App-ID

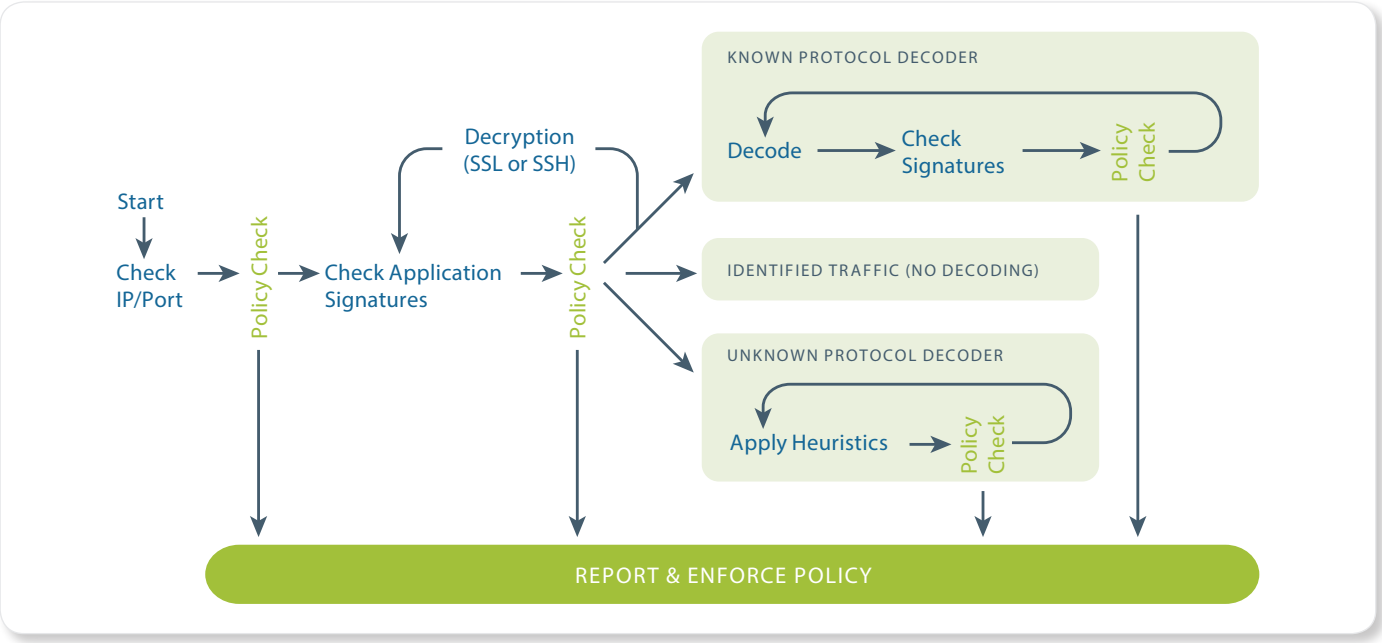

How App-ID classifies traffic.

App-ID™ is a patent-pending traffic classification technology that identifies applications traversing the network, irrespective of port, protocol, evasive characteristic or encryption (SSL or SSH).

- Facilitates more complete understanding of the business value and associated risk of the applications traversing the network.
- Enables creation and enforcement of secure application enablement policies.
- Brings application visibility and control back to the firewall where it belongs.

App-ID™ uses as many as four identification techniques to determine the exact identity of applications traversing the network—irrespective of port, protocol, evasive tactic, or SSL encryption. Identifying the application is the very first task performed by App-ID, providing administrators with the greatest amount of application knowledge and the most flexibility in terms of safe application.

As the foundational element of the Palo Alto Networks next-generation firewall, App-ID provides visibility and control over work-related and non-work-related applications that can evade detection by masquerading as legitimate traffic, hopping ports or sneaking through the firewall using encryption (SSL and SSH).

In the past, unapproved or non-work-related applications on the corporate network were summarily removed or blocked. However, in today's business environment, the response options are not nearly as clear because many of the same applications are helping employees get their jobs done.

App-ID enables administrators to see the applications on the network, learn how they work, their behavioral characteristics, and their relative risk. When used in conjunction with User-ID, administrators can see exactly who is using the application based on their identity, not just an IP address. Armed with this information, administrators can use positive security model rules to block unknown applications, while enabling, inspecting and shaping those that are allowed.

paloalto NETWORKS

the network security company™

Firewall Traffic Classification: Applications, not Ports

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 IPs. The strict adherence to port-based classification and control methodology is the primary policy element; It is hard-coded into the foundation and cannot be turned off. This means that many of today's applications cannot be identified, much less controlled by the firewall and no amount of "after the fact" traffic classification by firewall helpers can correct the firewall port-based classification.

Palo Alto Networks recognized that applications had evolved to where they can easily slip through the firewall and chose to develop App-ID, a new method of firewall traffic classification that does not rely on any one single element like port or protocol. Instead, App-ID uses multiple mechanisms to determine what the application is, first and foremost, and the application identity then becomes the basis for the firewall policy. App-ID has been created to be highly extensible and as applications continue to evolve, application detection mechanisms can be added to App-ID or updated as a means of keeping pace with the ever-changing application landscape.

App-ID Traffic Classification Technology

The first task that a Palo Alto Networks next-generation firewall executes is the identification of the applications traversing the network using App-ID. Using as many as four different techniques, App-ID determines what the application is, irrespective of port, protocol, encryption (SSL and SSH) or other evasive tactic employed. The number and order of identification mechanisms used to identify the application will vary depending on the application. The general flow is as follows:

- **Application Signatures:** Signatures are used first to look for unique application properties and related transaction characteristics to correctly identify the application regardless of the protocol and port being used. The signature also determines if the application is being used on its default port or it is using a non-standard port (for example, RDP across port 80 instead of port 3389, its standard port). If the identified application is allowed by security policy, further analysis of the traffic is done to identify more granular applications as well as scan for threats.
- **SSL and SSH Decryption:** If App-ID determines that SSL encryption is in use and a decryption policy is in place, the traffic is decrypted and then passed to other identification mechanisms as needed. If no policy is in place, then SSL decryption is not employed. Once the application is identified, and deemed acceptable by policy, threat prevention profiles are applied and the traffic is then delivered to its destination. A similar approach is used with SSH to determine if port forwarding is in use as a means to tunnel traffic over SSH. Such tunneled traffic is identified as ssh-tunnel and can be controlled via security policy.
- **Application Protocol Decoding:** Decoders for known protocols are used to apply additional context-based signatures to detect other applications that may be tunneling inside of the protocol (e.g., Yahoo! Instant Messenger used across HTTP). Decoders validate the traffic conforms to the protocol specification and provide support for NAT traversal and opening dynamic pinholes for applications such as VoIP or FTP. Decoders for popular applications are used to identify the individual functions within the application as well (e.g., webex-file-sharing). In addition to identifying applications, decoders also identify files and other content that should be scanned for threats or sensitive data.
- **Heuristics:** In certain cases, evasive applications still cannot be detected even through advanced signature and protocol analysis. In those situations, it is necessary to apply additional heuristics, or behavioral analysis to identify certain applications such as peer-to-peer file-sharing or VoIP applications that use proprietary encryption. Heuristic analysis is used as needed, with the other App-ID techniques discussed here, to provide visibility into applications that might otherwise elude positive identification. The actual heuristics used are specific to an application and include checks based on such things as the packet length, session rate, packet source, etc.

With App-ID as the foundational element for every Palo Alto Networks next-generation firewall, administrators can regain visibility into, and control over, the applications traversing the network.

App-ID: Dealing with Custom or Unknown Applications

Palo Alto Networks adds an average of five new applications to App-ID each week, yet there are cases where unknown application traffic will be detected. There are typically two scenarios where unknown traffic will appear: a commercially available application that does not have an App-ID or an internal, custom application is in use.

- **Unknown Commercial Applications:** Using ACC and the log viewer, users can quickly determine that the application is used commercially or not. Using the packet capture feature on the Palo Alto Networks firewall, customers can record the traffic and submit it for App-ID development. The new App-ID is developed, tested with the customer, then added to the database for all users in the form of a weekly update.
- **Internal or Custom Applications:** Once it has been determined with ACC and the log viewer, that the application in question is internal or custom, then customers have several options. First off, an application override can be applied, effectively renaming the application. Alternatively, customers can develop a custom App-ID for their application using the exposed protocol decoders. The protocol decoders that have been exposed include: FTP, HTTP, HTTPs (SSL), IMAP, SMTP, RTSP, Telnet, unknown-TCP, unknown-UDP, and file body (for html/pdf/flv/swf/riff/mov). Once developed, traffic identified by the custom App-ID is treated in the same manner as the previously classified traffic; it can be enabled via policy, inspected for threats, shaped using QoS and so on. Custom App-IDs are managed in a separate database on the device, ensuring they are not impacted by the weekly App-ID updates.

An important point to highlight is that Palo Alto Networks next-generation firewalls use a positive enforcement model, which means that all traffic can be denied except those applications that are expressly allowed via policy. This means that in some cases, the unknown traffic can be easily blocked or tightly controlled. Alternative offerings that are based on IPS will allow unknown traffic to pass through without providing any semblance of visibility or control.

How App-ID Works: Identifying WebEx

When a user initiates a WebEx session, the initial connection is an SSL-based communication. With App-ID, the device sees the traffic and the signatures determine that it is using SSL. The decryption engine and protocol decoders are then initiated to decrypt the SSL and detect that it is HTTP traffic. Once the decoder has the HTTP stream, App-ID can apply contextual signatures and detect that the application in use is WebEx. WebEx is then displayed within ACC and can be controlled via a security policy.

If the end user were to initiate the WebEx Desktop Sharing feature, WebEx undergoes a "mode-shift" to where the session has been altered from a conferencing application to a remote access application. In this scenario, the characteristics of WebEx have changed and App-ID will detect the WebEx Desktop Sharing feature which is then displayed in ACC. At this stage, an administrator has learned more about the application usage and can exert policy control over the use of the WebEx Desktop Sharing feature separately from general WebEx use.

Application Identity: The Heart of Policy Control

Identifying the application is the first step in learning more about the traffic traversing the network. Learning what the application does, the ports it uses, its underlying technology, and its behavioral characteristics is the next step towards making a more informed decision about how to treat the application. Once a complete picture of usage is gained, organizations can apply policies with a range of responses that are more fine-grained than allow or deny. Examples include:

- Allow or deny
- Allow but scan for exploits, viruses and other threats
- Allow based on schedule, users or groups
- Decrypt and inspect
- Apply traffic shaping through QoS
- Apply policy-based forwarding
- Allow certain application functions
- Any combination of the above

With App-ID as the foundational element, Palo Alto Networks next-generation firewalls are restoring visibility and control over the applications traversing the network to the firewall, the most strategic security component in the network security infrastructure.

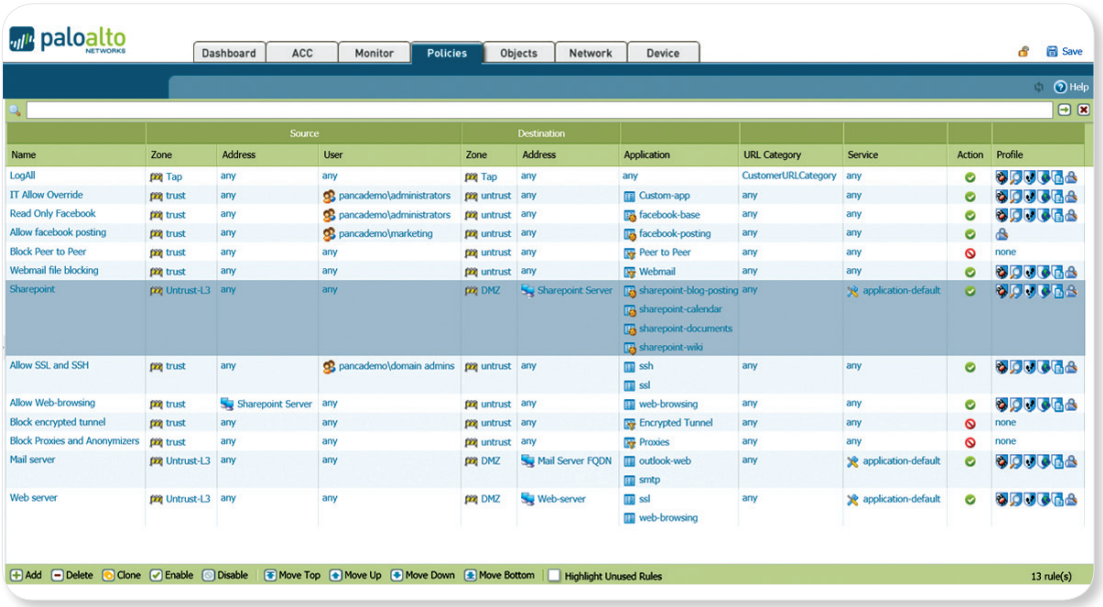

Application Function Control
Maximize productivity by safely enabling the application itself (Microsoft SharePoint) or individual functions.

Application Function Control

To many customers, secure application enablement means striking an appropriate security policy balance by enabling individual application functionality while blocking other functions within the same application. Examples may include:

- Allowing SharePoint Documents, but blocking the use of SharePoint Administration.
- Block Facebook-mail, -chat, -posting and -apps, but allow Facebook itself, effectively only allowing users to browse Facebook.
- Enable the use of MSN, but disable the use of MSN-file transfer and only allow certain file types to be transferred using the file blocking feature.

Using an application hierarchy that follows a container and supporting function model, App-ID makes it easy for administrators to choose which applications to allow, while blocking or controlling functions within the application. The graphic shows SharePoint as the container application, and the individual functions within.

Controlling Multiple Applications: Dynamic Filters and Groups

There are many cases where customers may want to control applications "in bulk", as opposed to controlling them individually. The two mechanisms that address this need are application groups and dynamic filters.

- **Application groups:** A group of applications is a static list of applications may that can be used to enable use for certain users while blocking their use for others. An example may be the use of remote management applications such as RDP, Telnet, and SSH. Each of these applications are known to be used by support and IT personnel, yet employees that fall outside of these groups are also known to use them as a means of accessing their home networks. A group of applications can be created and assigned to IT and support through User-ID, tying the groups to the policy. As new employees are added, they only need to be added to the directory group. No updates are needed to the policy itself.
- **Dynamic filters:** A dynamic filter is a set of applications that is created based on any combination of the filter criteria: category, subcategory, behavioral characteristic, underlying technology and risk factor. Once the desired results for the filter are achieved, a policy that blocks or enables and scans the traffic can be applied. As new App-IDs that fulfill the filter criteria are added in the weekly content updates, the filter is automatically updated as soon as the device is updated, thereby minimizing the administrative effort associated with policy management. The complete list of filter options are shown below.
  - **Category and Subcategory**
    - **Business:** Authentication services, database, ERP, general management, office programs, software updates, storage/backup
    - **General Internet:** File sharing, Internet utilities (web-browsing, toolbars, etc)
    - **Collaboration:** Email, instant messaging, Internet conferencing, social networking, social business, VoIP/video, web posting
    - **Media:** Audio streaming, gaming, photo/video
    - **Networking:** Encrypted tunnel, infrastructure, IP protocol, proxy, remote access, routing

Applipedia
Browse up-to-date application research and analysis at the Palo Alto Networks Application and Threat Research Center.

- Application Behavioral Characteristics
  - Able to transfer files from one network to another
  - Used to propagate malware
  - Consumes 1 Mbps or more regularly through normal use
  - Evades detection using a port or protocol for something other than its intended purpose with intent
  - Has been widely deployed
  - Application has had known vulnerabilities
  - Prone to misuse or is easily configured to expose more than intended
  - Tunnels other applications
- Underlying Application Technology
  - Client-server based
  - Browser-based
  - Peer-to-peer based
  - Network protocol

Expanding the List of Applications

The list of application App-IDs is growing rapidly with 3-5 new applications added weekly based on input from customers, partners, and market trends. Customers that find unidentified applications on their network can capture the traffic and then send the information back to Palo Alto Networks for App-ID development. Once a new App-ID is developed and tested, it is added to the list as part of the weekly content updates.

paloalto NETWORKS
the network security company™

3300 Olcott Street
Santa Clara, CA 95054

Main: +1.408.573.4000
Sales: +1.866.320.4788
Support: +1.866.898.9087

www.paloaltonetworks.com