APPLICATION-CENTRIC POLICIES

Technical Differentiator Brief

August 2018

paloalto NETWORKS®

Introduction

Before an organization can control access to its digital assets, IT teams must be able to identify what is actually happening on their network. The proliferation of mobile and cloud applications and encrypted traffic has made it very hard to identify what users are doing on the network. Attackers leverage the inability of IT teams to distinguish what is happening on a network to their advantage.

CONTENTS

THE CHALLENGE

Stateful inspection, the basis for most of today's firewalls, was created at a time when applications could be controlled using ports and source/destination IP addresses. The strict adherence to port-based classification and control is foundational and cannot be turned off. Even when augmented by "after the fact" classifiers, applications cannot be effectively controlled.

Security professionals are forced to open commonly used ports and then try to block access through those ports based on the application. This approach is risky, difficult to manage, and at the end of the day, ineffective. Further, attackers continue to nest encrypted command and control sessions within other encrypted tunnels and adapt the newest protocols to exfiltrate data. Stateful inspection based on ports is not only outdated, but it has become an attack vector. Lack of visibility into who is using what application on the network makes it impossible to ensure only desired traffic is occurring. When IT teams open trouble tickets based on ports, they allow or block applications with firewalls, not routers that block using IP addresses and ports. That port-centric approach uses legacy controls and thinking to address modern problems. With only 65,535 ports and millions of applications, we must acknowledge that opening a port is not the same as allowing a single application.

HOW IT SHOULD BE AND WHY

Instead, there should be visibility and policy control at the application level and down into the functions of applications that involve file movement, such as uploading and downloading. To control what is happening in an organization, IT teams should be able to control who accesses an application and when as well as what content is permitted to flow, and that should be true whether that application resides locally, in a cloud service, or elsewhere on the internet. The goal of policy control is to restrict the attack surface so IT teams can narrow their focus on what matters most.

To maintain that narrowed focus, IT teams should not be required to deny all of the bad applications; by default, they should be able to allow only the desired applications so they don't have to keep up with the proliferation of new applications that leverage common protocols and ports. In addition to enabling explicit protection an organization's approved applications and key data assets, this level of control allows IT teams to enable employee access to popular consumer applications, such as mail and cloud storage, without exposing organizational data to accidental release. Such restricted access prevents employees from trying to find a way around it and creating additional shadow IT problems.

How It Works

Palo Alto Networks App-ID™ is the solution. Our App-ID technology is independent of port and protocol. Application-centric policies based on App-ID are a giant leap over port-based policies because many applications share the same port or use non-standard ports. They simplify what an administrator must know, increase policy comprehension, align directly to desired business outcomes, and reduce training and maintenance. More importantly, App-ID changes how applications are permitted—it allows for safe application enablement because it does not assume that an application uses a specific port or range. No other vendor uses App-ID in initial security policy match and its base offering. Instead, other vendors open a port to all traffic and then require intentional blocking of every undesired application that uses the port. The policy definition strategy of blocking all undesired applications follows a negative-enforcement model, also known as *blacklisting*. Safe application enablement selectively allows network traffic based on applications, users, and content. It supports granular control over functionality within an application, allowing unique policies for different groups users and controlling what content they may share or move. Policy rules that permit only application-specific traffic tied to User-ID™ and Content-ID™ while denying all other traffic follow a positive-enforcement model, commonly referred to as *whitelisting*. Furthermore, we've found that policies without App-ID and User-ID are ineffective at defining a breach prevention posture, which is required to successfully protect critical applications and data. Inline IPS signatures do not replace the need to follow a positive enforcement model.

Figure 1 Requirements of an effective prevention posture

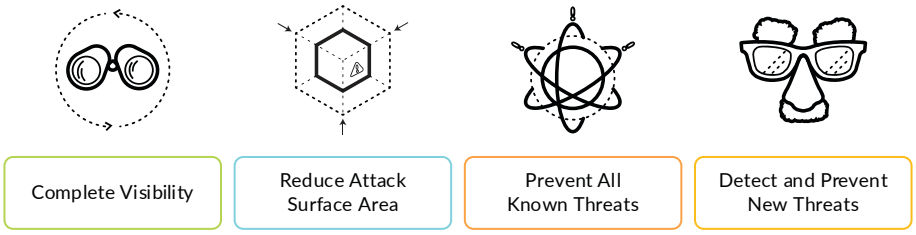

The Palo Alto Networks® Security Operating Platform classifies all applications regardless of network protocol, port, encryption, or any bypass techniques that applications employ. That information, the App-ID, becomes the basis for all policies and inspections that are performed, and because we can identify users, content, and data associated with each session, we can also identify "gray-area" applications that may be good or bad depending on the circumstances. Furthermore, App-ID is a key attribute of session logs, which can inform analysts about the applications that users and devices are using. This data also informs Application Framework apps, such as Magnifier, providing valuable clarity into network traffic, whether expected or not.

So how does it work? Palo Alto Networks defines an *application* as a specific program or feature for which communication can be labeled, monitored, and blocked if necessary. Applications can be delivered through a web browser, a client-server model, or a decentralized peer-to-peer design.

App-ID identifies the applications traversing the firewall—regardless of port or protocol. It even determines when the traffic is tunneled in GRE, uses evasive tactics, is encrypted, and distinguishes between base applications and application functions. This context brings understanding of the applications on a network, their business value, and their associated risks. Additionally, App-IDs are dynamically updated weekly with 3 to 5 new applications typically added based on input from customers, partners, and trends.

Figure 2 App-ID classification techniques

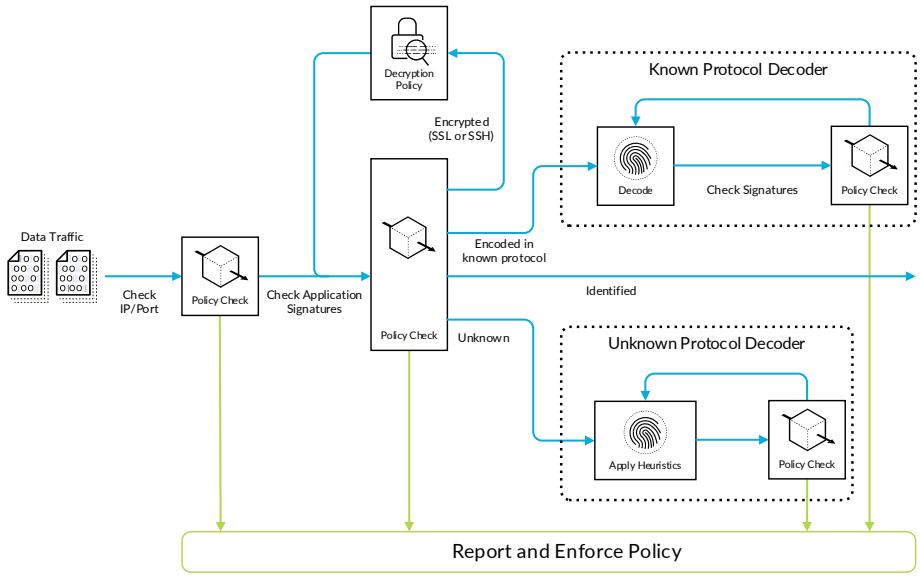

App-ID uses up to four techniques to identify the underlying application in traffic, including:

- **Application signatures**—To identify an application, signatures are used first to look for unique application properties and related transaction characteristics. These signatures also determine whether the application is using its default port or a non- standard port, which is critical because attackers commonly use non-standard ports to exfiltrate data. (We've found that 97% of day-zero exploits use FTP on non-standard port ranges.) If security policy allows the identified application, the traffic is analyzed further to identify granular applications and scan for threats. Customers can also define custom App-IDs to address their home-grown applications.

- **Application and protocol decoding**—For known protocols, decoders apply additional context-based signatures in order to detect applications tunneling inside the protocols. Decoders validate that traffic conforms to the protocol specifications, and they support NAT traversal and opening of dynamic pinholes for applications such as VoIP or FTP. Decoders for popular applications also identify the individual functions within an application, such as file upload and download in Dropbox. In addition to identifying applications, decoders identify files and other content to be scanned for threats or sensitive data.

- **Heuristics**—In certain cases, evasive applications cannot be detected using advanced signature and protocol decoding. In those cases, heuristic or behavioral analysis is used to identify applications that use proprietary encryption, such as peer-to-peer file sharing. Heuristic analysis is used with the other App-ID techniques to provide visibility into applications that might otherwise elude identification. The heuristics are specific to each application and include checks based on information such as the packet length, session rate, and packet source.

- **TLS/SSL and SSH Decryption**—If App-ID determines that TLS/SSL encryption is in use, it can decrypt and re-evaluate the traffic. A similar approach is used with SSH to determine whether port forwarding is being used to tunnel traffic over SSH.

Other key technologies that support App-ID and a positive enforcement model are User-ID, Content-ID, and quality of service (QoS). User-ID helps solve user mobility problems. Whether a user is working at home over VPN or roaming the campus (or another site where each connection assigns a local address), User-ID applies a consistent policy. It also only requires a single policy, not one for each location. User-ID allows customers to map a specific user or group of users to a specific application-based policy, providing users access to only the applications they need for their role and decreasing the attack surface.

Content-ID further reduces the risks associated with the transfer of unauthorized files and data. It extends our positive enforcement model by enabling:

- **URL filtering**—Control end-user web activities and block access to known malicious URLs and phishing links.

- **File blocking by type**—Block a range of file types by looking within the payload to identify the file type (as opposed relying on file extensions). File blocking works in either direction—upload or download.

- **Data filtering**—Filter sensitive data patterns, such as credit card or Social Security numbers, in application content or attachments.

- **File transfer function control**—Allow application use yet preventing undesired inbound or out- bound data transfers within an individual application.

- **Threat and known/unknown conversion**—A key capability is to block all “known bad” traffic (IPS, AV, command and control, spyware) in the context of the application, as well as leveraging WildFire® to rapidly turn “unknown bad” traffic or files into “known bad” signatures and preventions that address every possible link in the malware attack lifecycle.

Figure 3 Addressing each stage of an effective prevention posture

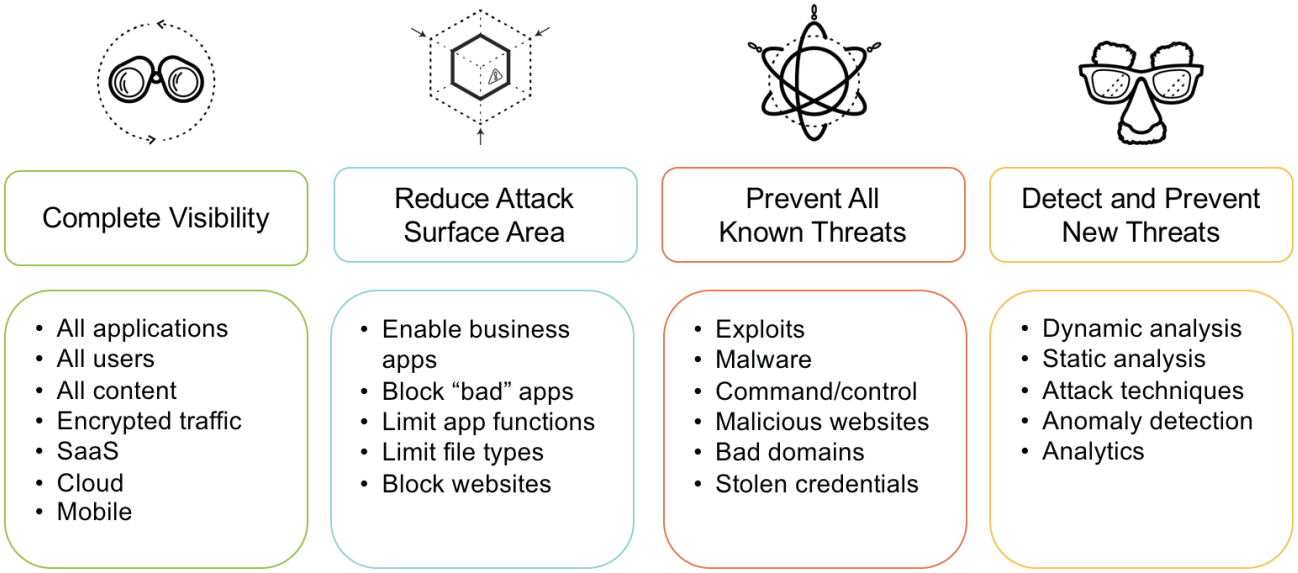

A security policy rule can include a number of conditions: application, origin of content by region, time of day, users or groups. All App-IDs are classified by category, subcategory, technology, and risk rating. Security policies can use these classifications as conditions. In addition, a rule can include a number of actions: allow or deny, allow but scan the content for exploits, viruses, and other threats, control file or sensitive data transfer, or allow or deny a subset of application functions. With a positive enforcement model, we can greatly reduce the attack surface and define an effective prevention posture.

QoS technologies work to dependably run high-priority applications and traffic under limited network capacity. Although many dedicated traffic-shaping products with application-knowledge exist, they fail to keep up with the high-speed queue management requirements. More importantly, they are disconnected from other policy decisions, such as inspection and filtering, which means that traffic is inspected twice, latency increases, and policy rules must be maintained across two or more separate products.

Our QoS policy rules allow customers to match critical business applications (including home-grown applications) and to prioritize that traffic relative to other applications. As with security rules, the Palo Alto Networks firewall provides this capability by integrating the features App-ID and User-ID with the QoS configuration. We maintain an application-centric view. App-ID and User-ID entries that exist to identify specific applications and users are available in the QoS configuration so that customers can easily specify applications and users for which to prioritize and/or guarantee bandwidth.

Technical Benefits

All App-ID signatures are applied against all traffic, all of the time. Our strength is in accurately identifying applications and application functions combined with our ability to turn on or off those functions within security policy. App-ID is not about blocking. It is about identifying, allowing, and controlling applications in a safe and understood way. No other vendor enables control as extensively or as simply as we do.

Technical benefit	Realized value
Improve prevention posture. Security policies that feature App-ID, User-ID, and Content-ID attributes are more effective at reducing the attack surface than policies that focus on ports and IP addresses. The result is that customers can define an effective breach prevention posture.	**Protection**. With effective breach prevention postures, customers experience much more effective protection. Also, the policies required to enable a positive enforcement model are easier to develop and maintain than port-based rules and result in fewer rules to maintain, which reduces complexity and risk.
Accuracy that's always on. To ensure accurate matching, App-ID looks at packet headers and the session payload, not just the session header. In addition, every App-ID is always enabled and classifies every packet, which ensures the detection of application shifts (e.g., open Messenger within Facebook). App-ID is a key part of our security policy selection, which allows multiple policies that rely on the same or share some underlying ports. Separate rules can allow multiple applications to share ports rather than opening the port to all traffic.	**Confidence**. Customers can be confident that their policies have the visibility and control to ensure that their business applications can operate without creating gaps that nefarious applications can hijack. It's not open ports; it is application-specific awareness. And because all App-IDs are always on with no performance hit, customers do not have to select which ones to disable.
Control the unknown. Rather than opening up protocol/port pairs for proprietary applications, the App-ID engine provides an accurate way to define custom apps that initially appear as unknown TCP, UDP, or P2P (peer-to-peer) traffic. Custom App-IDs also allow customers to control/prioritize short-lived applications, such as March Madness, World Cup, or Olympic Games. Custom App-IDs enable the same visibility and control as predefined App-IDs. It's not just unknown applications, but unknown content. WildFire provides advanced analysis and prevention for highly evasive zero-day exploits and malware, effectively detecting emerging threats and distributing preventions in under 5-minutes.	**Peace of mind.** Because App-ID excels at classifying traffic, any unknown traffic highlights what customers should investigate. With custom App-IDs, proprietary applications can be identified in order to reduce the amount of unknown traffic. Custom App-IDs also provide granular control over that traffic and can detect applications tunneling within your customers' applications. Our devices are automatically updated with new zero-day preventions via WildFire, ensuring that emerging threats, detected anywhere in the world, won't find their way into the customer's network. It all happens with no effort from their staff.
A single, consistent user interface. To define security policies at the firewall, across multiple firewalls, or across an organization, remote sites, or cloud implementations, Panorama™ and the PAN-OS® user interface presents a single, consistent view of the data and policy rules.	**Speed of change.** Customers neither have to coordinate to make a policy change nor make similar changes across different products to affect that change, whether it enables an application or prioritizes its availability. They make the change once, in a single user interface.

Table continued on next page

Continued table

Technical benefit	Realized value
Granularity. We offer granular controls, including visibility and control over tunneling applications and application functions (functional App-IDs). These controls ensure that customers can allow and block traffic as it relates to their business. For each allowed application, they enable QoS settings, threat filtering, URL filtering, file blocking, and data filtering. Using user and group information, they can also control who is allowed to use these applications.	**Control**. Customers can achieve the right level of control to ensure application access matches the needs of their organization and to quickly reduce the attack surface.
The language of today's networks is applications. Security policies need to speak the same language, rather than one based on protocols, ports, and technologies. It used to be about bits flowing over ports, but today, the language of the network is applications. App-ID enables customers to focus on applications.	**Aligned**. Customers can align the security policies to their business objectives and desired outcomes. Application-centric policies make conversations, traceability, and compliance about the applications and the user, not about the technologies.
Reduced complexity. App-ID simplifies the ability to control complex applications. It provides easy control over standard and nonstandard ports (i.e., any vs. application-default) and ensures that a single App-ID covers the application, which can include multiple releases/versions. It further reduces complexity via explicit dependencies (and in some cases, implicit) to reduce the number of rules required to enable complex applications. When App-IDs are updated, no policy change is required.	**Simplicity**. Customers don't have to figure this out. Easy port control helps them reduce the occurrences of unknown threats operating in their environment that use standard applications, such as FTP, in non-standard ways to avoid detection. Application-centric policies are central to positive application enablement. Customers create policy rules that are simple to define and use, easy to understand and learn, and consistent across all places in the network. Applications that have granular controls are organized and presented as application functions under a meaningful container.
Granular logging details. Because App-ID understands functions within an application, it provides more granular detail for network traffic logs. Customers can learn more about what is actually happening on their networks.	**Clarity.** The value of the right logs and the right log details can't be quantified. Leveraging the Application Framework and Logging Service, customers can extract new value from older logs.

Business Impact

Although improved security posture is our primary outcome, customers can realize other business impact through cost effectiveness, operational efficiency, and leveraging technology improvements. Palo Alto Networks addresses these drivers in the following ways.

Cost savings:

- **Operational cost savings**—The ability to control and free up network bandwidth by blocking unwanted applications that are non-business related can improve ROI. Furthermore, customers can prioritize traffic based on application and user, which allows critical business applications to more effectively consume available bandwidth.
- **Cost savings through consolidation**—The Palo Alto Networks Security Operating Platform provides operational efficiency through consolidation via a single platform and management interface and reduced hardware/software costs and maintenance. This requires fewer staff certifications, less lab equipment, and fewer API integrations.

Simplification:

- **Adapting security to address changing business models**—The Palo Alto Networks Security Operating Platform allows customers to quickly apply role-based policies across multiple deployment models in order to provide consistent security everywhere. These policies allow consistency whether it is private infrastructure, private cloud, public cloud, or physical or virtual.
- **Improve security operations**—Security staff can cross-train and quickly on-board because the Palo Alto Networks Security Operating Platform is easy to deploy and manage. This time savings is possible because of better policy-comprehension and reduced training and maintenance. The ability to migrate legacy rules is easier because we classify traffic and applications.
- **Improve visibility and awareness of traffic**—With the ability to view and classify all traffic, customers can create a more secure environment dynamically. They can prioritize what is allowed, not allowed, or tolerated to meet business requirements and risk tolerance.

Operational efficiency:

- **Deploy, scale, and secure new applications faster and consistently**—The ability to identify specific applications, whether off-the-shelf, SaaS, or custom applications, allows customers to quickly and securely deploy new business applications. Ease of operation and deployment is achieved through simple classification and through platform updates that identify new or changed applications.
- **Rapidly on-board new users and partners**—The ability to leverage users and user groups provides simple, role-based policies with minimal disruption or concern for user mobility. Customers can virtually segment networks based on user groups as the users transition to new locations on campus. This saves time and money and reduces complexity.
- **Simplify licensing**—Palo Alto Networks provides ease of licensing. This is realized through the ease of purchase and the ability to quickly deploy and maintain licenses through automatic updates. It's all backed by a world-class support organization.
- **Simplify application management**—The ability to consolidate and reduce complex security rule sets by up to 60% can save time and resources compared to stateful firewalls. The ability to intelligently organize policy with application filters and application groups further reduces the proliferation of redundant rules.

Example Scenarios

This section highlights some common uses cases where App-ID provides clear advantages and highlights why it useful in such scenarios.

- **Rapidly adopt new SaaS application models—without expanding the attack surface on the network**. Common enterprise applications are moving to cloud-only models. But defining these applications in terms of the ports and protocols is increasingly difficult. Usually, an application is accessed via a URL. However, content delivery networks and dynamic DNS have rendered control via the URL obsolete. To control traffic for these applications, we must use the application itself. App-ID matches against the traffic, not a URL or port. Customers must keep users safe no matter what intermediate technology occurs between users and their critical business-applications.

- **Reduce time to migrate and deploy applications in the data center**—Customers can control the unknown TCP/UDP until they develop custom App-IDs. Data center applications increasingly drive business. They improve efficiency and productivity, generate revenue from end customers, and facilitate communications and financial interconnection with business partners. Because providing security alongside application uptime is so difficult, the data center has become the most dangerous place in the network. It typically has the weakest security—either none or port-based firewalls only. At the application level, logical segmentation of east-west traffic requires security systems that dynamically integrate with external systems to provide resource definitions based on business need, simple IP, or VLAN segments. The lack of east-west control combined with the ever-growing list of custom, multi-tiered, dynamically distributed applications in the data center exacerbates this problem. An application-centric view of east-west traffic allows customers to control and protect these flows. With App-ID, customers can control the unknown TCP/UDP/P2P traffic until they can develop custom App-IDs to reduce the attack surface.

- **Day 1 protection of custom applications and data**—Custom App-IDs ensure that only desired applications run across commonly used ports. With custom App-IDs, customers can start instituting security controls—at the application level, not port level—during the application development process. This approach natively and unobtrusively allows customers to provide automated security policy and control when introducing a new application into production. They can secure it on day one, expedite the deployment, and protect sensitive data.

- **Reduce risk of shadow IT**—Safely tolerate SaaS applications, such as Facebook and Gmail. With App-ID, customers can provide users read-only access to such applications without letting them upload files (closing off a risky application function from the customer's perspective). Read-only access supports the users' need to keep up with those important in their lives—without putting the customer at risk. This limits the driver behind shadow IT and still achieves security objectives.

- **Simplify operations and training**—Customers can leverage consistent data center and public-cloud deployments. Application-centric policies can efficiently segment traffic into and out of the data center and between servers. Such policies support on-premise data centers and public-cloud deployments. With Palo Alto Networks, customers can leverage the same OS, policies, and features and manage using a single central management system. To prevent unauthorized users from instantiating VMs and unwittingly creating new avenues for data leaks and attacks, customers can use App-ID to treat the public-cloud management consoles as an application, allowing control over who can bring up public-cloud deployments in AWS and Azure containers.

- **Enable role-based applications**—When protecting their organization against large scale data leaks, customers can:

  - Enable critical-business functions, such as marketing, HR, and customer support teams.

  - Require unique levels of access to applications, such as Twitter, LinkedIn, and Facebook.

With role-based access to sanctioned SaaS marketing applications, customers can, for example, allow members of their marketing team to post and upload graphic files (but blocking PDF and Office files) to Twitter, LinkedIn, and Facebook, while allowing all other teams to have best-effort, read-only access on a best-effort basis. App-ID can control key sub-application functions, giving customers choice in control.

- **Reduce the attack vector for users who access public networks and resources**—The internet gateway determines how exposed a network is to the outside world. The only way to reduce the attack surface and still allow users to be productive is to focus on applications, not ports. Outbound port-hopping at the perimeter, as is common in malware, is not a factor if a security device is port-agnostic and designed to elevate applications to a key condition of the positive-enforcement model policy.

- **Enable multi-faceted applications with ease and security—**Enterprise applications, such as Microsoft Lync or Active Directory, often incorporate dozens of application functions and require dozens or hundreds of open ports. Controlling these applications at anything other than the application/function level has become so difficult that many organizations have simply given up. By decoupling the application from the protocol/port number, customers can define fewer policies that are easy to understand, maintain, and enable for new users. If the application developer makes changes, policies remain in effect with no interruptions to users.

- **Rapidly deploy seasonally affected or short-term applications**—Customers can securely and confidently burst internal data center applications and avoid expensive unnecessary hardware buildouts. When an internal application is affected by seasonal usage, such as holiday sales, customers can burst into public infrastructure to address temporary compute requirements. They can control access into multiple availability zones and provide protection among all container instances, mimicking the segmentation and controls in the data center. This scenario ensures consistent policies, user control, application tier controls, and secure remote access. Consistent application-centric policies, based on App-ID and User-ID, ensure identical traffic restrictions among application tiers and consistent user restrictions, regardless of location. This consistency remains even when bursting into multiple providers.

Headquarters

Palo Alto Networks
4401 Great America Parkway
Santa Clara, CA 95054, USA
www.paloaltonetworks.com

Phone: +1 (408) 753-4000
Sales: +1 (866) 320-4788
Fax: +1 (408) 753-4001
info@paloaltonetworks.com

B00905P-18A-1