that flow again) for better performance. This method pre-dates modern applications and prevents those firewalls from meeting this requirement.

6.

Your next firewall must systematically deal with unknown traffic by policy, not by just letting it through.

Business case: There will always be unknown traffic and it will always represent significant risks to any organization. There are several important elements to consider with unknown traffic – minimizing it, easily characterizing custom applications so they are “known” in network security policy, and having predictable visibility and policy control over traffic that remains unknown.

Requirements: First, by default, your next firewall must attempt to classify all traffic – this is one area where the earlier architecture and security discussion becomes very important. Positive (default deny) models classify everything, negative (default allow) models classify only what they’re told to classify. Second, for custom developed applications, there should be a way to develop a custom identifier – so that traffic is counted among the “known.” Third, the security model plays into these requirements again – a positive (default deny) model can deny all unknown traffic – so what you don’t know can’t hurt you. A negative (default allow) model allows all unknown traffic – so what you don’t know will hurt you.

For example, many botnets will use port 53 (DNS) for communication back to their control servers. If your next firewall lacks the ability to see and control unknown traffic, bots will be able to drive right through, unimpeded.

7.

Your next firewall must scan for threats in allowed collaboration applications (e.g., SharePoint, Box.net, Microsoft Office Live).

Business case: Enterprises continue to adopt collaborative applications hosted outside their physical locations. Whether it’s hosted SharePoint, Box.net, Google Docs, or Microsoft Office Live, or even an extranet application hosted by a partner, many organizations have a requirement to use an application that shares files – in other words, is a high-risk threat vector. Many infected documents are stored in collaboration applications, along with some documents that contain sensitive information (e.g., customers’ personal information). 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 along with the (potential) associated threats.

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 and HTTPS - see requirement #3), and require a more sophisticated policy than “block application.” The first step is to identify the application (regardless of port or encryption), allow it, and then scan it for any of the appropriate threats – exploits, viruses/malware, or spyware... or even confidential, regulated, or sensitive information.

8.

Your next firewall must enable the same application visibility and control for remote users as for on-premise users.

Business case: Users are increasingly outside the four walls of the enterprise. Once the domain of road warriors, now a significant portion of the enterprise user population is capable of working remotely. Whether working from a coffee shop, home, or a customer site, users expect to connect to their applications via WiFi, wireless broadband, or any means necessary. Regardless of where the user is, or even where the application they’re employing might be, the same standard of control should apply. If your next firewall enables application visibility and control over traffic inside

the four walls of the enterprise, but not outside, 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 – inside or outside. This is not to say that enterprises will have the exact same policy for both – some organizations might want employees to use Skype when on the road, but not inside headquarters, where others might have a policy that says if outside the office, 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 or undue operational hassle for the administrator, or significant cost for the organization.

9.

Your next firewall must make network security simpler, not more complex with the addition of application control.

Business case: Many enterprises struggle with incorporating more information feeds and more policies, and more management into already overloaded security processes and people. In other words, if teams cannot manage what they've already got, adding more management, policies, and information doesn't help. Furthermore, the more distributed the policy is (e.g., port-based firewall allows port 80 traffic, IPS looks for/blocks threats and applications, secure web gateway enforces URL filtering) – the harder it is to manage that policy. Where do admins go to enable WebEx? How do they resolve policy conflicts across these different devices? Given that typical port-based firewall installations have rulebases that include thousands of rules, adding thousands of application signatures across tens of thousands of ports (see #3 above) is going to increase complexity by several orders of magnitude.

Requirements: Firewall policy should be based on user and application. Subsequent content analysis can be performed on allowed traffic, but fundamental access control should be based on relevant elements (i.e., application and user or group). This can have a significant simplifying effect. Firewall policy based on port and IP address, followed by subsequent analysis to understand the application makes things more complicated than they are today

10.

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

Business case: Many enterprises struggle with the forced compromise between performance and security. All too often, turning up security features in the network security realm means turning down 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) 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.

Firewalls Should Safely Enable Applications – and Business

Users continue to adopt new applications and technologies – and the threats carried by them. In some organizations, obstructing the adoption of new technologies can be a career-limiting move. Even when it isn't, applications are how employees get their jobs done, or maintain productivity in the face of competing personal and professional priorities. Because of this, safe enablement is increasingly the correct policy stance. But to safely enable these applications and technologies, and the business that rides atop them, network security teams need to put in place the appropriate policies governing use, but also controls capable of enforcing them.

The 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 step is to translate those requirements into actionable steps; selecting a vendor through an RFP process, evaluating, that will culminate 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 Gartner, “the changing threat conditions and changing business and IT processes will drive network security managers to look for next-generation firewall capabilities at their next firewall/IPS refresh cycle.” 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 things 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, not 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 attacks 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. It is important to determine if the application identification mechanisms port-agnostic or are they dependent upon specific application ports. Are the signatures dependent on a specific port, or range of ports – or are they applied automatically to all ports, all the time?
- When traffic first hits the device, is it first classified based on port (this is port 80, therefore it is HTTP) or application (this is Gmail)?
- Describe in detail how the firewall can accurately identify applications. Are signatures the only mechanism, or are other elements such as decoders, heuristics, and decryption used as a means of ensuring that all applications are identified?
- What 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? 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 to manage?
- How often is the application database updated and is it a dynamic update or a system reboot upgrade?

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.

Controlling Evasive Applications, SSL and SSH

A wide range of applications can be used to circumvent security controls. Some, such as external proxies, and encrypted tunnels are designed with circumvention as a goal. Others, such as remote-desktop access 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. In either case, it is important to determine how the potential vendors will address this category of applications. Our own data from the Application Usage and Risk Report (May 2011) shows that 25% of the applications found are capable of using SSL in some way, shape or form. Consider the following questions and statements when issuing an RFP for next-generation firewalls:

- Describe the process by which applications and protocols are identified on all, including non-standard ports.
- 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.
- 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 at 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.

Policy-based Application Enablement

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. Many "platforms" (Google, Facebook, Microsoft) enable different applications once the user initially logs in, therefore, it is imperative to determine how the vendor offering monitors the state of the application,

looking for changes on the application, and more importantly, is the change in state classified correctly and how can it be used within policy. Consider the following questions and statements when issuing an RFP for next-generation firewalls:

- Describe how the application database hierarchy (flat, multi-level, other) exposes functions within the parent application for more granular enablement policies.
- Is stateful inspection traffic classification performed separately, prior to application identification and if so, describe how, once an application is identified, the changes in application state are monitored, tracked and made use of within the policy.
- 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, port;
 - allow and scan for viruses, application exploits, spyware, drive-by downloads;
 - allow and shape;
 - deny.
- Can port-based controls be implemented for all applications in the application database so that an administrator can for example, force Oracle database developers over a specific port or range of ports?
- List all the enterprise repositories supported for user-based controls.
 - Is an API available for custom or non-standard repository integration?
- Describe how policy-based controls are implemented by users and groups for terminal services environments.

Systematically Managing Unknown Applications

Every network will have 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 is a specific description of how the vendor enables users to systematically manage the unknown traffic, which by its very nature, represents a higher business and security risk. Consider the following questions and statements when issuing an RFP for next-generation firewalls:

- Provide specifics on how unknown traffic is identified, categorized and managed.
 - 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 custom application signatures 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 ex-ploits 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, anti-malware, URL filtering, data loss prevention, 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 carried 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 the approach to controlling unknown vulnerabilities and unknown malware.
- Describe the threat prevention research and development process.
- List all threat discoveries over the last 12 months.

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 these instances where they are beyond the network. 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 product transparent to the client?
- Describe how performance is ensured for geographically distributed users.

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.

- 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.)

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?
- What management options are supported: CLI? Browser? Fat-client? Centralized Server?
- What visibility tools, outside of the log viewer and reporting, are available to enable a clear picture of the application traffic 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 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 for specific applications such as BitTorrent, SharePoint and MS-Exchange.
 - Is full log analysis available on-box, or is it an extra cost/added license/separate device?
- Are 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/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.

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, if it is an appliance, to confirm appropriate processing power is applied to the task at hand. Using repurposed servers will result in poor performance when multiple services are enabled.
- What has the actual performance result been in a test environment that is representative of the target network environment?
- What is the rated throughput?
- What is the throughput based on a real-world mix of traffic with application control enabled?
- What is the throughput based on a real-world mix of traffic with application control enabled or disabled?
- What is the throughput based on a real-world mix of traffic with all application control, user and threat prevention options enabled?

Additional RFP Considerations

Every organization will have varied requirements over and above the items listed within this document. Examples may include company viability, customer references, ease of deployment, as well as networking and routing support. The recommended best practices for an RFP is 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 organizations’ business. This section provides some recommendations on how to physically evaluate a next-generation firewall. The evaluation will give organizations 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, determine that the first task that the device under test (DUT) executes is traffic classification based on the application identity; not the network port. Second, determine that the DUT classifies applications on any port, even those that can hop ports, use 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.

Application Identification

- Confirm that the firewall can identify various applications. The ideal way to execute this verification 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 high level and ground level visibility and analysis tools.
- 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 Skype, AIM, or one of many P2P applications.

Application Identity as Basis of Firewall Security Policy

- Confirm that when creating a firewall policy, the application, not the port, is used as the primary policy element. Termed another way, does the application control policy require a port-focused rule first and 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.

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 applications.
- Validate that if the decrypted application is “allowed”, it will be re-encrypted and sent on its way.
- Confirm the ability to perform SSL decryption and inspection both inbound and outbound.
- 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).

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.

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.

- Using an application such as Gmail or SharePoint, first confirm that the DUT identifies the initial application (such as Gmail or SharePoint).
- Without logging out of the application, switch to a separate function (Google Docs, Google Chat, 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 application usage, yet exerting some level of control to address potential business and security risks. File transfer is a common example, but other examples may include administrative, VoIP, email, blog posting, or chat functions within the parent application.

- Confirm that the DUT provides visibility into the application hierarchy (base 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 will have a small amount of unknown traffic, and you need to determine how well the DUT can manage it (application or threat).

- Validate that visibility into unknown traffic is available (users, IP addresses, etc.) and can be used in policy control (allow, block, inspect, etc.)
- Verify, using an internal application and known IP addresses, that the traffic is identified as unknown.
- 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 custom identification mechanism and if so, how quickly?

Threat Prevention

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

- Verify that threat prevention techniques (IPS, malware, content filtering) are consistently applied even on 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.
- 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 with the same policy used internally and second, determine the management effort and deployment complexity.

- Verify 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.
- Close the test out by monitoring remote users via the log viewer.

Secure Application Enablement With Next-Generation Firewalls

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.
- Validate that application policy controls, firewall policy controls, and threat prevention features can all be enabled from the same policy editor.

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 appropriate processing power is applied to the task at hand.
- Evaluate the actual performance in a test environment using traffic patterns that are representative of the target network environment.

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 ease of deployment (tap mode, transparent mode, other?), networking (layer 2, layer 3, mixed mode) and routing (RIP, OSPF, BGP) 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.

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, often times, melding personal and work-related usage. Blindly blocking these applications is equivalent to blocking the business.

The 10 Things Your Next Firewall Must Do section validates the fact that the best location to execute secure application enablement is at the firewall 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 network security company. Its next-generation firewalls enable unprecedented visibility and granular policy control of applications and content—by user, not just IP address—at up to 20Gbps with no performance degradation. Based on patent-pending App-ID™ technology, Palo Alto Networks firewalls accurately identify and control applications—regardless of port, protocol, evasive tactic or SSL encryption—and scan content to stop threats and prevent data leakage. Enterprises can for the first time embrace Web 2.0 and maintain complete visibility and control, while significantly reducing total cost of ownership through device consolidation. Most recently, Palo Alto Networks has enabled enterprises to extend this same network security to remote users with the release of GlobalProtect™.



Reinventing Network Security | www.paloaltonetworks.com