



App-ID™ Application Classification Technology Overview

June, 2007

Palo Alto Networks
2130 Gold Street, Suite 200
Alviso, CA 95002-2130
Main 408.786.0001
Fax 408.786.0006
Sales 866.207.0077
www.paloaltonetworks.com

Table of Contents

Introduction.....	3
The Need for Application-Centric Traffic Classification.....	3
App-ID Application Classification Technology.....	4
How App-ID Works	4
An App-ID Example	6
Policy-Based Application Usage Control	6
Conclusion.....	7

Introduction

Security experts will admit to the fact that without visibility into which applications are traversing their network, they are unable to control, much less protect the network as well they would like. Enterprise networks are being populated by a new generation of end-user applications, both personal and business oriented, that are designed to evade detection by existing firewalls. At the same time, well-meaning corporate applications are utilizing similar tactics to accelerate deployment, facilitate wide spread access and minimize disruption. The ramifications resulting from the inability to identify, control and inspect the applications traversing the network include:

- Business liability: Regulatory and internal policy compliance, unseen data leakage
- Operational expenses: Rising bandwidth consumption, added IT operational expenses, lost user productivity from personal application usage
- Increased propagation of threats: Viruses, spyware, worms and application vulnerabilities can quickly spread across the network

Most network administrators are aware that these applications are roaming the network and are managing as best they can with a patchwork of existing technologies. The key challenge they face is that their firewall, as originally defined, uses port and protocol to classify and control what gets in and out of the network. What's needed is a fresh approach to the traffic classification within the firewall, one that identifies the actual application and is capable of bringing policy-based controls back to the network security team.

The Need for Application-Centric Traffic Classification

The key challenge facing network administrators is that their firewall, as originally defined, uses protocol and port (a key component of Stateful Inspection) to identify and control what gets in and out of the network. This port-centric design is relatively ineffective when faced with new applications that are equipped with increasingly sophisticated security evasion techniques, such as dynamic or random port numbers, application emulation and SSL encryption. The fundamental reasons that Stateful Inspection can be easily evaded include:

- Fixed ports: Stateful Inspection assumes that the port(s) used for an application are well known, fixed and unique, therefore evasive applications can easily live up to their namesake, dynamically selecting an open port and passing quietly through the firewall, circumventing all manner of inspection.
- First packet inspection: Stateful Inspection looks only at the first packet to identify the application, therefore, applications can more easily pass themselves off as HTTP, then switch in mid stream, bypassing any attempted controls or inspection.
- Encrypted traffic: Stateful firewalls do not have decryption capabilities, therefore any SSL encrypted traffic passes through the firewall without being identified, controlled or inspected.

Current attempts to solve the application classification challenge have focused primarily on add Deep Packet Inspection (DPI) technologies to the firewall. Using DPI, which is primarily an attack detection technology, to identify and block "bad" applications is simply a patch and not the full solution. DPI is typically designed to look only at a partial set of traffic to avoid impeding performance, is not able to cover the breadth of applications needed and is managed through separate policies. Simply put, the current bolt-on

solutions do not have the accuracy, policy, or performance to solve today's application visibility and control requirements.

App-ID Application Classification Technology

Rather than classify traffic solely by protocol and port, or by bolt-on point product offerings that do not address the problem, Palo Alto Networks App-ID technology takes a completely new approach, classifying traffic from an application-centric perspective.

Whereas traditional port-based solutions use a single classification technique (protocol/port) to identify traffic, App-ID uses four of them, operating in concert, to determine exactly what applications are traversing the network irrespective of port number. Application-centric traffic classification addresses security evasion tactics such as the use of non-standard ports, dynamically changing ports and protocols, emulating other applications, and tunneling to bypass existing firewalls.

With multiple classification techniques, App-ID is able to accurately identify the application—even those that use common evasive tactics. With increased visibility into the actual identity of the application, administrators can deploy comprehensive, policy-based application access control for both inbound and outbound network traffic. Key benefits of App-ID include:

- Addresses key weaknesses in existing solutions by using multiple traffic classification techniques to more accurately identify what type of traffic is traversing the network.
- Improves security by dictating access rights based upon the actual application traffic as opposed to only the protocol and port.
- Delivers ability to enforce corporate application usage policy compliance.
- Accurate identification makes malware threat inspection more effective.

How App-ID Works

As traffic flows through the Palo Alto Networks PA-4000 Series, App-ID establishes the application session and maintains session state, while additional application identification mechanisms more accurately classify, and therefore control, the traffic. App-ID identifies traffic using as many as four or as few as one of the following classification engines.

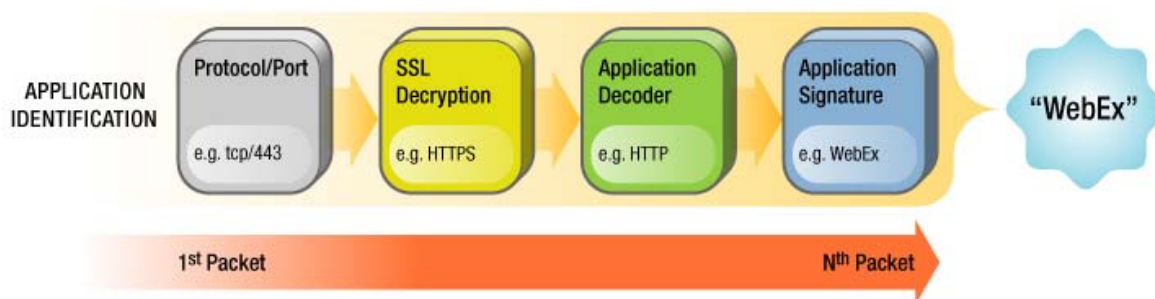


Figure 1: App-ID uses four traffic classification techniques to accurately identify the application, irrespective of port or protocol.

- **Protocol/Port:** Very simply, the protocol (such as TCP or UDP) and the port number (such as port 80) of the traffic. For App-ID, Protocol/Port information is primarily used for policy enforcement, such as allowing or blocking a specific

application over a specific protocol or port number, but is sometimes used in classification, such as ICMP traffic where the protocol is the primary classification method used.

- **SSL Decryption:** The PA-4000 Series is the first firewall on the market to enable policy-based identification, control and inspection of SSL traffic. When SSL traffic flows through the PA-4000 Series, App-ID invokes an SSL proxy to identify, decrypt and control the application traffic within the SSL tunnel. The decision to invoke the SSL proxy is performed on a per policy basis, allowing some traffic to pass through un-inspected and other traffic to be decrypted, controlled, and inspected. The policy table to determine which SSL traffic is decrypted uses a combination of source IP, destination IP, and URL category of the destination IP. The SSL proxy currently supported with App-ID is a forward proxy (defined below) which provides SSL decryption of outbound connections to uncontrolled servers.
 - Forward proxy: Inserts a proxy into an “outbound” connection where an internal user is connecting to a server, ie. an employee establishing a connection to <https://www.wellsfargo.com> or an Intranet server. The SSL decryption engine responds to the client as if it was the server and initiates a connection to the server as if it were the client. In responding to the client, an SSL certificate is dynamically created to match the destination server. This certificate will be signed by a root certificate on the forward proxy. To keep the client browser from seeing a certificate error, the root public certificate should be loaded on the client’s browser.
 - Reverse proxy: Not initially supported in the PA-4000 Series, reverse proxy is the process used to insert a proxy into an “inbound” connection where the client is not an internal user, but the server is a controlled internal server, ie. an customer connecting to a secure web server. The SSL decryption engine would proxy the SSL connection the same as the forward proxy, but in this case, would use the actual valid certificate on the destination server instead of a dynamically created certificate.
- **Application Decoders:** Application decoding in App-ID serves two purposes. First, it identifies the more complex and/or evasive applications such as Skype. It not only contains the ability to apply application signatures, but it is also able to perform more complex pattern matching operations on the traffic. Second, it is used for continuous application decoding to perform threat detection throughout the session, and can look for anomalies and changes in applications during this process.
- **Application Signatures:** Context-based signatures focus on identifying the specific applications and can apply to all traffic. These signatures look for the unique application properties and related information exchange to correctly identify the traffic regardless of the protocol and port being used. The application signatures are capable of identifying a wide range of applications even when they are tunneling over non-standard ports or emulating carrier applications such as HTTP.

The application-centric nature of App-ID means that it can not only identify and control traditional applications like HTTP, FTP, SNMP but it can also accurately identify over 400 other applications across 14 categories.

When App-ID identifies the application, the information is displayed in plain English, complete with information about the application, such as if it is pervasive, prone to misuse, able to transfer files, and other key stats. With a click of the mouse, an administrator can

drill down into details on the application itself, the threats it poses, who is using it, and how much bandwidth it is consuming. Once accurately identified, the applications can then be controlled and inspected based upon the security policy.

An App-ID Example

Using HTTP as an example, App-ID accurately identifies the specific application within HTTP using a combination of the four classification techniques. Because App-ID sees the specific such as Yahoo!IM or any one of the 400+ applications, the administrator can be very specific about what the security policy should do with that traffic. And if the application shifts to TCP port 443 to use SSL encryption to bypass the controls, App-ID can decrypt the SSL connection to allow control and inspection of Hotmail inside of SSL.

In contrast, a Stateful Inspection security policy might include a rule that states “allow all HTTP” which is telling the firewall, “allow all traffic on TCP port 80”. This would not cause a problem if all traffic crossing port 80 were web traffic, but there are many applications that use port 80: Yahoo!IM, Kazaa, Bittorrent and Hotmail are just a few of the applications that use HTTP/Port 80. Because the only thing that Stateful Inspection sees is TCP port 80, it will not identify the individual applications, and therefore, cannot control them.

Policy-Based Application Usage Control

Using App-ID to more accurately identify application traffic means that administrators can regain control over their network by deploying application usage control policies. Through the process of blocking the usage of these applications at the gateway, an IT department can reduce the manpower associated with monitoring desktops, thereby gaining operational efficiencies while mitigating security risks.

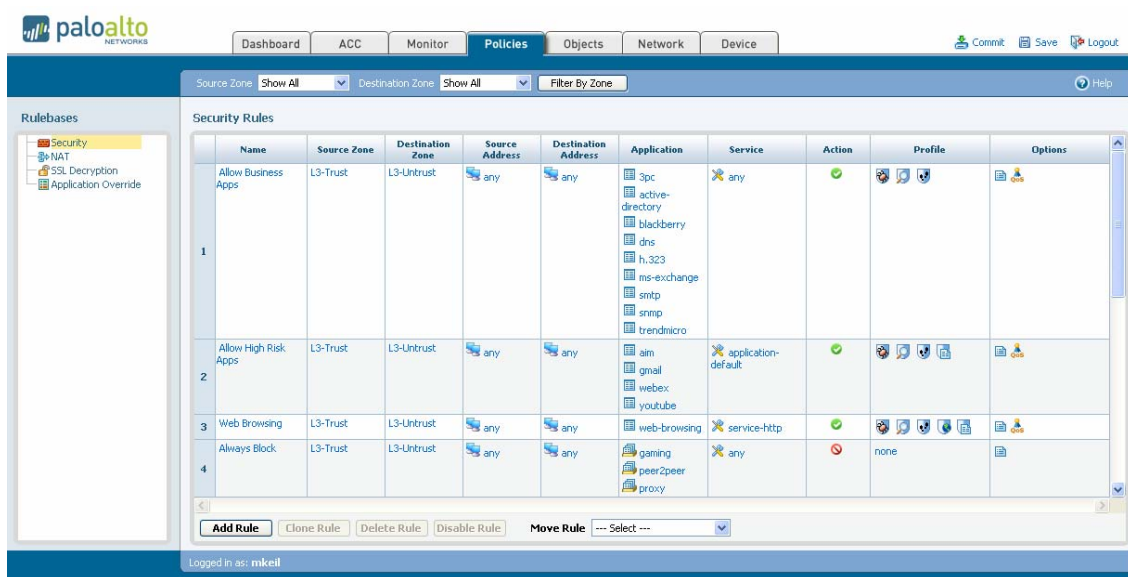


Figure 2: A familiar look and feel takes full advantage of existing firewall policy editing skills, accelerating the deployment of application usage control policies.

Application usage control policies can be enforced using a combination of the following parameters:

- By source and destination IP address, source and destination security zones, and time-based schedules
- Application type or specific application: use application categories to define policy (such as all peer2peer), or select from over 400 applications that are dynamically updated, including traditional enterprise, networking and Internet applications
- Protocol and port number: restrict and/or enforce applications to use default or specific protocols and port numbers
- File blocking: block file transfer capabilities within an application as well as block application transfers by file type

Application access control policies can be implemented to enforce appropriate usage, and traffic can be more completely inspected because of the accuracy of the application identification.

Conclusion

An IT administrator cannot implement security controls and threat prevention policies against traffic that cannot be identified. And as more and more new applications are developed with security evasion characteristics, App-ID brings new levels of visibility and control to the security team, enabling the development and enforcement of application usage control policies.

Copyright 2007, Palo Alto Networks, Inc. All rights reserved. Palo Alto Networks, the Palo Alto Networks Logo, PAN-OS, FlashMatch and App-ID are trademarks of Palo Alto Networks, Inc. in the United States. All specifications are subject to change without notice. Palo Alto Networks assumes no responsibility for any inaccuracies in this document or for any obligation to update information in this document. Palo Alto Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.