It's Time to Fix The Firewall

Re-Establishing the Firewall as The Cornerstone of Enterprise Network Security

February, 2009

Palo Alto Networks
232 East Java Dr.
Sunnyvale, CA 94089
Sales 866.207.0077
www.paloaltonetworks.com

TABLE OF CONTENTS

EXECUTIVE SUMMARY

For the last 15 years, port-blocking firewalls have been the cornerstone of enterprise network security. But much like a stone, they've stood still in the face of rapidly evolving applications and threats. It's no secret that modern applications and threats easily circumvent the traditional network firewall – so much so that enterprises have deployed an entire crop of "firewall helpers" to help try to manage applications and threats. But that hasn't worked – applications and threats still easily make their way around these "helpers," frustrating enterprise IT groups who have taken on additional complexity and costs without fixing the problem.

Palo Alto Networks was founded to fix the firewall – starting over with a clean sheet of paper, and focusing on applications, users, and content as the key elements to deliver visibility and control. Built on a high-performance architecture, Palo Alto Networks' next-generation firewalls enable enterprises to safely enable a new breed of applications, without taking on the unnecessary risks that accompany them in traditional organizations. By "fixing the firewall," enterprises can increase visibility and control, and reduce risk. Enterprises also experience an additional benefit of fixing the firewall – streamlined and simplified network security infrastructure resulting in a substantial reduction in complexity and cost.

NEW APPLICATIONS AND THREATS ARE EXTREMELY EVASIVE

Over the past several years there have been a number of significant changes to both the application and threat landscapes.

PERVASIVE PERSONAL APPLICATIONS

To begin with, user-centric applications have become pervasive. Internet-oriented and originally intended primarily for personal communications, this class of applications includes instant messaging, peer-to-peer file sharing, web mail, and the plethora of social networking sites that have emerged in recent years. The issue is that their presence on enterprise networks is practically guaranteed, even if an organization's policies indicate otherwise. Not only are these applications extremely popular, but they've also been designed to evade traditional countermeasures, such as firewalls, by dynamically adjusting how they communicate. Common tactics include:

- Port hopping, where ports/protocols are randomly shifted over the course of a session;
- Use of non-standard ports, such as running Yahoo! Messenger over TCP port 80 instead of TCP port 5050;
- Tunneling within commonly used services, such as when P2P file sharing or an IM client like Meebo is running over HTTP; and,
- Hiding within SSL encryption.

BUSINESS APPLICATIONS MIMIC PERSONAL APPLICATIONS

Two closely related developments complicate matters further. First is the fact that many of these next-generation applications have proven to be extremely useful for more than just personal communications. These days enterprises worldwide are routinely employing them for legitimate business purposes as well – helping to accelerate key processes, improve customer service, and enhance collaboration, communications, and employee productivity in general.

The second development is that new business applications are often being designed to take advantage of the same types of evasion techniques. The intentions are typically positive in this case: to facilitate operation in the broadest set of scenarios and with the least amount of disruption for customers, partners, and the organization's own security and operations departments. However, the unintended side effect of IT further losing control over network communications is clearly negative.

Another relevant trend is the webification of enterprise applications. To improve accessibility and reduce administrative effort and costs, standard client-server applications are steadily being re-designed to take advantage of Web technologies. Alternately, a wide range of conventional applications are being displaced all together in favor of hosted, Web-based services such as Salesforce.com, WebEx, and Google Apps. The result is that HTTP and HTTPS now account for approximately two thirds of all enterprise traffic. By itself this is not a problem, per se, but it does exacerbate an inherent weakness of traditional security infrastructure. Specifically, for most older products, the wide variety of higher-order applications riding on top of this universal protocol, whether or not they serve a legitimate business purpose, are practically indistinguishable.

THREAT EVOLUTION: PROFIT MOTIVE + APPLICATION PIGGYBACKING

Turning to the threat landscape, there have been significant changes there too. In particular, a shift in motivation – from building reputations to actually making money – means that hackers are now focused on evasion as well. In this regard, one of the general approaches they are pursuing is to build threats that operate at the application layer. This allows their creations to pass right through the majority of enterprise defenses, which have historically been designed to provide network-layer protection.

Today's hackers are also paying considerable attention to the growing population of user-centric applications. This is supported by the SANS Institute routinely including instant messaging and peer-to-peer programs on its list of the SANS Top-20 Security Risks. Not only are such applications interesting targets due to their high degree of popularity, but also because their evasion capabilities can be leveraged to provide threats with "free passage" into enterprise networks.

IT IS NO LONGER IN CONTROL

The impact of all the ongoing changes to the application and threat landscapes is that IT has lost control. In reality, however, the inability of their security infrastructure to effectively distinguish good/desirable applications from those that are bad/unwanted leaves most shops with no reasonable option. One possibility is to continue with business as usual, an approach that ensures the availability of desirable applications by allowing sessions associated with all types of next-generation applications to proceed unchecked. Alternately, organizations can attempt to crank down on bad and unwanted sessions as best they can with the tools they have on hand. Not only is this second approach highly unlikely to be successful, but it also suffers from the propensity to throw the good out with the bad.

To rectify this situation, enterprises need security technology with sufficient visibility and intelligence to discern:

- which network traffic corresponds to applications that serve a legitimate business purpose;
- which network traffic corresponds to applications that can serve a legitimate business purpose but, in a given instance, are being used for unsanctioned activities; and,
- which communications traffic, even though it corresponds to legitimate business activities, should be blocked because it includes malware or other types of threats.

LEGACY PORT-BLOCKING FIREWALLS ARE INEFFECTIVE

Providing highly granular access control is functionality that would normally be expected of the enterprise firewall. Based on its ability to control the flow of communications traffic, this long-standing pillar of enterprise security has historically been used in strategic locations to establish the boundary between domains characterized by different levels of trust – such as at the Internet gateway, on connections to partner networks, and, more recently, at the logical front door to the data center.

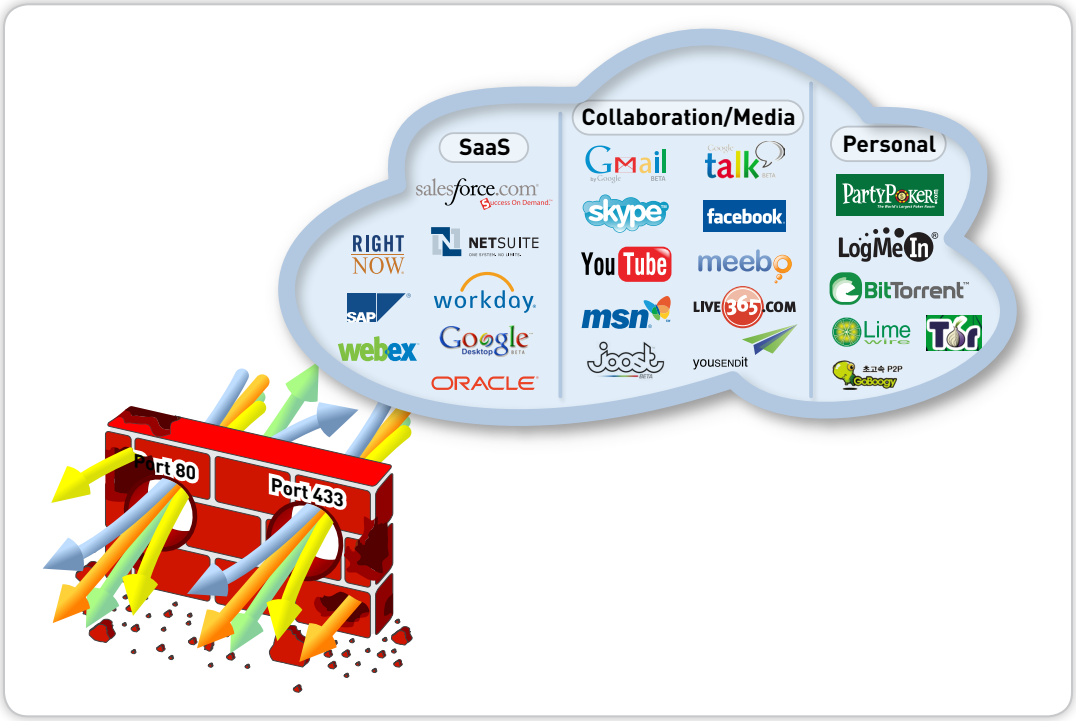

Figure 1: Port Blocking Firewalls Can't See or Control Applications

The problem, though, is that most firewalls are far-sighted. They can see the general shape of things, but not the finer details of what is actually happening. This is because traditional firewalls operate by inferring the application-layer service that a given stream of traffic is associated with based on port numbers. They rely on a convention – not a requirement – that a given port corresponds to a given service (e.g., TCP port 80 corresponds to HTTP). As such, they are also incapable of distinguishing between different applications that use the same port/service.

Consequently, traditional "port-blocking" firewalls are basically blind to the new generation of applications. They can't account for common evasion techniques such as port hopping, protocol tunneling, and use of non-standard ports. And, therefore, they can't even begin to address the visibility and intelligence requirements identified above. For enterprises that continue to rely on these products – as well as other countermeasures that suffer from the same limitations – the result is that their networks are becoming like the wild, wild west: users have free rein to do whatever they want with whichever applications they choose.

FIREWALL REMEDIES HAVE FAILED

It doesn't really help matters that the two most common steps taken to address the inadequacies of traditional firewalls have, for all intents and purposes, been completely unsuccessful.

BOLTING-ON DEEP PACKET INSPECTION IS FUNDAMENTALLY FLAWED

Many purveyors of traditional firewalls have attempted to correct the myopic nature of their products by incorporating deep packet inspection (DPI) capabilities. On the surface, adding a measure of application-layer visibility and control in this manner may appear to be a reasonable approach. However, the boost in security effectiveness that can be achieved in most cases is only incremental because (a) the additional capability is being "bolted on", and (b) the foundation it is being bolted to is weak to begin with. In other words, the new functionality is integrated rather than embedded, and the port-blocking firewall, with its complete lack of application awareness, is still used for initial classification of all traffic. The problems and limitations this leads to include the following:

- Not everything that should be inspected necessarily gets inspected. Because the firewall is unable to accurately classify application traffic, deciding which sessions to pass along to the DPI engine becomes a hit or miss proposition.
- Policy management gets convoluted. Rules on how to handle individual applications essentially get "nested" within the DPI portion of the product – which itself is engaged as part of a higher/outer level access control policy.
- Inadequate performance forces compromises to be made. Inefficient use of system resources and CPU and memory intensive application-layer functionality put considerable strain on the underlying platform. To account for this situation, administrators can only implement advanced filtering capabilities selectively.

DEPLOYING FIREWALL "HELPERS" DOESN'T SOLVE THE PROBLEM, AND LEADS TO COMPLEX AND COSTLY APPLIANCE SPRAWL

Left with no choice, enterprises have also tried to compensate for their firewall's deficiencies by implementing a range of supplementary security solutions, often in the form of standalone appliances. Intrusion prevention systems, antivirus gateways, Web filtering products, and application-specific solutions – such as a dedicated platform for instant messaging security – are just a handful of the more popular choices. Unfortunately, the outcome is disappointingly similar to that of the DPI approach, with one additional and often painful twist.

Not everything that should get inspected does because these firewall helpers either can't see all of the traffic, rely on the same port- and protocol-based classification scheme that has failed the legacy firewall, or only provide coverage for a limited set of applications. Policy management is an even greater problem given that access control rules and inspection requirements are spread among several consoles. And performance is still an issue as well, at least in terms of having a relatively high aggregate latency.

Then comes the kicker: device sprawl. As one "solution" after another is added to the network, the device count, degree of complexity, and total cost of ownership all continue to rise. Capital costs for the products themselves and all of the supporting infrastructure that is required are joined by a substantial collection of recurring operational expenditures, including support/maintenance contracts, content subscriptions, and facilities costs (i.e., power, cooling, and floor space) – not to mention an array of "soft" costs such as those pertaining to IT productivity, training, and vendor management. The result is an unwieldy, ineffective, and costly endeavor that is simply not sustainable.

IT'S TIME TO FIX THE FIREWALL

To be clear, because they are deployed in-line at critical network junctions, firewalls essentially see all traffic and, therefore, are the ideal resource for enforcing control. The challenge, as discussed, is that legacy firewalls are basically blind to the latest generation of applications and threats. This is only one part of the problem though. The other part is that attempts to remedy this situation have only focused on compensating for this deficiency. The far-from-stellar track record of these approaches raises a question however. Why not fix the problem at its core instead?

Indeed, why not avoid the need for "helpers" of any type by delivering a solution that natively addresses the essential functional requirements for a truly effective, modern firewall:

- The ability to identify applications regardless of port, protocol, evasive tactics or SSL encryption;
- The ability to provide granular visibility of and policy control over applications, including individual functions;
- The ability to accurately identify users and subsequently use identity information as an attribute for policy control;
- The ability to provide real-time protection against a wide array of threats, including those operating at the application layer; and,
- The ability to support multi-gigabit, in-line deployment with negligible performance degradation.

INTRODUCING PALO ALTO NETWORKS AND THE NEXT-GENERATION FIREWALL

Having recognized the challenges posed by the latest generation of applications and threats, Nir Zuk, security visionary and the co-inventor of Stateful Inspection, founded Palo Alto Networks in 2005. Backed by top-tier investors and a management team with extensive experience in the network security industry, the Palo Alto Networks' engineers set out to restore the effectiveness of the enterprise firewall by "fixing the problem at its core." Starting with a blank slate, the team took an application-centric approach to traffic classification in order to enable full visibility and control of all types of applications running on enterprise networks, new-age and legacy ones alike. The fruit of this effort is the Palo Alto Networks family of next-generation firewalls – the only firewall solution available in the market that fully delivers on the essential functional requirements identified in the previous section.

The key to this distinction and the next-generation firewall's market-leading capabilities is the combination of three innovative identification technologies, a high-performance design, and additional foundational features that yield a robust enterprise-class solution.

UNIQUE IDENTIFICATION TECHNOLOGIES RESTORE VISIBILITY AND CONTROL

The starting point for the enhanced visibility and control achieved by the Palo Alto Networks family of next-generation firewalls is three unique technologies: App-ID, User-ID, and Content-ID. These are the underlying components that enable enterprises to focus on business relevant elements such as applications, users, and content for policy controls, instead of having to rely on nebulous and often misleading attributes such as ports and protocols.

APP-ID: POSITIVELY IDENTIFY APPLICATIONS REGARDLESS OF PORT/PROTOCOL OR SSL ENCRYPTION

App-ID is the patent-pending traffic classification technology at the heart of the next-generation firewall. Using four distinct techniques it is able to determine the exact identity of more than 800 applications flowing across the network, irrespective of port, protocol, SSL encryption, or evasive tactics.

Application Protocol Detection and Decryption. This initial step determines the application protocol (e.g., HTTP) and, if SSL is in use, decrypts the traffic so that it can be analyzed further. Re-encryption is performed, as needed, after all of the identification technologies have had an opportunity to operate.

Application Protocol Decoding. This technique determines whether the initially detected application protocol is the "real one", or if it is being used as a tunnel to hide the actual application (e.g., Yahoo! Instant Messenger might be wrapped in HTTP).

Application Signatures. In this step of the process, context-based signatures look for unique properties and transaction characteristics to correctly identify the application regardless of the port and protocol being used. This includes the ability to detect specific functions within applications – e.g., file transfers within IM sessions, or desktop sharing within conferencing applications.

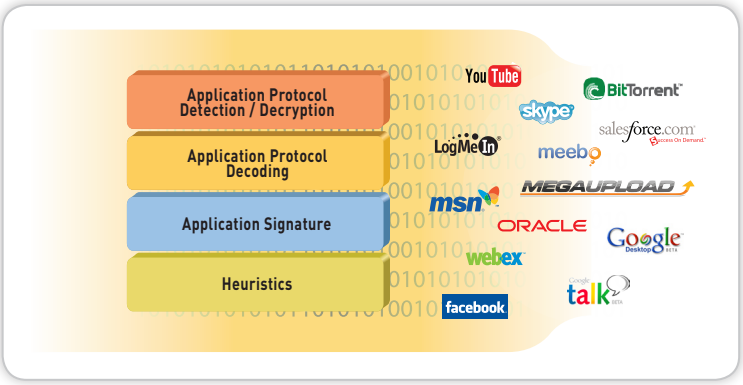

Figure 2: App-ID Identifies Applications Regardless of Port, Protocol, Evasive Tactic, or SSL Encryption

Heuristics. For traffic that eludes identification by signature analysis, additional heuristic, or behavioral, processes are applied. This enables identification of any troublesome applications, such as peer-to-peer or VoIP tools that use proprietary encryption.

Recognizing that identification is only part of the problem, Palo Alto Networks complements App-ID with an application browser. This powerful research tool provides administrators with a wealth of intelligence on over 800 applications so that they can make informed decisions on how to control them. Applications can be viewed by category, subcategory, underlying technology and a wide range of characteristics, including: file transfer capabilities, known vulnerabilities, ability to evade detection, and propensity to consume bandwidth, transmit malware, or otherwise be misused.

With App-ID, IT departments gain the visibility and intelligence needed to create and enforce policies that effectively control the applications traversing their networks.

USER-ID: ENABLE VISIBILITY AND CONTROL BY USER OR GROUP, NOT JUST IP ADDRESS

A standard feature on every Palo Alto Networks firewall platform, User-ID technology links IP addresses to specific user identities, enabling visibility and control of network activity on per-user basis. Tightly integrated with Microsoft Active Directory (AD), the Palo Alto Networks User Identification Agent supports this objective in two ways. First, it regularly verifies and maintains the user-to-IP address relationship using a combination of login monitoring, end-station polling, and captive portal techniques. Next, it communicates with the AD domain controller to harvest relevant user information, such as role and group assignments. These details are then available to:

- Gain visibility into who specifically is responsible for all application, content, and threat traffic on the network;
- Enable the use of user identity as a variable within access control policies; and,
- Facilitate troubleshooting/ incident response and be used in reports.

With User-ID, IT departments get another powerful mechanism to help control the use of applications in an intelligent manner. For example, a social networking application that would otherwise be blocked because of its risky nature can now be enabled for individuals or groups that have a legitimate need to use it, such as the human resources department.

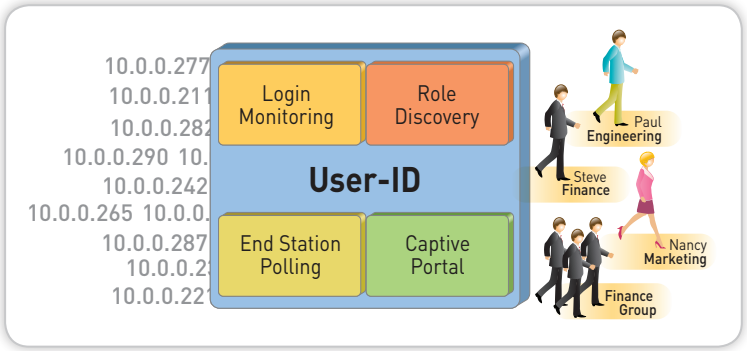

Figure 3: User-ID Integrates Enterprise Directories for User-based Policies and Reporting

CONTENT-ID: HIGH-PERFORMANCE CONTENT SCANNING PREVENTS THREATS, INAPPROPRIATE WEB CONTENT, AND SENSITIVE DATA LEAKS

Like its counterpart technologies, Content-ID infuses the Palo Alto Networks next-generation firewall with capabilities previously unheard of in an enterprise firewall. In this case, it's real-time prevention of threats within permitted traffic, granular control of web surfing activities, and file and data filtering.

Threat Prevention. This component of Content-ID leverages several innovative features to prevent spyware, viruses, and application vulnerabilities from penetrating the network, regardless of the type of application traffic – legacy or new-age – with which they hitch a ride.

- Application decoder. Content-ID leverages this App-ID component, using it to pre-process data streams that it then inspects for specific threat identifiers.
- Stream-based virus and spyware scanning. Scanning traffic as soon as the first packets of a file are received – as opposed to waiting until the entire file is loaded into memory – maximizes throughput while minimizing latency.

- Uniform threat signature format. Performance is further enhanced by avoiding the need to use separate scanning engines for each type of threat. Viruses, spyware, and vulnerability exploits can all be detected in a single pass.
- Vulnerability attack protection (IPS). Robust routines for traffic normalization and de-fragmentation are joined by protocol-anomaly, behavior-anomaly, and heuristic detection mechanisms to provide comprehensive protection from the widest range of both known and unknown threats.

URL Filtering. A fully integrated, on-box URL database is available so that administrators can monitor and control the web surfing activities of employees as well as guest users. Employed in conjunction with User-ID, web usage policies can even be set on a per-user basis, further safeguarding the enterprise from a broad spectrum of legal, regulatory, and productivity related risks.

File and Data Filtering. Taking advantage of the in-depth application inspection performed by App-ID, this set of features enables enforcement of policies that reduce the risk associated with unauthorized file and data transfer. Specific capabilities include the ability to block files by their actual type (i.e., not based on just their extension), and the ability to control the transfer of sensitive data patterns such as credit card and social security numbers. This complements the granularity of App-ID, which for many applications offers the ability to control the file transfer functionality within an individual application (e.g., an IM client).

The bottom line is that with Content-ID, IT departments gain the ability to stop known and unknown threats, reduce inappropriate use of the Internet, and help prevent data leakage – all without having to invest in a pile of additional products and risk falling victim to appliance sprawl.

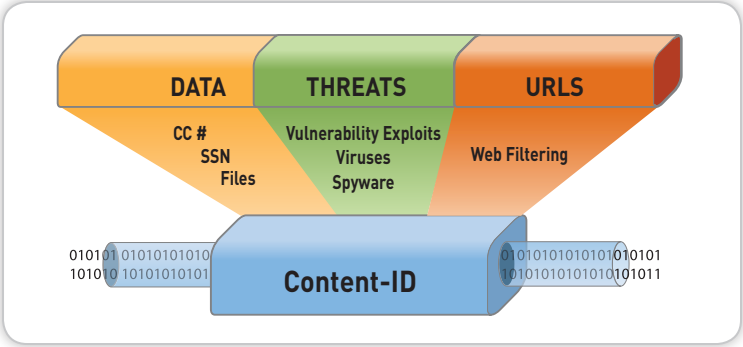

Figure 4: Content-ID Unifies Content Scanning for Threats, Confidential Data, and URL Filtering

HIGH-PERFORMANCE SP3 ARCHITECTURE DELIVERS "SECURITY WITHOUT COMPROMISES"

Having a comprehensive suite of application awareness and content inspection capabilities makes little difference if administrators are unable to fully engage them due to performance constraints. And to be clear, the issue is not just that these capabilities are inherently resource intensive. There's also the tremendous traffic volume confronting today's security infrastructure, not to mention the latency sensitivity of many modern applications.

Recognizing these challenges, Palo Alto Networks set out right from the start to deliver a high-performance solution – something that cannot be said for competing products sporting bolted-on feature sets. To begin with, consideration was given to how individual capabilities could be optimized to achieve greater efficiency. Results of this effort included, among others, the decisions to employ stream-based scanning and a uniform threat signature format. Intent on realizing even greater gains, the engineering team also took a couple of notable steps at the system/platform level. In particular, they designed the next-generation firewall to have single-pass software (packet flow) and feature function-specific parallel processing. The result is Palo Alto Networks' Single Pass Parallel Processing (SP3) Architecture.

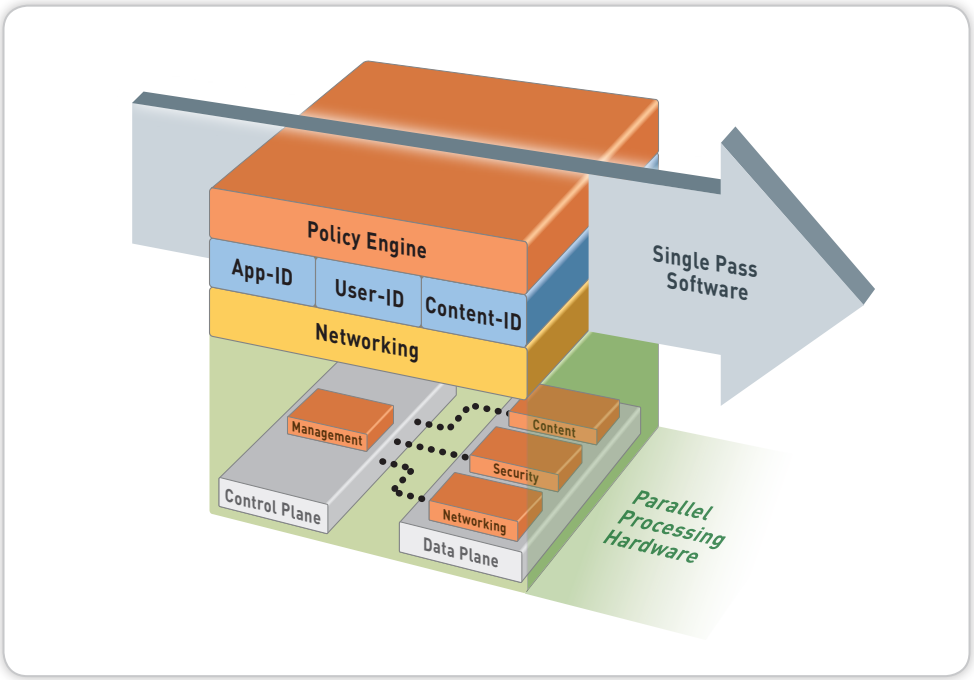

Figure 5: Single Pass Parallel Processing Architecture Marries Software and Hardware for Enterprise Performance

For conventional security products, especially those with bolted-on capabilities, each high-level security function is performed independently. The resulting multi-pass approach requires low-level packet handling and stream reassembly routines to be repeated numerous times. System resources are used inefficiently and a relatively large amount of latency is introduced. In contrast, Palo Alto Networks' next-generation firewall uses a single-pass. By design, the processing model is highly structured, essentially linear. This eliminates repetitive handling of packets and streams, dramatically reducing the burden placed on system hardware and minimizing latency.

The other major advantage of the Palo Alto Networks SP3 Architecture is that it features function-specific processing. Each next-generation firewall appliance has a control plane with dedicated CPUs, memory, and disk to support management functions. All other processing is executed by a separate data plane, which includes:

- a network processor for initial packet handling and network-layer functions;
- a multi-core security processor and hardware acceleration capabilities for standardized functions; and,
- a content scanning hardware engine.

Designed from the outset to be a high-performance solution, the next-generation firewall is able to deliver the full suite of functionality enabled by App-ID, User-ID, and Content-ID without having to make any compromises – low-latency performance is achieved for all services, even at line rate.

ADDITIONAL CAPABILITIES ENSURE AN ENTERPRISE-CLASS SOLUTION

Palo Alto Networks is keenly aware that a complete solution, in addition to overcoming the inadequacies of traditional firewalls, must also address the numerous practical issues confronting enterprises when it comes to deployment and ongoing operations. Key considerations in this regard include having compatibility with existing infrastructure, the flexibility to support a wide variety of use cases, and a high degree of reliability – not to mention being straightforward and easy to use. This is why our solution has been developed to also feature:

- A strong networking foundation, including support for L2/L3 switching, dynamic routing (OSPF, RIPv2), 802.1Q VLANs, and trunked ports;
- Flexible deployment options, including an out-of-band "visibility-only" mode, transparent in-line operation, and a fully active in-line "firewall replacement" configuration;
- Active/passive high availability with full configuration and session synchronization; and,
- Intuitive and flexible firewall management, including a command line interface, a web interface and centralized console that share the same look and feel, support for Syslog and SNMP, and extensive logging and reporting capabilities.

With a rich set of networking, integration, and systems management capabilities, the Palo Alto Networks next-generation firewall ensures IT organizations are getting exactly what they need: a robust, enterprise-class security solution.

THE NEW CORNERSTONE FOR ENTERPRISE SECURITY

As a ground-breaking, enterprise-class security solution, the next-generation firewall affords today's organizations with the opportunity to realize a number of significant benefits. From a technological perspective it helps CIOs tackle a broad range of increasingly substantial challenges by:

- Enabling user-based visibility and control for all applications across all ports;
- Stopping malware and application vulnerability exploits in real time;
- Reducing the complexity of security infrastructure and its administration;
- Providing a high-speed solution capable of protecting modern applications without impacting their performance;
- Helping to prevent data leaks; and
- Simplifying PCI compliance efforts.

Of course, it's also important to consider matters from a business perspective. In this regard, the advantages of the Palo Alto Networks next-generation firewall are that it helps organizations:

- Better and more thoroughly manage risks and achieve compliance – by providing unmatched awareness and control over network traffic;
- Enable growth – by providing a means to securely take advantage of the latest generation of applications and new-age technologies; and,
- Reduce costs – by facilitating device consolidation, infrastructure simplification, and greater operational efficiency.

The net result is that Palo Alto Networks provides today's enterprises with precisely what they need to take back control of their networks, to stop making compromises when it comes to information security, to put an end to costly appliance sprawl, and to get back to the business of making money. Providing unmatched visibility and control over applications and the threats that seek to exploit them, the Palo Alto Networks next-generation firewall is poised to take its place as the new cornerstone of enterprise security.