Re-Inventing Network Security

The Next-Generation Firewall Forms a New Foundation

June 2010

Palo Alto Networks
232 E. Java Dr.
Sunnyvale, CA 94089
408.738.7700
www.paloaltonetworks.com

Table of Contents

EXECUTIVE SUMMARY

For the last 15 years, port-blocking firewalls have been the cornerstone of network security. It's no secret, however, 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 remedy the situation. But that hasn't really worked. Neither have attempts to bolt application awareness and control onto existing firewall products, or to consolidate firewall helpers with a Unified Threat Management (UTM) device. Applications and threats are still making their way around these so-called solutions, frustrating IT groups that have only managed to incur additional cost and complexity without fixing the problem.

Palo Alto Networks, the network security company, is re-inventing network security, starting with the firewall. By focusing on applications, users, and content – not ports and protocols – as the key elements to deliver visibility and control, Palo Alto Networks' next-generation firewalls allow enterprises to safely enable modern applications, without taking on the unnecessary risks that accompany them. Built on a high-performance architecture and featuring a rich set of networking, availability, and management capabilities, our next-generation firewalls also help streamline and simplify network security infrastructure, often delivering a substantial reduction in cost and complexity by eliminating the need for enterprises to deploy a wide variety of additional network security products.

NETWORK SECURITY IS NO LONGER BLACK AND WHITE

The old model for network security was simple because everything was black and white. Business applications constituted good, low-risk traffic that should be allowed, while threats – and pretty much everything else – constituted bad traffic that should be stopped. The problems with this approach today are basically threefold:

- Applications have become increasingly gray – classifying types of applications as good or bad is not a straightforward exercise;
- Applications have become increasingly evasive; and,
- Applications have become the predominate target of today's threat developers.

TODAY'S APPLICATIONS ARE MOSTLY GRAY – NOT BLACK OR WHITE

Over the past decade, the application landscape has changed dramatically. Corporate productivity applications focused on automating key business processes have been joined by a plethora of personal productivity applications for improving the efficiency of individual users and "lifestyle" or consumer-oriented applications that enable people to handle their non-work affairs and maintain online personas. The problem with this is not the growing diversity of applications itself, but the inability to strictly and consistently classify them as good or bad. Although a few are clearly good (low risk, high reward), and a few are clearly bad (high risk, low reward), most lie somewhere in between. Moreover, which end of the spectrum they fall on can vary from scenario to scenario, even user to user or session to session.

For example, using a personal productivity application to share product documentation with a prospective customer would be "good" (medium risk, high reward), while using the same application to forward details of an upcoming release to a "friends list" that includes employees of a competitor would, most likely, not be so good. Indeed, most organizations now use a variety of social networking applications to support a wide range of legitimate business functions, such as recruiting, research and development, marketing, and customer support – and many are even inclined to allow use of lifestyle applications to some extent as a way to provide an "employee friendly" work environment and improve morale.

A modern network security solution, therefore, must be able not only to distinguish one type of application from the next, but also to account for other contextual variables surrounding its use *and* to vary the resulting action that will be taken accordingly.

TODAY'S APPLICATIONS ARE EVASIVE

The challenge does not end there. Although "distinguishing one type of application from the next" sounds simple on paper, it really isn't – for a number of reasons. The first of these is that to maximize their availability and effectiveness, many personal productivity and lifestyle applications were designed from the outset to circumvent traditional 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.

A second, vexing issue is that many new business applications are now being designed to take advantage of these same 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.

Finally, there is also the webification of enterprise applications. Standard client-server applications are steadily being re-designed to take advantage of Web technologies at the same time that enterprises are increasingly embracing cloud-based Web 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 network security 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.

TODAY'S THREATS ARE COMING ALONG FOR THE RIDE

As if there weren't already enough problems for conventional network security tools, threat developers have also turned up the heat. Intent on making money – not just building reputations – it is no wonder today's hackers have set their sights on applications. Threats that are designed to operate at the application layer can pass right through the majority of enterprise defenses, which have historically been built to provide network-layer protection. And let's face it, any resource that involves a couple hundred million users and promotes information sharing – which is the case with many of today's most popular social networking applications – is going to be an attractive target. Not only that, but the evasion techniques built into these and many other modern applications are being leveraged to provide threats with "free passage" into enterprise networks.

IT IS NO LONGER IN CONTROL

The impact of all these changes is that within many organizations IT has lost control – not to mention that the security department is running the risk of being marginalized. The inability of their existing security infrastructure to effectively distinguish good/desirable applications from those that are bad/unwanted leaves most shops with no reasonable option. Basically, they are stuck with an all or nothing decision. Do they take a permissive stance, an approach that ensures the accessibility of important applications but also allows unwanted ones, along with many threats, to proceed as well? Or do they fall into a pattern of always saying "no" – including to the growing population of highly useful gray-area applications – in order to maintain a high state of security?

What they need instead is the confidence to say "yes" to requests made by the business based on having the ability to exert granular control and provide in-depth protection down to the level of individual applications. Unfortunately, however, traditional network security devices, especially firewalls, have failed to keep pace with the continuing changes to the application and threat landscapes.

LEGACY PORT-BLOCKING FIREWALLS FAIL TO MEASURE UP

Because they are deployed in-line at critical network junctions, firewalls see all traffic and, therefore, are the ideal resource to provide granular access control. The problem, however, 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 they 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.

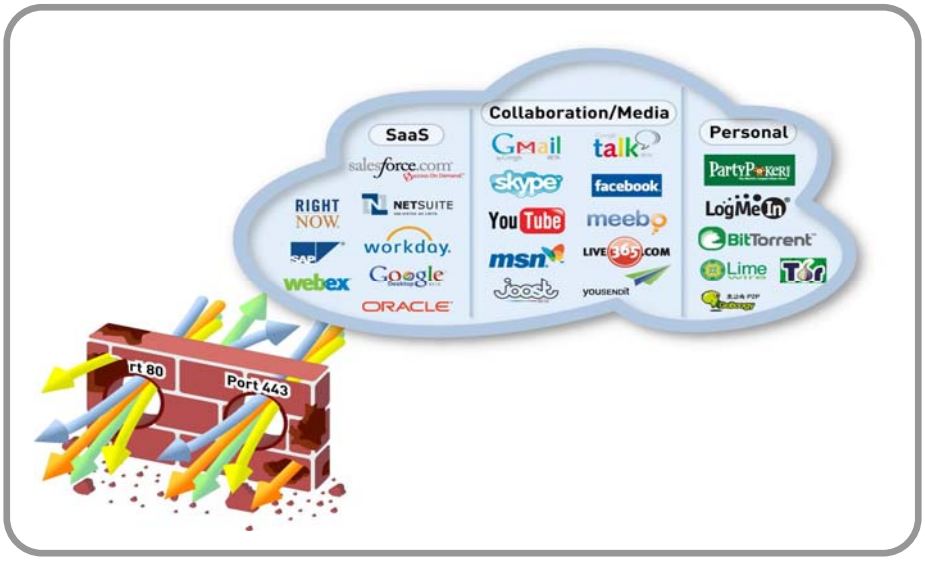

Figure 1: Port Blocking Firewalls Can't See or Control Applications

The net result is that traditional, "port-blocking" firewalls have basically gone blind. Besides being unable to account for common evasion techniques such as port hopping, protocol tunneling, and the use of non-standard ports, they simply lack the 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 network traffic, even though it corresponds to legitimate business activities, should be blocked because it includes malware or other types of threats.

On top of everything else, their control model is typically too coarse-grained. They can either block or allow traffic, but offer little variation in between to craft a more appropriate response for all of the "gray" applications enterprises would ultimately like to support – for example, by allowing certain functions within an application but not others, allowing but also applying traffic-shaping policies, or allowing based on users, groups, or time of day.

FIREWALL REMEDIES HAVE FAILED

It doesn't really help matters that the 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 appears 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

Over the years, 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 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 involve multiple policy models. 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.

UTM ONLY MAKES WHAT IS BROKEN CHEAPER

Another potential remedy that has emerged in recent years is Unified Threat Management devices. The primary advantage of the UTM solution is that it typically does a reasonable job of addressing the issues associated with device sprawl. Instead of having all of the "helper" countermeasures deployed as separate devices, with UTM they all come in one physical package.

But so what? The result is really no different than the bolted-on approach and, therefore, exhibits the same deficiencies. Inadequate application classification and resulting blind spots in the inspections that are performed remain as fundamental problems, while performance and policy management issues are compounded even further based on having to account for multiple additional countermeasures instead of just one (i.e., DPI).

IT'S TIME TO RE-INVENT NETWORK SECURITY

The bottom line is that network security in most enterprises is fragmented and broken, exposing them to unwanted business risks and ever-rising costs. Traditional network security solutions have simply failed to keep pace with changes to applications, threats, and the networking landscape in general. Furthermore, the remedies put forth to compensate for their deficiencies have, for the most part, proven ineffective as well.

The need to address this situation is the reason security visionary Nir Zuk founded Palo Alto Networks. Indeed, the mission of Palo Alto Networks is to re-invent network security by delivering a new generation of solutions that:

- Enable IT to confidently say "yes" to whatever applications are needed to best support the business – by giving them the ability to accurately identify and granularly control applications while also preventing a broad array of threats;
- Achieve comprehensive coverage – by providing a consistent set of protection and enablement capabilities for all users, regardless of their location; and,
- Simplify security and networking infrastructure – by obviating the need for numerous standalone products.

INTRODUCING THE NEXT-GENERATION FIREWALL

As the first step to fulfilling its mission, Palo Alto Networks set out to restore the firewall as the cornerstone of enterprise network security infrastructure by "fixing the problem at its core." Starting with a blank slate, its world-class engineering 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 result of this effort is the Palo Alto Networks family of next-generation firewalls – the only solution that fully delivers on 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 extensive visibility of and granular, policy-based 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 deployments with negligible performance degradation.

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 enhanced visibility and control provided by Palo Alto Networks next-generation firewalls is made possible by 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 1000 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.

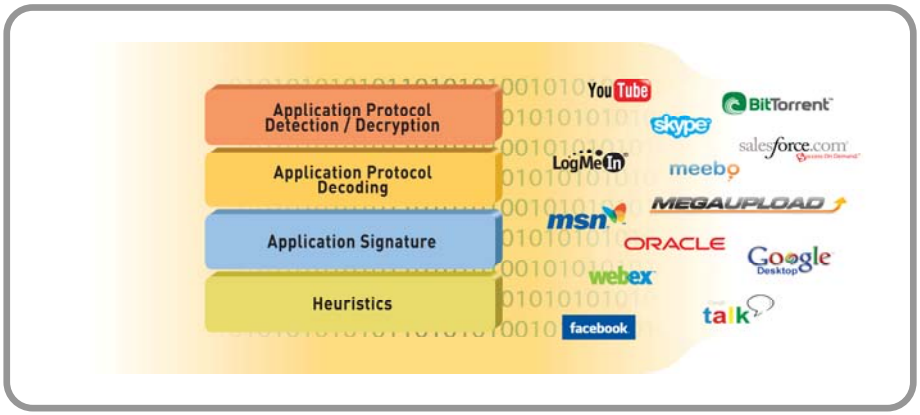

Figure 2: App-ID Identifies Applications Regardless of Port, Protocol, Evasive Tactic, or SSL Encryption

- **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 so they can make informed decisions. 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) and other LDAP directories, 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.

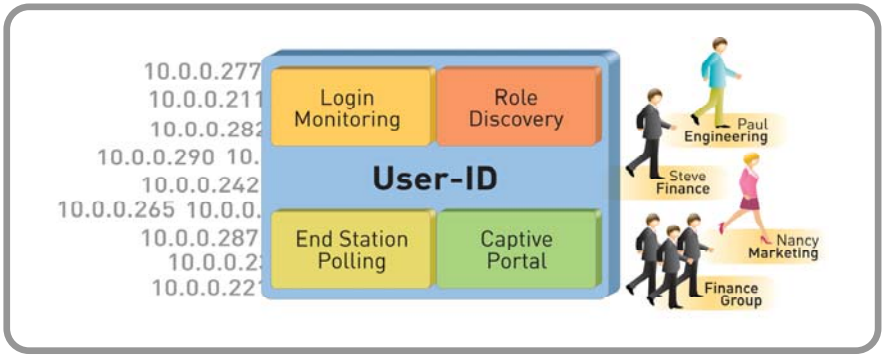

Figure 3: User-ID Integrates Enterprise Directories for User-based Policies, Reporting and Forensics

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 allows administrators to 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.

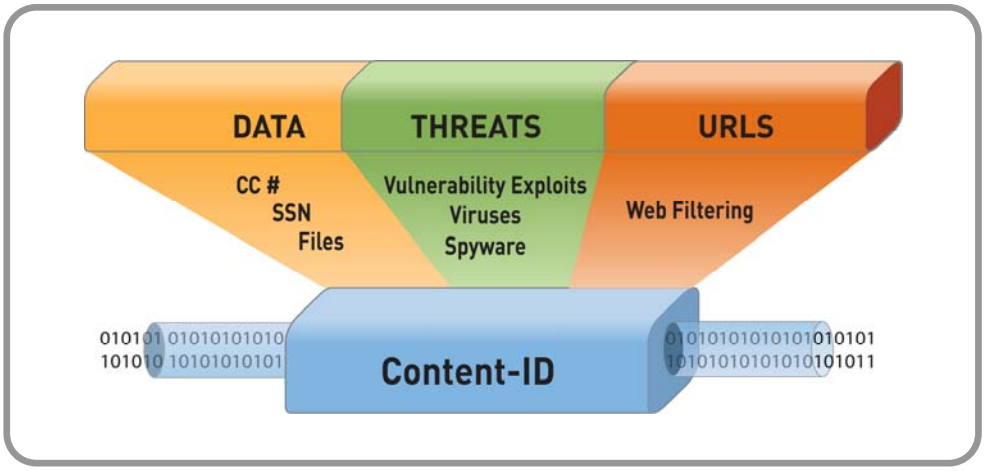

Figure 4: Content-ID Unifies Content Scanning for Threats, Confidential Data, and URL Filtering

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 from the start to deliver a high-performance solution – something that cannot be said for competing products with bolted-on feature sets. To start with, thought was given to how individual capabilities could be optimized to achieve greater efficiency. This led to the decisions to employ stream-based scanning and a uniform threat signature format. But the engineers didn't stop there. They also designed the next-generation firewall to have single-pass software and to feature function-specific parallel processing. The result is Palo Alto Networks' Single Pass Parallel Processing (SP3) Architecture.

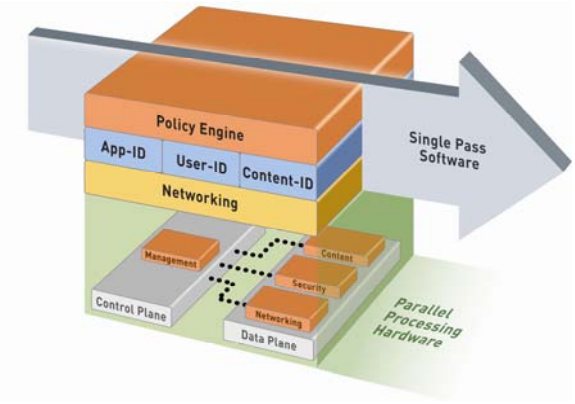

Figure 5: Single Pass Parallel Processing Architecture Marries Software and Hardware for Enterprise Performance

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, the Palo Alto Networks next-generation firewall uses a single-pass. By design, the processing model is highly structured and essentially linear. This eliminates repetitive handling of packets and streams, dramatically reducing the burden placed on system hardware and minimizing latency.

As for the function-specific parallel processing capability, 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 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 (BGP, OSPF, RIPv2), 802.1Q VLANs, trunked ports, and a full set of bandwidth monitoring and traffic shaping capabilities;
- Standards-based IPSec and SSL VPN capabilities, for secure site-to-site connectivity and remote user access, respectively;
- 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. Not only that, but in conjunction with the core identification technologies and a high-performance architecture, these are what allow enterprises to simply their network security infrastructure – to effectively get by with fewer, standalone products if they so choose.

WHAT ABOUT MOBILE AND REMOTE USERS?

As a second step in the fulfillment of its core mission, Palo Alto Networks is next taking on the challenge of providing visibility and control for users that are operating remotely, beyond the boundary established by enterprise firewalls. The goal in this case is to deliver a solution that provides the same degree of protection and application enablement received by users on the local network without having to manage a completely independent set of policies. Another major objective is to avoid the limitations and disadvantages associated with the current crop of solutions in this area.

- Endpoint security suites – Distribution and installation are often problematic, while overloaded feature sets typically create challenges in terms of client-side performance, resource requirements, and ongoing administration.

- Cloud or CPE-based proxies – Associated Web services and products typically focus on a narrow traffic stream (e.g., port 80/HTTP only), can have a limited set of services/countermeasures (e.g., URL filtering only, malware filtering only), and – because they rely on a proxy architecture – often have to allow many applications to bypass their filters to avoid breaking them.

- Backhaul via VPN technology – Whether it's IPSec or SSL-based makes little difference. There is an inevitable bump in latency as client traffic is directed back to one of a few central sites where the VPN gateways are typically located. Of even greater concern, however, is the lack of application visibility and control characteristic of the head-end devices that are subsequently used to disposition and filter this traffic.

In comparison, the Palo Alto Networks solution that is planned for delivery in the second half of 2010 relies on a persistent client that can be installed on-demand. Like the VPN-based approach, remote traffic will be sent over a secure tunnel. The difference in this case is that the connection will automatically be made to the *nearest* Palo Alto Networks next-generation firewall – whether it's deployed at one of an organization's hub facilities, out in a regional or branch office location, or as part of a public/private cloud implementation. The latency impact will thus be minimized, and the user's session will be protected and controlled by the full portfolio of Palo Alto Networks application-, user-, and content-oriented identification and inspection technologies – exactly like it would be if the user were operating on the local network instead of remotely. The net result is an easy-to-implement, easy-to-implement solution that provides remote and mobile users with the same degree of application enablement and protection as their in-office counterparts.

NETWORK SECURITY: RESTORING EFFECTIVENESS = RESTORING VALUE

With the introduction of its family of next-generation firewalls, Palo Alto Networks began the process of re-inventing network security, of restoring effectiveness and simplifying security infrastructure. The result is a market-leading solution that allows CIOs to 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; and,

- Helping to prevent data leaks.

Considering matters from a business perspective, the Palo Alto Networks next-generation firewall also 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.

And with its forthcoming client security solution, Palo Alto Networks continues the re-invention of network security, effectively extending the scope of its solution and the above benefits such that they are also applicable to remote and mobile users.

The net result is that Palo Alto Networks is providing 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. By delivering unmatched visibility and control over applications and the threats that seek to exploit them, network security solutions from Palo Alto Networks are substantially raising the bar for effectiveness and efficiency while establishing a new foundation for enterprise security.