Revolutionize NGFWs and CASB App-ID with Machine Learning

According to Gartner, "Next-Generation Firewalls"' or NGFWs are deep-packet inspection firewalls that move beyond port and protocol inspection and blocking to include application-level inspection, intrusion prevention, and intelligence from outside the firewall."[1]

Application-level inspection is one of the key technology foundations that define the NGFW product and Cloud Access Security Broker (CASB) technology. At Palo Alto Networks, we have termed this technology, App-ID™. Ever since we invented App-ID in 2007, we have led the industry in terms of both coverage and depth of application-level inspection and supported more than 3,500 different enterprise applications over the last 14 years, including the Oracle database and CRM. App-ID's highly accurate capabilities have enabled numerous customers to upgrade from traditional static port- and protocol-based policy enforcement methods to more dynamic and powerful application-based policy enforcement.

How Traditional App-ID Works

All NGFW and CASB vendors offer some form of application-level inspection today. And the common approach is to rely on application identification signatures.

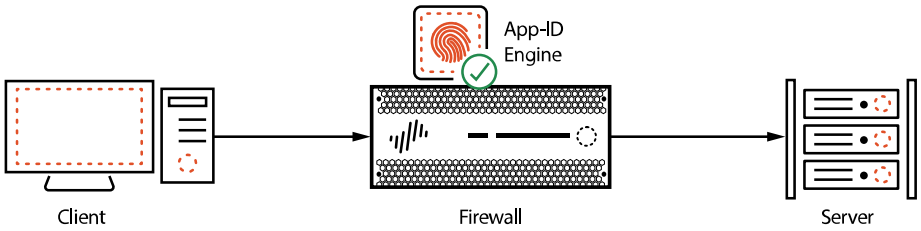

Figure 2: ML-Based App-ID Operating Model

An enterprise application exposes certain characteristics on the network, i.e., runs on a specific port, exposes specific traffic patterns or contains certain types of metadata. Security researchers and developers can observe and extract these characteristics and then come up with pattern match rules to detect representative characteristics. These pattern match rules are called signatures. NGFW and CASB administrators can then leverage these signatures to detect applications inside the enterprise and then define application-specific security policies.

Below is an example signature which matches pattern in the http url, host, and headers to identify the upload activity from the dropbox personal account:

alert tcp $EXTERNAL_NET any -> $HOME_NET $HTTP_PORTS (msg:"dropbox upload personal"; flow:to_server,established; content:"POST"; http_method; content:"reported_block_size="; fast_pattern:only; http_uri; content:"dl-web.dropbox.com|od oa|"; nocase; http_header; content:"last_active_role=personal:"; nocase; http_header; metadata:service http;sid:500001;)

Application-based policy enforcement offers a lot more value and flexibility to customers than the traditional static port or IP address-based approach. But a key foundational consideration here is that the application recognition has to be accurate, which means those signatures have to be accurate.

The Pandemic Has Accelerated Cloud Adoption in the Enterprise

The pandemic has forced the enterprise to accelerate cloud adoption. As a result, we have started seeing two new megatrends:

1. More and more cloud native business-based SaaS applications are showing up inside enterprises, such as Zoom, GitHub, and Workday. According to Blissfully, enterprises have, on average, 288 different sanctioned SaaS apps in use across their businesses, with an upward trend of 30% year over year.[2]

1. Gartner
2. Blissfully

2. More and more unsanctioned consumer-based applications are showing up inside enterprises, such as Facebook, Netflix, and PayPal. While the use of a large majority of these applications may be tolerated, many SaaS applications can introduce serious cybersecurity risks such as data loss, malware entry, and non-compliance.

Compared to traditional enterprise applications, cloud native applications and consumer applications have different characteristics:

1. **Cloud-native applications change much more frequently and dynamically:** A traditional enterprise application usually has version upgrades multiple times per year. But a cloud native application or consumer application can upgrade multiple times every day! And each new version may significantly change the application's characteristics. An App-ID signature that used to work in a previous version might become outdated within hours in the new version.
2. **Cloud-native applications are becoming more and more encrypted:** If the signature relies on certain metadata or fields inside the application, the signature becomes less effective or outdated.
3. **Cloud-native applications run on public clouds like AWS, Azure, or GCP:** This means the IP addresses of these applications are constantly changing, and any signature based on IP addresses becomes much less accurate.
4. **Cloud-native applications outnumber traditional enterprise applications:** The number of cloud native applications and consumer applications in order of magnitude is much bigger than traditional enterprise applications: About 3,500 enterprise App-IDs provide decent coverage for a traditional enterprise. However, cloud native applications and consumer applications run in the tens of thousands or even millions.
5. **Cloud-native applications and consumer applications can carry different levels of risks** based on application categories such as collaboration and productivity, marketing, ERP, etc., or based on specific app characteristics that vary from app to app, such as data privacy violations.
6. **Lack of visibility into all sanctioned and unsanctioned apps** and their risks creates a problem called Shadow IT, given that the unsanctioned apps can be accessed by employees without the explicit knowledge and approval of the IT department.

   To make matters even worse, fast-changing applications are making the traditional signature-based application recognition approach outdated and inaccurate. Undeniably, application-based policy enforcement is facing fundamental challenges in the post-pandemic era.

A Machine Learning Approach for Next-Generation App-ID

The traditional signature-based approach has fundamental challenges when it comes to addressing changing demands. Traditional firewalls and CASB products carry such limitations and therefore cannot keep up with SaaS hypergrowth, lacking the intelligence to automatically identify new apps. We need to find a new way that can adapt to such a dynamic and fast-changing "new" application environment and figure out a new method that can learn new applications on the fly and scale up automatically.

The answer is **machine learning (ML)**. As with imaging recognition, natural language processing to self-driving cars, machine learning has proved to be an effective methodology to learn about new applications in an uncertain and complex environment, and it scales up automatically.

The challenge of application recognition is similar in nature to image recognition. Dynamic and fast-changing cloud applications and consumer applications are just like new images that pop up on the internet.

The number of applications worldwide is probably much smaller than the number of images on the internet. From a scalability perspective, machine learning can handle the scale of applications. But application recognition probably has higher demands in speed and accuracy than image recognition. This is due to the fact that security policies rely on the result of application recognition, and those policies need to work in real time as traffic passes through the firewall or other security enforcement points.

Application recognition has a different data source from image recognition. Its data source comes from traffic generated by each application, while the data source for the image recognition is the image itself. Image belongs to a static data source, while application traffic is dynamic and flow-based.

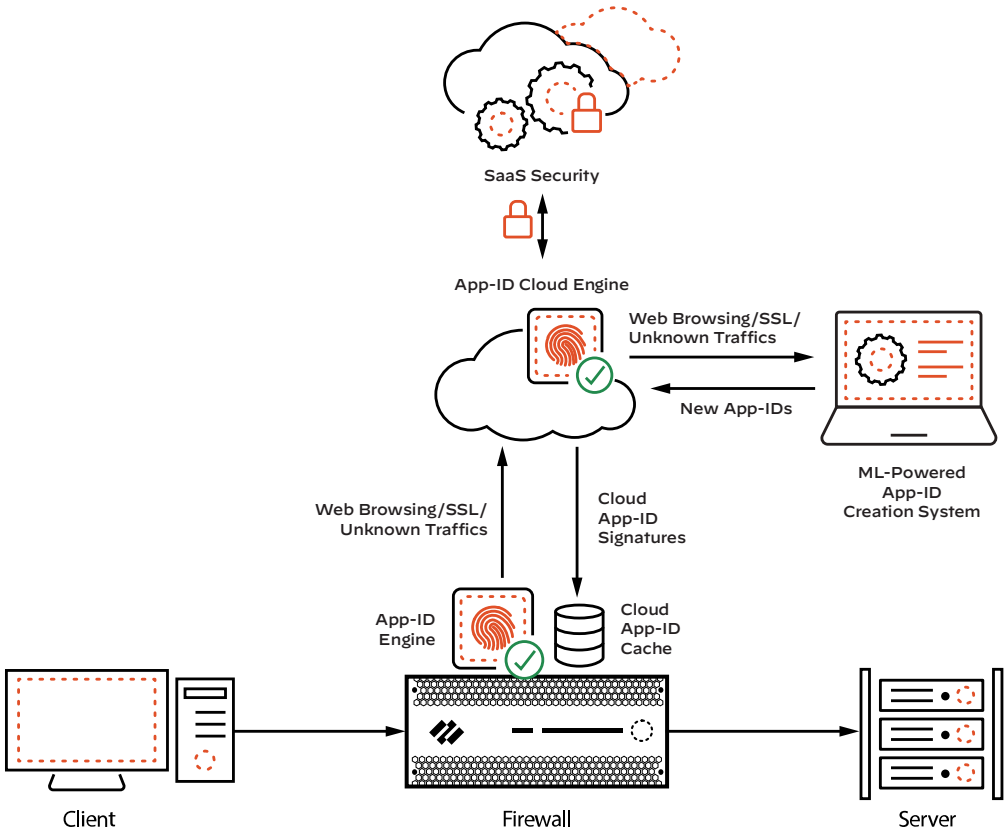

Figure 2: ML-Based App-ID Operating Model

To meet real-time requirements and achieve high accuracy, the machine learning process can start with metadata extraction—extract representative and expressive metadata from the traffic flow of each application automatically, and then feed the metadata to various machine learning techniques, such as Deep Neural Networks (DNN). One advantage of application recognition is that, just like image recognition, we have a large data set to train our machine learning algorithms. Each new application we see will make our machine learning algorithms a bit more smarter, quicker, and accurate.

Depending on the goals, machine learning can run inline (embedded inside the firewall or other security enforcement points) in the cloud or hybrid. Inline ML provides better real-time protection. Cloud-based ML will likely provide more accurate results, while hybrid can blend both real-time protection and accuracy to satisfy different needs.

Advantages of ML-Based App-ID

Compared to the traditional signature-based App-ID technologies, the ML-based App-ID approach offers several distinct advantages:

1. **ML-based App-ID technology can adapt to dynamic and fast-changing cloud applications and consumer applications and offers accurate real-time recognition.**

   ML-based App-ID can automatically detect new version upgrade changes, behavior changes and other characteristic changes on the fly and then automatically adapt to those changes and ensure the application recognition is consistently accurate.

2. **ML-based App-ID technology can scale up automatically.**

   ML relies on data; more data means better and quicker results. As the number of applications grows, ML will become more efficient and accurate.

3. **ML-based App-ID technology can provide more values and enable more use cases.**

   ML-based App-ID technology can not only accurately recognize dynamic applications, but also provide more values, and enable more use cases in the future. For instance, this approach can help the IT team analyze and understand the behavioral changes of a specific application so that it can better secure the application or offer a better user experience or performance. As another example, this approach can also help the customer better predict and plan for application onboarding and retirement, capacity planning, and troubleshooting.

How Is App-ID Related to CASB?

App-ID was traditionally designed for enterprise applications. But as more and more apps have now shifted to the cloud and are easily delivered from the cloud, the need for enterprises to create their own apps has greatly reduced. In today's times, both enterprise and SaaS apps need to be thought of as together and holistically from a business and security perspective.

Both the Next-Generation Firewall for the enterprise and our integrated CASB for SaaS apps use our patented App-ID technology to continuously identify new SaaS applications, ensuring applications are discovered automatically as they become popular. As for CASB specifically, one can say that App-ID is the foundation. This is because one key value of CASB is to accurately detect cloud applications and then secure those applications. CASB relies on App-ID to recognize those dynamic cloud applications accurately and in real time.

When to Use App-ID vs. an Integrated CASB

It all comes down to your management and the number of applications. App-ID will help you identify all apps, enterprise or SaaS, inline from your NGFW. Once you have visibility, it's about determining the count. If you have an overwhelming number of enterprise apps, then NGFW controls will suffice. But if you have an overwhelming number of SaaS apps, then while NGFW controls will allow security to be delivered, they will not scale to deal with the upcoming SaaS explosion. This is where an integrated CASB comes in for inline. The other pieces of CASB are API security and DLP.

Summary

The application-level inspection is the foundation of Next-Generation Firewalls and Cloud Access Security Broker products. Enterprise application is rapidly evolving and changing as the enterprise moves to the cloud, and the pandemic changes how enterprises consume applications.

Traditional signature-based application-level inspection technologies have fundamental challenges that don't adapt well to changes. ML-based approaches, on the other hand, can adapt to new changes inside the enterprise, offer accurate application-level inspection, and enable new use cases.

Palo Alto Networks SaaS Security is the only integrated CASB that automatically keeps pace with the SaaS explosion. Natively integrated with the Palo Alto Networks NGFWs, it uses ML-based App-ID technology to deliver continuous discovery, categorization, and control of new and emerging SaaS applications. In addition, it provides best-in-class protection and the fastest time to value for all sanctioned and unsanctioned SaaS applications.

paloalto NETWORKS

3000 Tannery Way
Santa Clara, CA 95054

Main: +1.408.753.4000
Sales: +1.866.320.4788
Support: +1.866.898.9087

www.paloaltonetworks.com