SECURING DATA IN THE PRIVATE DATA CENTER AND PUBLIC CLOUD WITH ZERO TRUST

RELEASE 1
AUGUST 2018

Table of Contents

Purpose of This Guide

This guide describes how your organization can use the Palo Alto Networks® Security Operating Platform in the design of a Zero Trust security policy to protect your sensitive data and critical applications.

This guide provides architectural guidance on Zero Trust for solution architects and engineers who are familiar with the Security Operating Platform. Use this guide as a roadmap for architectural discussions between Palo Alto Networks and your organization.

AUDIENCE

This design guide is written for technical readers, including system architects and design engineers, who want to deploy the Palo Alto Networks Security Operating Platform in support of a Zero Trust security model. It assumes the reader is familiar with the basic concepts of applications, networking, virtualization, and security, as well as a basic understanding of network, data center, and public cloud architectures.

RELATED DOCUMENTATION

The following documents support this guide:

- Palo Alto Networks Security Operating Platform Overview—Introduces the various components of the Security Operating Platform and describes the roles they can serve in various designs.
- Zero Trust Overview—Introduces the concepts of the Zero Trust security model and how to implement it using the Palo Alto Networks Security Operating Platform.
- Reference Architecture Guide for Cisco ACI—Presents a detailed discussion of the available design considerations and options for the next-generation firewall in a Cisco ACI-based private data center deployment.
- Reference Architecture Guide for VMware NSX—Presents a detailed discussion of the available design considerations and options for the next-generation firewall in a VMware NSX-based software defined data center deployment.

Introduction

Stories of breaches and data loss that expose private information are in the news almost every week. When these events happen, there can be a significant personal impact on those who have information exposed, as well as a loss of trust in the applications and companies who were breached. It doesn't matter if the loss occurred because of accidental exposure or malicious act, the impact to an organization that has a breach or data loss event is real.

These events have become so common, at a rapid pace industry is developing new standards and governments are developing new regulations that are forcing organizations to evaluate their security posture and do everything they can to prevent these events from occurring. This process can be challenging for several reasons:

- **Location of security infrastructure**—Security infrastructure requires visibility and control of the relevant activities to prevent data breaches. Many organizations deploy security as a function at the Internet perimeter. Although security at the Internet perimeter is important, sensitive data can be widely dispersed within an organization and in the public cloud and accessible to users without traversing a security device.
- **Capabilities of security infrastructure**—Even when relevant activities are visible to the security infrastructure, many organizations have a security infrastructure that cannot fully characterize relevant traffic and activities. The infrastructure must have the ability to characterize traffic and activities based on business-relevant attributes such as application (not port and protocol) and user and group (not IP address). It also must be able to identify and stop threats as they happen.
- **Security component coordination**—When threat and policy information is siloed in multiple security devices, coordinating a comprehensive security posture to protect against breach and data loss requires manual coordination and operation. This reduces the effectiveness of the security infrastructure because the threats evolve faster than security operations can manually coordinate the security infrastructure.
- **Security framework**—The biggest challenge for many organizations is defining a security model that provides the required security across the organization holistically. Most security architectures focus on the place of protection and policy that is based on what is known to be a risk.

Zero Trust is a security model developed specifically to address the security of sensitive data and critical applications in an enterprise organization. Zero Trust remedies the deficiencies of perimeter-centric strategies and the legacy devices and technologies used to implement them. Zero Trust policy leverages the capabilities of an automated security platform to define policy based on what is required, reducing the number of ways attacks can happen.

The most critical and sensitive data often resides on resources within your private data center or public cloud environment. This guide discusses obtaining visibility and control through segmentation and designing security policy for the private data center and public cloud environment using the Palo Alto Networks Security Operating Platform.

Segmenting Data and Applications

Securing applications and services in the Zero Trust model depends upon the ability of the next-generation firewall to have visibility and control of the inbound traffic to the application, outbound traffic from the application, and traffic between applications components. To provide the required visibility and control, you should segment data and applications in the private data center and public cloud provider behind a next-generation firewall.

The data-sensitivity level of the application does not define if it requires protection by a firewall. Instead, it informs you on how to group applications and services with common security and traffic flow requirements. When the data sensitivity increases, additional policies and protections are necessary, including a stricter definition of what is permitted to access the application. Although applications of low data-sensitivity can be grouped behind the firewall, the increased policy detail required for moderate and high levels of data sensitivity means fewer applications can be grouped even when they are at the same sensitivity level. In fact, an application or service that is at the highest level of sensitivity should not be grouped with any other application. High-sensitivity services should even be separated from other components of their application if those other components have a reduced security requirement. The sensitivity levels are:

- **Low**—Applications and information whose loss of availability would have limited impact on the organization or its customers.
- **Moderate**— Infrastructure, applications, and systems whose loss of integrity and availability would impact the organization or its customers.
- **High**— Any information falling under statutory requirements for notification in the case of a breach.

How you create the network segments for an application depends upon the infrastructure on which it is built. In traditional network infrastructure, the Layer 2 and Layer 3 forwarding tables define which endpoints can communicate, and a next-generation firewall at the Layer 2 and Layer 3 boundaries is used to control and limit the traffic. The simplest form of a network segment is a Layer 2 VLAN that doesn't have infrastructure Layer 3 services. Instead, the firewall provides Layer 3 services to the servers and services through an interface (or subinterface) in the VLAN. This traditional configuration works in almost any private data center infrastructure. However, this type of deployment can be challenging to scale.

Figure 1 Layer 2 VLAN with next-generation firewall default router

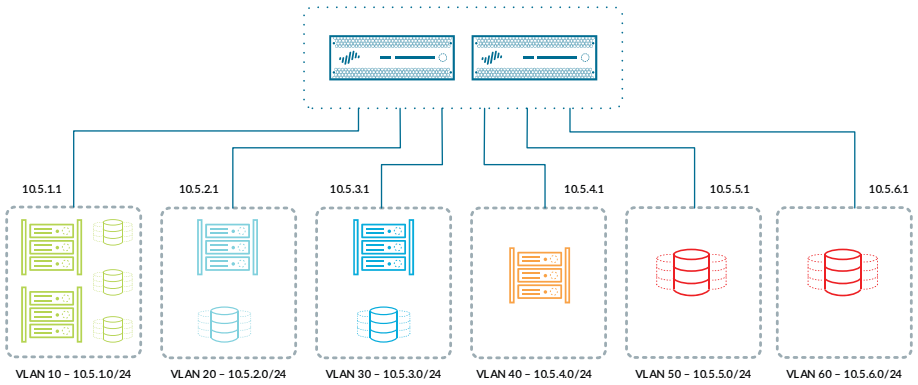

To reduce the number of VLANs and firewall interfaces required in the data center as the number of servers and services scale instead of using VLANs to create network segments, you can use virtual routing and forwarding (VRF) instances to aggregate multiple networks while keeping them separate from others. When using a VRF to create a network segment, the next-generation firewall integrates through Layer 3 static or dynamic routing. It is important to remember, though, that everything within a segment should have common security and traffic flow requirements. VRFs allow the scale of the number of devices in a segment but do not ease the challenges associated with scaling the number of segments within the data center.

Figure 2 Layer 3 VRF with next-generation firewall integrated through OSPF

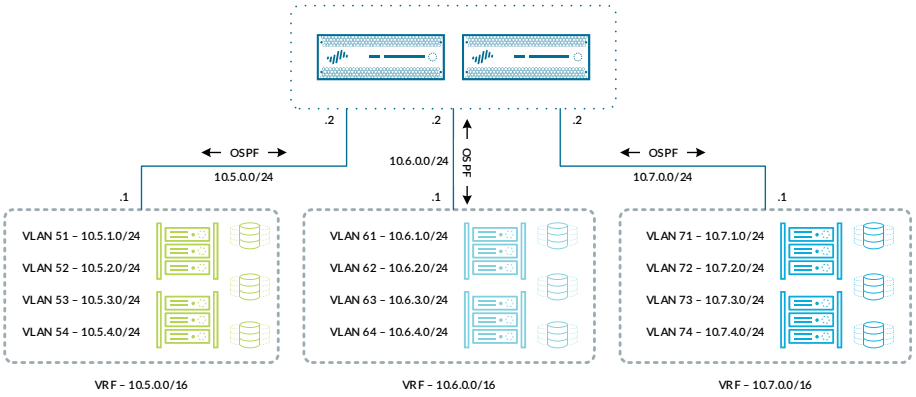

Because fine-grained segmentation is difficult to achieve at scale with traditional infrastructure tools, many organizations that have this requirement are transitioning to software-defined data center (SDDC) infrastructure. SDDC infrastructure gives organizations that ability to define the segments not just through attachment to a specific VLAN or VRF but also through policy that can group servers and services dynamically based on policy and forward the traffic as appropriate to the next-generation firewall. Examples of software-defined data center infrastructure include:

- **Cisco ACI**—Application profiles and endpoint groups define the groups of servers and services that have common security requirements. Contracts, filters, and policy-based redirect ensure that the infrastructure forwards traffic between the groups to the firewall for security policy.
- **VMware NSX**—Security tags dynamically associate virtual machines with a security policy through security group membership. Service profiles are used in NSX security policies to provide network introspection services in the next-generation firewall.

Software-defined data center infrastructure also makes it easier to use different firewalls for different traffic flows. The ability to use different firewalls for north-south traffic versus east-west traffic, for example, makes it easier to scale out the security infrastructure as the amount of traffic grows.

Little changes when the applications and services do not live in a private data center but instead live in a public cloud provider's infrastructure. All applications and services should be segmented behind a next-generation firewall. One additional consideration in the public cloud is that there might be different paths into the infrastructure for traffic sourced from the Internet versus traffic sources from within the organization. If there are multiple paths into the infrastructure, take care to ensure all paths must traverse a next-generation firewall before reaching the application or service.

Like with the private data center, how you accomplish the segmentation is dependent on the public cloud provider. For example:

- **Amazon Web Services**—Virtual Private Cloud instances define the groups of servers and services. Route tables direct traffic to the VM-Series firewall for security policy.
- **Microsoft Azure**—VNets and subnets define the groups of servers and services. User-defined routes direct traffic to the VM-Series firewall for security policy.

Although public cloud providers use the term *route table* when describing how to define the traffic steering policies, public cloud environments more closely resemble software-defined data center infrastructure. For example, in Microsoft Azure, the user-defined routes can direct traffic leaving a subnet directly to the firewall. The firewall does not have to have an IP address or interface attached to the subnet. This closely resembles the way Cisco ACI works when using policy-based redirect.

Figure 3 Microsoft Azure subnets and user-defined routes

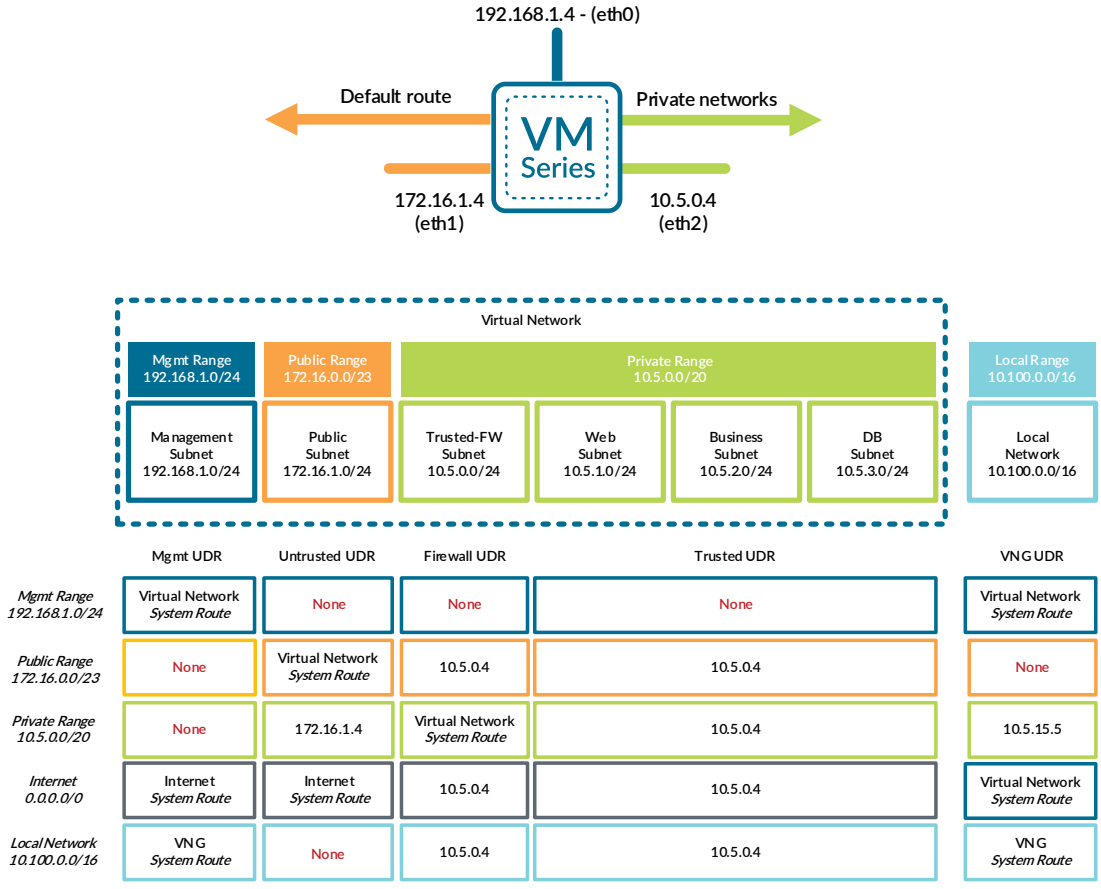

Security Policy

The concept of least privileged access forms the basis of a Zero Trust security policy. You can define least-privileged access by using the following set of the questions:

- **Where**—From where is the access sourced?
- **Who**—Who are the individuals or groups who must have access?
- **What**—What data and services are being accessed?
- **How**—How are the users accessing the data and services?

DEFINING SOURCE AND DESTINATION WITH ZONES AND DYNAMIC ADDRESS GROUPS

Primarily, the next-generation firewall applies Zero Trust security policy to network segments through zones. A *zone* is a group of one or more firewall interfaces that connect to the network segments that separate the different applications, services, and private and public connectivity. Zones do not contain security policy. Instead, Zero Trust security policies use zones to define at a basic level which network segments can communicate. Zones are used instead of referencing interfaces directly because they give a level of abstraction that allows you to define common policy in Panorama™ that applies to multiple firewalls, independent of how the firewalls are physically connected to the infrastructure. In the private data center and public cloud environments, this abstraction is especially important when there is a mix of physical and VM-Series next-generation firewalls.

Every Zero Trust policy should include zone information to identify where traffic should be coming from and going to. Zone names are arbitrary, with conventions of “trust” and “untrusted” commonly used in examples to describe connections to private and public network segments on perimeter firewalls. In the Zero Trust model, it is best to name the zones descriptively based on the connected functionality instead of a label denoting policy. As the number of network segments and zones increases, this naming convention makes policy easier to maintain.

Dynamic Address Groups

In some environments, especially software-defined data centers and public cloud providers where infrastructure policy defines network segments, not VLANs and VRFs, multiple private segments might connect to the firewall on a single interface and zone. This is possible because the infrastructure policy ensures the two segments can’t communicate without going through the firewall, even though they share a common IP network. In these environments, the firewall can use dynamic address groups to define and separate the private segments so that separate policy can still be applied even though they share a zone.

Dynamic address groups allow you to create an object on the firewall that contains the devices known to be in a network segment and automatically adapts to adds, moves, or deletions of devices. Dynamic address groups use tags to determine which IP addresses are members of the group. In most infrastructures, the firewall polls the software-defined data center or public cloud provider for application server IP addresses and their associated tags through the VM Monitoring Agent. In some instances, like Cisco ACI and Microsoft Azure, IP addresses and tags are dynamically registered on the firewall through the XML API.

Because the members of a dynamic address group are automatically updated when application administrators make changes, you can use address groups to adapt to those changes without relying on the firewall administrator to make policy changes and commit them.

Figure 4 Dynamic address groups in the public cloud and software defined data center

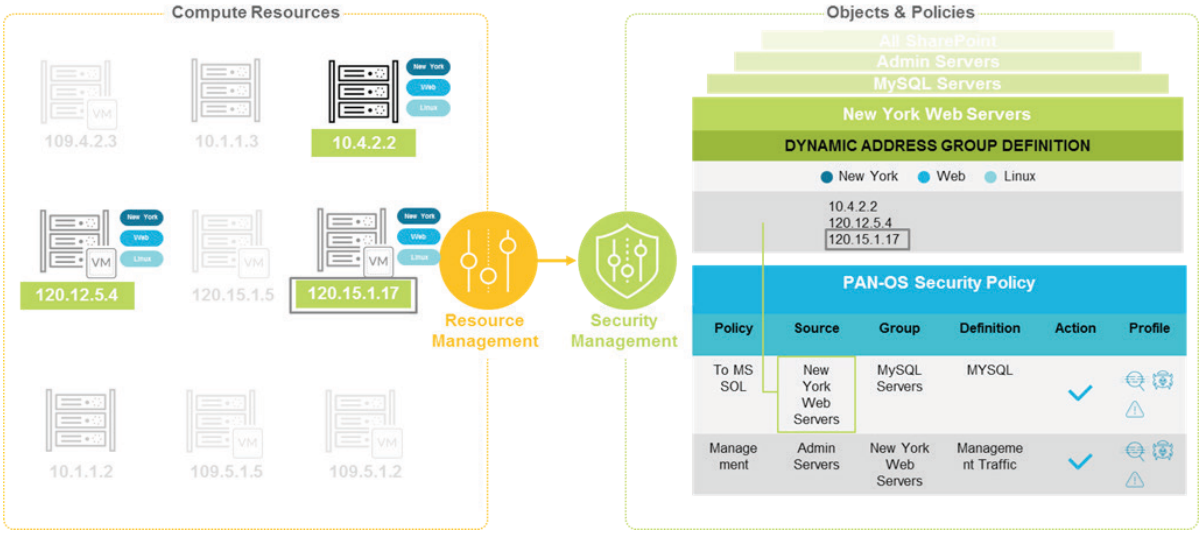

Using Zones and Dynamic Address Groups in Policy

In the private data center, the number of network segments you define to group your applications and services drives the number of zones and dynamic address groups on the firewall. From a Zero Trust security policy perspective, zones are the most basic form of check on if access is allowed and the first building block of a least privileged access policy. When designing a Zero Trust security policy to protect an application, for each of the other zones on the firewall consider if any device in those network segments should be able to initiate communication with the application and vice versa. As the level of sensitivity of a zone increases, the number of zones that can communicate with it should decrease. For example:

- **Low**—All private zones
- **Moderate**—Private zones that have regular employees
- **High**—Only zones that contain front-end application services

Table 1 Using zones and dynamic address groups in the security policy

	Source zone	Source address	Destination zone	Destination address	Action
Services inbound from internal	WAN Application A Application B Application C Application D	Organization IPs	Services	DNS DAG DHCP DAG NTP DAG	Permit
Services outbound to internet	Services	DNS DAG NTP DAG	Internet	Any	Permit
Low inbound from internal	WAN	Organization IPs	Application A	Any	Permit
Low outbound to internet	Application A	Any	Internet	Any	Permit
Moderate inbound from internal contractors	WAN	Contractor IPs IoT IPs	Application B	Any	Deny
Moderate inbound from internal	WAN	Organization IPs	Application B	Any	Permit
Moderate outbound to internet	Application B	Any	Internet	Any	Permit
High inbound from moderate	Application B	Any	Application C	Any	Permit
High inbound from moderate	Application B	Any	Application D	Any	Permit
Interzone-Default	—	—	—	—	Deny

MAPPING IP ADDRESSES TO USERS WITH USER-ID

Although zones and dynamic address groups help define the source and destination of a traffic flow, alone they don't have enough detail to properly define a least-privileged access policy. Designing a Zero Trust security policy to protect applications and services relies heavily on being able to identify who should be able to use that application or service. Zones and IP addresses help inform the security policy about the location of a user, but with the mobility of users, they are ineffective identifiers of the user's group membership. To provide context to IP addresses, User-ID™ on the next-generation firewall dynamically maps IP addresses to users and their associated group membership and allows you to define security and authentication policy based on specific users or groups.

User-ID integrates with a wide range of user repositories and terminal service environments. Depending on your environment, you can configure multiple techniques for user and group mapping. However, not all integrations provide the same level of dependability and timeliness. When users move around the network or go remote as they leave the net-

work, it is important that the firewall has the most accurate representation of the IP address to user mapping. In fact, a Zero Trust security policy works the best when the IP-to-user mapping is in place before any inbound traffic from the users reaches the firewall.

Local Users in the Campus and Remote Site

For local users, User-ID integration into WLAN and wired authentication and authorization servers provides high fidelity IP-to-user mappings. When endpoints move around inside the network, the infrastructure authenticates each move before allowing access to the network. Authentication is also independent of the endpoint operating system. This gives the AAA servers very timely and accurate IP address and user information. The integration of User-ID and the AAA servers vary based on the vendor. Some integrations, like Aruba ClearPass, rely on the AAA server pushing IP and user information to the firewall through the firewall's User-ID API. Other integrations, like with Cisco ISE, rely on the AAA server sending RADIUS account information to the firewall though SYSLOG, which User-ID then parses for username and address information.

Figure 5 Infrastructure User-ID source

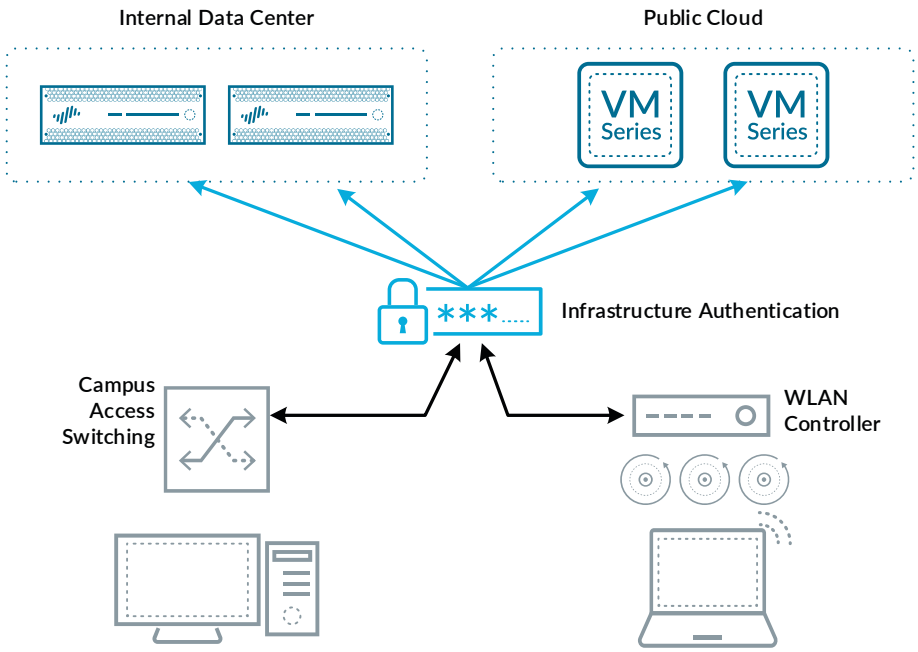

However, although authentication of WLAN is very common, authentication of wired connections has less adoption. In environments without authentication of wired users, GlobalProtect™ internal gateways are a high-fidelity source of User-ID information for endpoints on the wired networks. Although GlobalProtect is primarily used to attach remote users to a next-generation firewall so they have the same level of security as on-site users, GlobalProtect internal gateways are primarily used to identify internal users without terminating user traffic.

Figure 6 AAA and GlobalProtect internal gateway

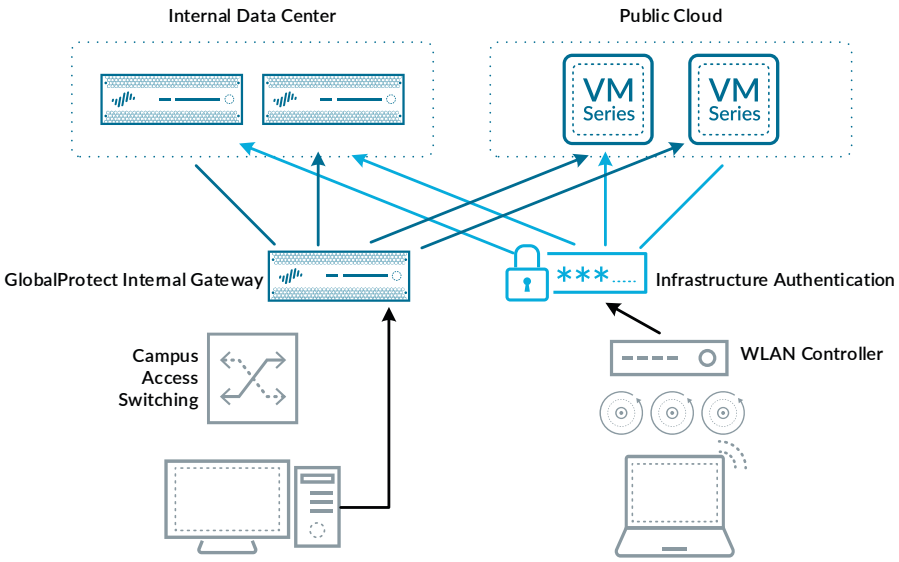

If AAA server or GlobalProtect isn't a viable User-ID source for an organization, another (but lower fidelity) choice is Active Directory domain controller monitoring. User-ID monitors domain log-on events and uses them to associate IP addresses with users. Active Directory monitoring is a lower fidelity User-ID source for two reasons. First, Active Directory log-on events do not happen every time an endpoint changes its IP address, and with a default configuration, the endpoints can go up to ten hours before renewing their login with Active Directory. Second, endpoints must be joined to the Active Directory domain, which can be challenging when there are MacOS, iOS, and Android endpoints in the organization. Even though Active Directory is a very common integration that allows User-ID to bring visibility of the user into traffic logs, because of these challenges, in a Zero Trust security policy, higher-fidelity sources of User-ID are strongly recommended.

Figure 7 Active Directory User-ID source

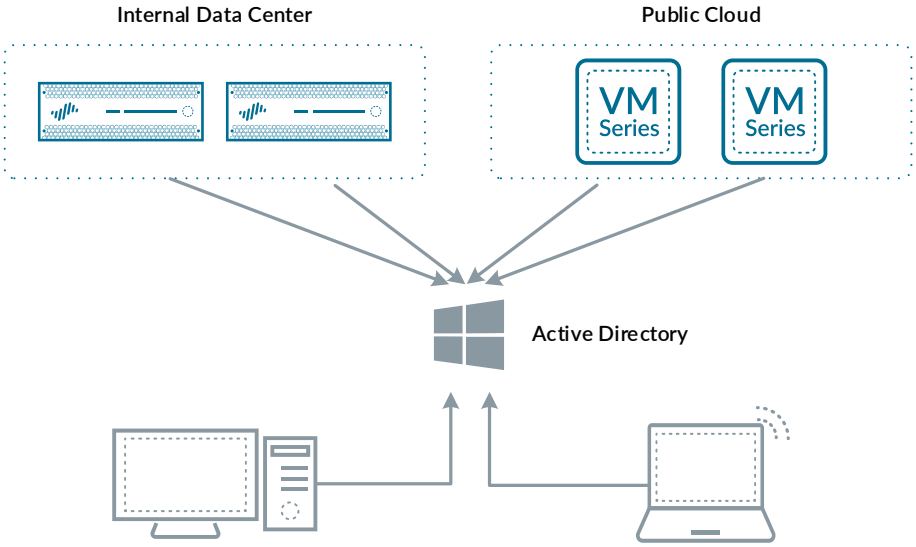

Remote Users

Because remote users must connect to GlobalProtect in order to access private applications and services from the Internet, GlobalProtect (either the external gateway and cloud service for mobile users) provides the best source of IP-to-user mapping. Remote users provide login credentials to GlobalProtect when they connect, providing a high-fidelity IP-to-user mapping. User-ID can redistribute the IP-to-user mappings from the firewall running GlobalProtect to the firewalls in the private data center and public cloud.

Considerations for User-ID in a Public Cloud Environment

User-ID is just as important for a Zero Trust security policy in a public cloud environment as in the private data center. However, there are a few design considerations.

First, configure a Zero Trust security policy that includes User-ID for connections originating within your organization and entering your public cloud environment through your private connection (Azure Express Route, AWS Direct Connect, etc). Deriving IP-to-user mappings for connections entering from the Internet can be extremely challenging.

Second, network address translation can have a detrimental effect on being able to map IP addresses to users. Although User-ID supports mapping multiple IP addresses to a single user, it cannot map more than one user to an IP address. So, if traffic from multiple users entering the firewall in a cloud environment has been translated to a single address (for example by a load balancer), then the next-generation firewall can't distinguish between the users. To combat this, many load balancers support inserting X-Forwarded-For (XFF) fields in HTTP headers. XFF allows the firewall to identify the originating IP address of a client connecting to a web application. If you are hosting non-HTTP applications in the public cloud environment, consider resiliency designs that do not translate the source IP address of private traffic destined to the firewall.

Validating User-ID with Authentication Policies

Authentication policy enables you to authenticate users before evaluating the security policy to determine if they are allowed access. Authentication policies are useful tools in a couple of different scenarios:

- **Authenticating unknown users**—Invariably a user will eventually need to reach an application protected with a Zero Trust security policy and be unable to because of a missing IP-to-user mapping. Authentication policies can be used on traffic of "unknown users" from the public zone to an application zone. The authentication policy can use several authentication sources, including SAML single sign-on. Once authenticated through the captive portal, the firewall updates the IP-to-user mapping.
- **Validating User-ID before allowing access to sensitive applications**—For your most sensitive applications and services, you should use authentication policies to validate that the expected user is initiating the traffic to the application. Authentication policies that match the traffic from specific users to the application zone can force a multi-factor authentication of the user before allowing access and ensure that the originator isn't using stolen credentials. The MFA is done transparently to the application and is especially useful for administrative access that doesn't natively support or is difficult to configure for MFA.

Using User-ID in Policy

Although IP-to-user mapping is critical to being able to implement a Zero Trust security policy, in most instances the policy will not be described by specific users. Instead, for ease of management, a Zero Trust security policy typically uses group information to describe who can access an application or service. To determine group membership, regardless of the source of the IP-to-user mapping, after the firewall gathers user information, User-ID uses LDAP to integrate with the directory service in order to obtain group information for that user.

Authentication policies and User-ID provide important detail to the Zero Trust security policy. To protect applications and services in the private data center and public cloud, security rules for traffic inbound from the internal network should be limited to the appropriate user groups. When the level of sensitivity for an application increases, fewer users should have access to the application. For the most sensitive applications, enforce multi-factor authentication policies before the security policy evaluation of user information.

Figure 8 MFA authentication for high-sensitivity segments

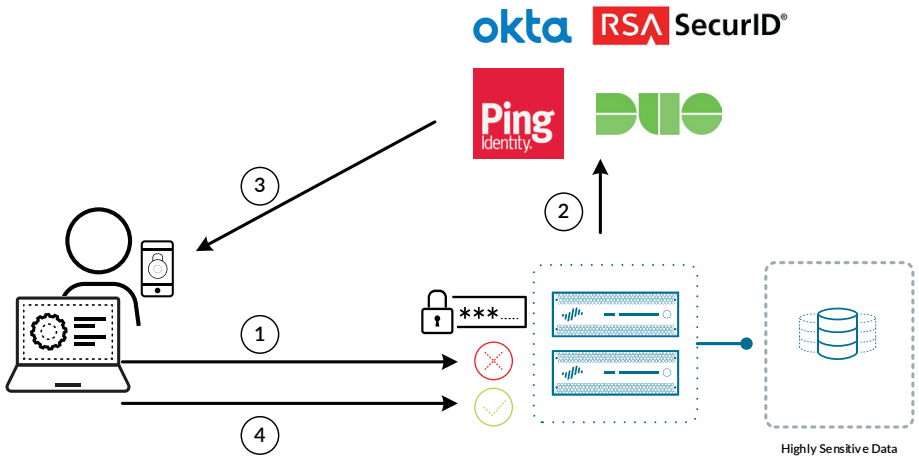

In most cases, the Zero Trust security policy for outbound traffic to the Internet initiated from the application servers should not evaluate User-ID. Servers typically do not have active users and instead should be identified through dynamic address groups.

Table 2 Authentication policy

	Source zone	Source user	Destination zone	Action
Force authentication is unknown user attempts traffic to web application	WAN	unknown	Application B	Captive Portal
Force MFA for administrators that want access to application	WAN	group.admins	Application C	MFA

Table 3 Security policy with User-ID

	Source zone	Source user	Source address	Destination zone	Destination address	Action
No user required for services	WAN Application A Application B Application C Application D	Any	Organization IPs	Services	DNS DAG DHCP DAG NTP DAG	Permit
No user required for traffic from servers	Services	Any	DNS DAG NTP DAG	Internet	Any	Permit
All known users	WAN	known-user	Organization IPs	Application A	Any	Permit
No user used to block	WAN	Any	Contractor IPs IoT IPs	Application B	Any	Deny
Permit Employee group	WAN	group.em-ployees	Organization IPs	Application B	Any	Permit
Permit admin group	WAN	group. admins	Admins IPs	Application C	Any	Permit
Interzone-Default						Deny

DEFINING APPLICATIONS WITH APP-ID

For a Zero Trust security policy to be able to protect applications and services, those applications and services must be described in the policy. Zones, dynamic address groups, and User-IDs all help apply policy based on a network segment or IP address, but they do not describe the application traffic itself, just where it is coming from, where it is going to, and who initiated it. Traditionally, the security policy describes application traffic as a set of TCP or UDP ports numbers. However, this is a coarser descriptor, especially with the movement to HTTP and HTTPS based applications that all would share the same TCP port numbers. App-IDs bring the visibility and control of the applications traversing the firewall that is required to define Zero Trust security policies.

The App-ID™ traffic classification system in the next-generation firewall identifies the applications (and their functions) when the traffic traverses the firewall. App-IDs identify applications through several techniques—regardless of the traffic port or protocol—even when it is encrypted, tunneled, or uses evasive tactics.

App-IDs identify applications traversing the next-generation-firewall, including both client to server applications as well as intra-component applications. Palo Alto Networks dynamically updates the list of App-IDs, based upon input from partners and customers, as well as market trends.

When an application has multiple functions that might need to be treated differently in policy (for example, editing, uploading and downloading), the functionality may be covered by multiple App-IDs. There are three types of App-IDs:

- **Container App-ID**—Groups multiple App-IDs related to a single application (Example in Figure 8: boxnet)
- **Base App-ID**—Identifies the core functionality of the application (Example in Figure 8: boxnet-base)
- **Functional App ID**—Identify specified functionality of the application (Examples in Figure 8: boxnet-sharing, boxnet-editing, and so on)

For most applications with multiple App-IDs, the base App-ID matches the vast majority of traffic for that application and the functional App-IDs depend upon it.

Figure 9 Container App-ID, base App-ID, and functional App-IDs

NAME	CATEGORY	SUBCATEGORY	RISK	TECHNOLOGY
boxnet				
└ boxnet-sharing	general-internet	file-sharing	2	browser-based
└ boxnet-base	general-internet	file-sharing	3	browser-based
└ boxnet-editing	general-internet	file-sharing	3	browser-based
└ boxnet-consumer-access	general-internet	file-sharing	3	browser-based
└ boxnet-enterprise-access	general-internet	file-sharing	3	browser-based
└ boxnet-downloading	general-internet	file-sharing	3	browser-based
└ boxnet-uploading	general-internet	file-sharing	4	browser-based

App-IDs are as unique as the applications they describe. Some applications have App-IDs that differentiate between consumer and enterprise-level logins to an application. Other applications have App-IDs that provide visibility when a user renames a folder or uses screen-sharing.

App-ID functionality is always enabled in the next-generation-firewall and is classifying all traffic, not just HTTP/S traffic, all of the time. Because App-IDs look at all the traffic passing through the next-generation firewall (business applications, consumer applications, and network protocols), you do not need to configure an App-ID to look for specific traffic. Instead, App-ID is used in the Zero Trust security policy to further define the least-privileged access by including application identification in with defining policy based on where the traffic is coming from (zones, dynamic address groups), as well as who initiated it (User-ID).

Using App-ID in policy

Use App-ID to control the traffic inbound traffic into the segment. For segments that contain multiple low sensitivity applications, use dynamic address groups to microsegment the application resources. So, instead of allowing DNS and NTP into all servers in a shared services segment, limit DNS traffic to the DNS address group in the shared services segment and NTP to the NTP address group.

Table 4 Security Policy with App-ID

	Source zone	Source user	Source address	Destination zone	Destination address	App-ID
Anyone can access infrastructure services	WAN Application A Application B Application C Application D	Any	Organization IPs	Services	DNS DAG DHCP DAG NTP DAG	DNS DHCP NTP
Services can access the internet for lookups	Services	Any	DNS DAG NTP DAG	Internet	Any	dns ntp
Internal users can access internal web servers	WAN	known-user	Organization IPs	Application A	Any	web-browsing
Web servers can do software updates	Application A	Any	Any	Internet	Any	apt-get
Block contractors from exchange servers	WAN	Any	Contractor IPs IoT IPs	Application B	Any	Any
Employees can contact exchange servers	WAN	group.em-ployees	Organization IPs	Application B	Any	mapi-over-http ms-ex-change rpc-over-http activesync
Exchange servers can contact SQL servers	Application B	Any	Any	Application C	Any	mssql-db

About Decryption

App-ID uses traffic payload information to identify applications. Because many applications encrypt the traffic between client and server, to have granular visibility and control to the application, the next-generation firewall must decrypt the traffic. The next-generation firewall can decrypt and inspect both inbound and outbound SSL/TLS connections traversing the firewall.

Inbound Decryption

Use SSL Inbound Inspection to decrypt and inspect inbound SSL traffic from internal endpoints to private servers in the data center and public cloud. SSL Inbound Inspection requires that you import the server certificate and key on to the firewall. Because it has the server certificate and key, the firewall can access the SSL session between the server and the client and decrypt and inspect traffic transparently. After traffic is decrypted, the firewall can apply security policies to the decrypted traffic, detecting malicious content and controlling applications running over this secure channel.

Outbound Decryption

With SSL Forward Proxy, the firewall can decrypt outbound traffic from servers in the data center and public cloud encrypted using SSL/TLS.

To secure the connection, SSL uses certificates to establish trust between the client and server. Certificates that are signed by a trusted certificate authority are installed on the firewall and used to establish the firewall as a trusted third party to your internal server during the connection setup. Your PKI infrastructure must trust the firewall's SSL certificate, or you must deploy the SSL certificate on each server in your data center that will be participating in sessions decrypted by the firewall.

The next-generation firewall preserves the integrity of the SSL/TLS session by using the cryptographic settings of the original SSL/TLS negotiation as mandated by the client and the server. It does not change the cryptographic parameters after the session has been negotiated. Further, to reduce risks associated with older versions of the protocols, PAN-OS® allows you to specify the supported SSL/TLS protocol versions and cipher suites. CRL/OCSP checks ensure that certificates presented during SSL decryption are valid.

Figure 10 SSL forward proxy

Using Decryption in Policy

Decryption is policy-based. Such policies can be used to decrypt, inspect, and control both inbound and outbound SSL and SSH connections. Decryption policies allow you to specify traffic for decryption according to destination, source, or URL category. Because the App-ID of the session is determined in part by the payload of the packets, it cannot be used to define decryption policies.

The best-practice security policy dictates that you decrypt all traffic except sensitive categories, such as Health, Finance, and Government. However, that practice is not always possible or desired by organizations, even when they lose fine-grained application granularity and threat and file blocking. When you can't decrypt everything, but still want those capabilities for a SaaS application, you can decrypt the application by integrating the IP and domain information that many SaaS vendors publish into the decryption policy dynamically or through static definitions.

The ability to decrypt or not decrypt depends not just on the firewall decryption policy but also on the design of the client software. Even when the operating system and web browser trust the certificate of the next-generation firewall, instances may exist where the client application software does not. Some client applications have a copy of the server certificate information built-in, which is called a *pinned certificate*. When the firewall attempts to negotiate the SSL/TLS session and sends its certificate to the client, the client compares the stored information with the information in the certificate. Because the information doesn't match, the client does not complete the connection. SSL/TLS sessions for client applications that use certificate pinning cannot be decrypted by the next-generation firewall or any other decryption vendor.

The client applications that use certificate pinning vary. Some SaaS applications ensure all standalone clients (those that are outside of the web browser) use pinned certificates. Other SaaS applications may have the macOS client pinned, but not the Windows client. In a few instances, it may even depend on the hardware on which the client is running.

Because clients with pinned certificates won't connect when the firewall attempts to decrypt their traffic, it is best to avoid such clients. In most cases, accessing the SaaS application through a web-browser instead of the standalone client allows the use of the application and provides you with granular application control, as well as threat and file blocking. If you must use a client application that has a pinned certificate, you must exclude the entire application from decryption. For such an application, you lose the granular App-IDs and content security for all clients, even web-browser clients. You can use server certificate information and URLs to exclude servers from SSL decryption.

INSPECTION AND ANALYTICS

The Zero Trust model prevents threats in traffic allowed by the least privileged access policy through inspection and analytics. Inspection is used to identify malware, vulnerabilities, data exfiltration, and threats previously identified by the security operating platform. Analytics are used to identify previously unknown threats to the organization.

Inspection

Antivirus Security Profiles

In the next-generation firewall, preventing known malware is the function of security profiles, specifically the antivirus profile. Antivirus profiles protect against downloading files that contain viruses, worms, trojans, and spyware. Using a stream-based malware prevention engine, the antivirus solution provides protection without significantly impacting the performance of the firewall. This profile scans for a wide variety of malware in executables, Office and PDF files, HTML, and JavaScript, including support for scanning inside compressed files and data encoding schemes.

The two sources of malware signature information for the next-generation firewall are antivirus and WildFire® content updates. Antivirus content updates are released by Palo Alto Networks daily. However, malware can propagate globally in a much shorter timeframe. WildFire reduces the window of vulnerability through content updates that are released every five minutes. Firewalls with an active WildFire license can retrieve the latest WildFire signatures. If you do not have a WildFire subscription, signatures are made available within 24-48 hours as part of the antivirus update for firewalls with an active Threat Prevention license.

The antivirus profile can protect FTP, HTTP, IMAP, POP3, SMB, and SMTP traffic. In an antivirus profile, you can configure the action for each protocol and specify how the firewall responds to a threat event. In a Zero Trust security policy, you should configure all Inbound, Outbound and East-West rules (regardless of sensitivity level) so that the traffic is scanned for known antivirus signatures. The best practice antivirus profile uses the default action when it detects traffic that matches either an antivirus signature or a WildFire signature. The default action differs for each signature and follows the most up-to-date recommendation from Palo Alto Networks for how to best prevent malware from propagating.

Figure 11 Best practice antivirus policy

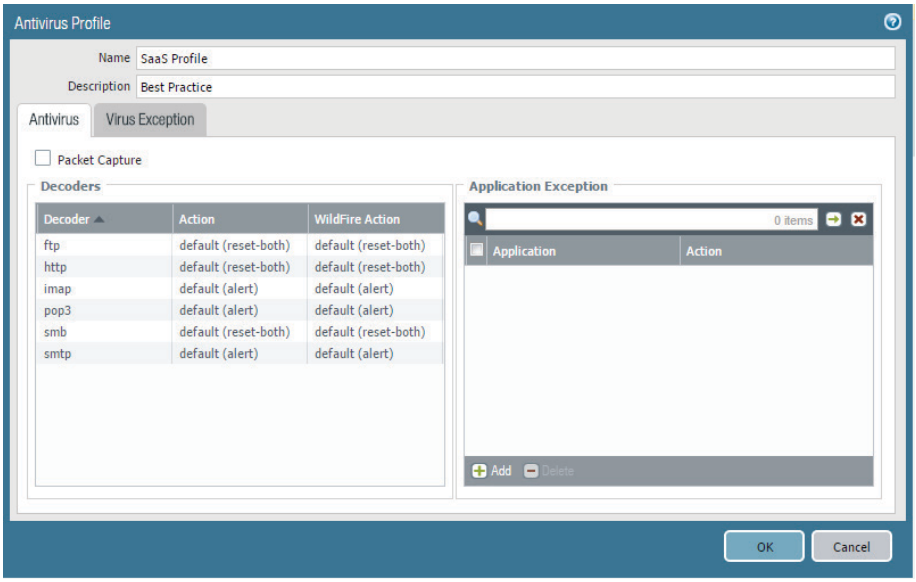

Anti-spyware Security Profiles

Anti-spyware profiles block spyware on compromised servers from contacting external command-and-control systems. To prevent infected servers from sending malicious traffic to command and control systems anti-spyware, you should enable profiles on all rules allowing outbound traffic to the Internet. Because traffic from internal servers destined to the Internet should be well known, Anti-spyware profiles on traffic inbound to and outbound from applications servers should be as strict as possible.

In addition, enable the DNS Sinkhole action on the profile so that the firewall generates responses to DNS queries for known malicious domains. DNS sinkholing causes the malicious domain name to resolve to an IP address that you define in policy and helps to identify infected hosts because any host that attempts to connect to the sinkhole IP address is most likely infected with malware.

Vulnerability Protection Security Profiles

The next-generation firewall's vulnerability protection and intrusion-prevention capabilities detect and block exploit attempts and evasive techniques at both the network and application layers, including port scans, buffer overflows, remote code execution, protocol fragmentation, and obfuscation.

In a Zero Trust security policy, configure all Inbound, Outbound and East-West rules (regardless of sensitivity level) so that the firewall resets critical and high severity threat. Leave low and informational threats to the default action.

URL Filtering

Most attacks and exposure to malicious content occurs during the normal course of web-browsing activities, which requires the ability to allow safe, secure web access for all users. URL filtering with PAN-DB automatically prevents attacks that leverage the web as an attack vector, including phishing links in emails, phishing sites, HTTP-based command and control, malicious sites, and pages that carry exploit kits.

Although the web-browsing App-ID should rarely be enabled for outbound connections from servers destined to the Internet, enable URL filtering on all rules allowing traffic to the Internet from within the data center or public cloud. At a minimum the following categories should be blocked in the URL filtering profile:

- Malware
- Phishing
- Hacking
- Proxy avoidance and anonymizers

Because traffic to the Internet from servers should be well known, you should block as many additional URL categories as possible. Limiting URLs to a specific whitelist should also be considered for moderately sensitive servers.

File Blocking

File-blocking profiles allow you to identify specific file types that you want to want to block or monitor. For inbound and east-west traffic, block files that are known to carry threats or for which the application has no use. At a minimum, these include batch files, DLLs, Java class files, help files, Windows shortcuts (.lnk), and torrent files. Additionally, to stop the exfiltration of data, outbound rules should only allow file types that are known to be required by the application.

Analytics

WildFire identifies previously unknown malware and generates signatures that the next-generation firewall, Aperture™, and Traps™ use to detect and block malware.

When a Palo Alto Networks firewall detects an unknown sample (a file or a link included in a POP or SMTP email), the firewall can automatically forward the sample to WildFire for analysis. Deciding which samples to forward to WildFire for analysis is the function of the WildFire Analysis Security Profile. Like all security profiles, the WildFire analysis profile is applied to a rule in the security policy. Each rule can have a unique profile.

For outbound traffic to the Internet, use the default WildFire Analysis profile to define the traffic that the firewall should forward for WildFire analysis. The default WildFire Analysis profile ensures complete WildFire coverage for all traffic that the security policy rule allows. It specifies that all supported file types across all applications are forwarded for WildFire analysis, regardless of whether the files are uploaded or downloaded. The file types that the next-generation firewall supports for forwarding to WildFire include Android APK, Adobe Flash, Java JAR, Microsoft Office (such as DOCX, PPTX, and XLXS), portable executable (PE), and Adobe PDF, as well as macOSX Mach-O, DMG, PKG, and application bundles.

Before forwarding the sample to WildFire, the firewall analyzes the structure and content of the sample and generates a hash. Because WildFire analyzes content from thousands of customers globally and through multiple places in the organization, the next-generation firewall then checks to determine if WildFire has already seen the hash. If it has, the firewall applies the default action defined for the signature (allow, alert, or block) and does not upload the file for analysis. If the sample remains unknown after comparing it against known WildFire signatures, the firewall forwards it to WildFire for analysis.

WildFire renders a verdict on the sample based on the properties, behaviors, and activities the sample displays when analyzed and executed in WildFire. The verdict identifies a sample as malicious, grayware, phishing, or benign. WildFire then generates signatures to recognize the newly-discovered malware, and it makes the latest signatures globally available through WildFire content updates every five minutes and the antivirus content updates within 24 to 48 hours.

Samples that firewalls submit for WildFire analysis are displayed, along with WildFire verdict, in the WildFire Submissions log. For each WildFire entry, you can open an expanded log view, which displays log details and the WildFire analysis report for the sample. For all samples, the WildFire analysis report displays file and session details. For malware samples, the WildFire analysis report also includes details on the file attributes and behavior that indicated the file was malicious.

Figure 12 WildFire analysis report

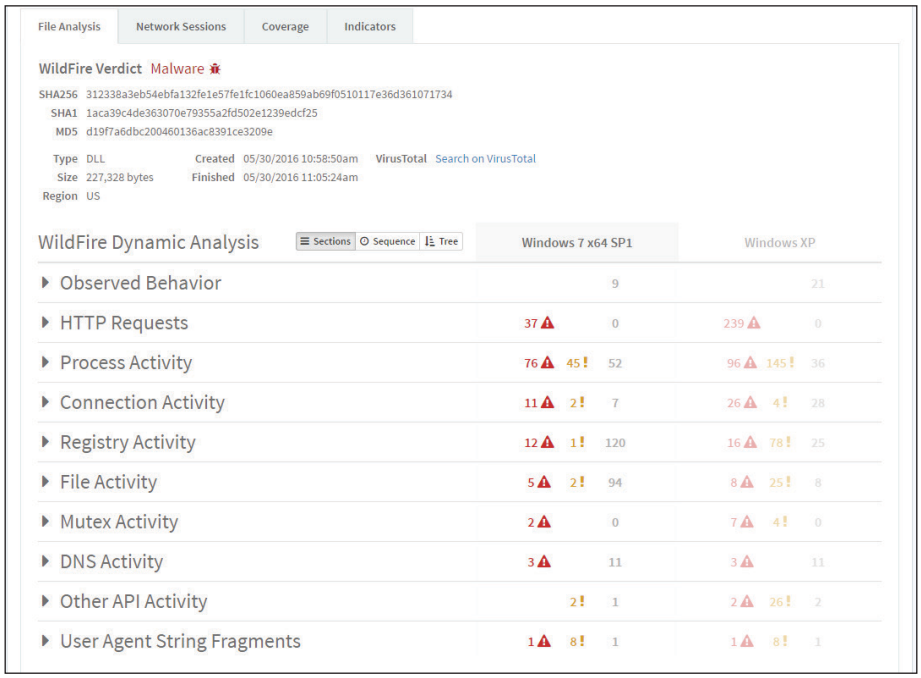

Any sessions that had a verdict of Malicious should be investigated further because the firewall does not cache the sample and wait for the result of the analysis before forwarding the traffic to the user. There will be a small period between when the first sample of malware is seen on a Palo Alto Networks security device and when the signatures are created and available globally for prevention. During this period between submission and signatures' availability to the next-generation firewalls, Traps endpoint protection is critical to ensuring prevention on the endpoint.

Logging to the Logging Service

Investigating historical or real-time security-relevant events depends heavily on the data captured in the firewall logs. The right data gives you the ability to derive actionable intelligence. It lets you identify known attacks that occurred in the past in order to understand what they took, added, or modified; to identify known attacks occurring presently and what they are targeting; and to identify known indicators of possible future attacks. It gives you this past, present, future view of known attacks so you can disclose, clean up, or lock down as needed. It also provides the ability to analyze this rich data source for previously unknown attack techniques, new vectors, and previously unidentified targets allowing you or others to develop new indicators of compromise, preventions, and remediation techniques. From this data, you can also develop an understanding of your adversaries and their playbooks and adapt your security posture to best match that of those who target your data.

The Palo Alto Networks Logging Service is a cloud-based log collector service designed to collect logs; perhaps more accurately, it is designed to collect data. The Logging Service collects and stores large amounts of log data, helping to solve challenges around security-relevant log storage and retention. That data can then be the basis for developing and delivering a whole new set of security applications. It also serves as the foundation for the Palo Alto Networks Application Framework.

The Logging Service acts as a centralized logging destination instead of dedicated logging to a Log Collector appliance or retaining logs on the local firewall itself. The Logging Service leverages a secure, multi-tenant cloud model to provide high availability and fault tolerant data protection. The cloud-based service provides massive search capabilities at scale. Using secure SSL/TLS connections, logs are sent by one or more firewalls and GlobalProtect cloud service instances to the Logging Service. These secure connections are established between the service and the sender by using a signed certificate provided by the Logging Service itself. These sessions are re-negotiated periodically to change the encryption key in use. Firewalls also encrypt the logs before sending them over a secured session.

Protecting Servers with Traps

Traps advanced endpoint protection stops threats on the server and coordinates enforcement with cloud and network security to prevent successful cyberattacks. Traps minimizes server infections by blocking malware, exploits, and ransomware.

Attackers often blend two primary attack methods to compromise organizations: targeting application vulnerabilities through exploits and deploying malicious files—including ransomware. These methods can be used individually or in various combinations, but they are fundamentally different in nature:

- *Exploits* are the results of techniques used against a system that are designed to gain access through vulnerabilities in the operating system or application code.
- *Malware* is a file or code that infects, explores, steals, or conducts virtually any behavior an attacker wants.
- *Ransomware* is a form of malware that holds valuable files, data, or information for ransom, often by encrypting data, with the attacker holding the decryption key.

Due to the fundamental differences between malware and exploits, effective prevention requires an approach that protects against both. Traps combines multiple methods of prevention at critical phases within the attack lifecycle to halt the execution of malicious programs and stop the exploitation of legitimate applications.

MULTI-METHOD MALWARE PREVENTION

Traps prevents the execution of malicious files with an approach tailored to combating both traditional and modern attacks. Additionally, administrators can utilize periodic scanning to identify dormant threats, comply with regulatory requirements, and accelerate incident response with endpoint context.

The primary prevention capabilities include:

- **WildFire threat intelligence**—In addition to third-party feeds, Traps leverages the intelligence obtained from tens of thousands of subscribers to the WildFire cloud-based threat analysis service to continuously aggregate threat data and maintain the collective immunity of all users across endpoints, networks and cloud applications.
    - Traps queries WildFire with the hash of any Windows or macOS executable file, DLL, or Office file before the file runs to assess its standing within the global threat community. WildFire returns a near-instantaneous verdict on whether the file is malicious or benign. If the file is unknown, Traps proceeds with additional prevention techniques to determine whether it is a threat that should be terminated.
    - If the file is deemed malicious, Traps automatically terminates the process and optionally quarantines it.

- **Local analysis via machine learning**—If a file remains unknown after the initial hash lookup and has not been identified by administrators, Traps uses local analysis via machine learning on the endpoint—trained by the rich threat intelligence of WildFire—to determine whether the file can run, even before receiving a verdict from the deeper WildFire inspection. By examining hundreds of file characteristics in real time, local analysis can determine whether a file is likely malicious or benign without relying on signatures, scanning, or behavioral analysis.
- **WildFire inspection and analysis**—In addition to local analysis, Traps sends unknown files to WildFire for discovery and deeper analysis to rapidly detect potentially unknown malware. WildFire brings together the benefits of independent techniques for high-fidelity and evasion-resistant discovery that go beyond legacy approaches. These techniques include:
  - **Static analysis via machine learning**—A more powerful version of local analysis, based in the cloud, that detects known threats by analyzing the characteristics of samples before execution.
  - **Dynamic analysis**—A custom-built, evasion-resistant virtual environment in which previously unknown submissions are detonated to determine real-world effects and behavior.
  - **Bare metal analysis**—A hardware-based analysis environment specifically designed for advanced threats that exhibit highly evasive characteristics and can detect virtual analysis.

If WildFire determines a file to be a threat, it automatically creates and shares a new prevention control with Traps and other components of the platform in as few as five minutes, ensuring the threat is immediately classified as malicious and prevented, should it be encountered again.

Additional prevention capabilities include:

- **Granular child process protection**—Traps prevents script-based and fileless attacks by default with out-of-the-box, fine-grained controls over the launching of legitimate applications, such as script engines and command shells, and continues to grow these controls through regular content updates directly from the Palo Alto Networks threat research team, Unit 42. Administrators have additional flexibility and control with the ability to whitelist or blacklist child processes, along with command-line comparisons, to increase detection without negatively impacting process performance or shutting them down.
- **Behavior-based ransomware protection**—In addition to existing multi-method prevention measures—including exploit prevention, local analysis and WildFire—Traps monitors the system for ransomware behavior. Upon detection, it immediately blocks attacks and prevents encryption of customer data.
- **Scanning**—Administrators can scan endpoints and attached removable drives for dormant malware, with an option to automatically quarantine it for remediation when found. Periodic or on-demand scanning can be configured as part of a security profile on one or more endpoints.
- **Admin override policies**—Traps enables organizations to define policies based on the hash of an executable file to control what is or isn't allowed to run in their environments. This not only reduces the attack surface but eliminates negative impact on homegrown or heavily customized applications.

- **Malware quarantine**—Particularly useful in preventing the inadvertent dissemination of malware in organizations where network- or cloud-based data storage and SaaS applications automatically sync files across multiple users and systems, Traps immediately quarantines malicious executable files, DLLs, and Office files in order to prevent propagation or execution attempts of infected files.
- **Grayware classification**—Traps enables organizations to identify non-malicious but otherwise undesirable software, such as adware, and prevent it from running in their environments.
- **Execution restrictions**—Traps enables organizations to easily define policies to restrict specific execution scenarios in order to reduce the attack surface of any environment. For example, Traps can prevent the execution of files from the Outlook "temp" directory or a particular file type from a USB drive.

MULTI-METHOD EXPLOIT PREVENTION

Rather than relying on signatures or behavior-based detection to identify exploit-based attacks, Traps targets the techniques any exploit-based attack must use to manipulate a software vulnerability. By preventing the techniques instead of identifying each attack, Traps can protect unpatched systems, unsupported legacy systems, and zero-day exploits. Traps delivers exploit prevention using multiple methods:

- **Pre-exploit protection**—Traps prevents the vulnerability-profiling techniques that exploit kits use prior to launching attacks. By blocking these techniques, Traps prevents attackers from targeting vulnerable endpoints and applications, effectively preventing the attacks before they begin.
- **Technique-based exploit prevention**—Traps prevents known, zero-day and unpatched vulnerabilities by blocking the exploitation techniques attackers use to manipulate applications. Although there are thousands of exploits, they typically rely on a small set of exploitation techniques that change infrequently. Traps blocks these techniques, thereby preventing exploitation attempts before they can compromise endpoints.
- **Kernel exploit prevention**—Traps prevents exploits that leverage vulnerabilities in the operating system kernel to create processes with escalated (i.e., system-level) privileges. Traps also protects against new exploit techniques used to execute malicious payloads, such as those seen in 2017's WannaCry and NotPetya attacks. By blocking processes from accessing the injected malicious code from the kernel, Traps can prevent the attack early in the attack lifecycle without affecting legitimate processes. This enables Traps to block advanced attacks that target or stem from the operating system itself.

By blocking the techniques common to all exploit-based attacks, Traps provides customers three important benefits:

- **Protects unpatchable applications and shadow IT**—Providing a positive work experience is critical to the productivity of any organization, but running unsupported legacy applications or granting users the flexibility to download and run programs as they please introduces risk. Traps enables organizations to run any applications, including those developed in-house, no longer receiving updates or security support, or running in their environment without IT's awareness, without opening the network to the threat of exploit-based attacks.

- **Eliminates the urgency to patch applications as soon as possible**—Organizations using Traps can apply security patches when it is appropriate for the business and after sufficient testing. Traps prevents the exploitation of application vulnerabilities regardless of when an organization applies security patches issued by application vendors.
- **Prevents zero-day exploits from succeeding**—Because Traps blocks the limited set of exploitation techniques zero-day exploits typically use, Traps protects organizations against attacks that utilize zero-day exploits.

TRAPS ON INTERNAL AND PUBLIC CLOUD SERVERS

Traps advanced endpoint protection stops threats on the server and minimizes server infections by blocking exploits, malware and ransomware. Like with the next-generation firewall, you can use a Zero Trust policy in Traps to extend the least-privileged access model to protecting the server operating system.

Figure 13 Traps protecting servers

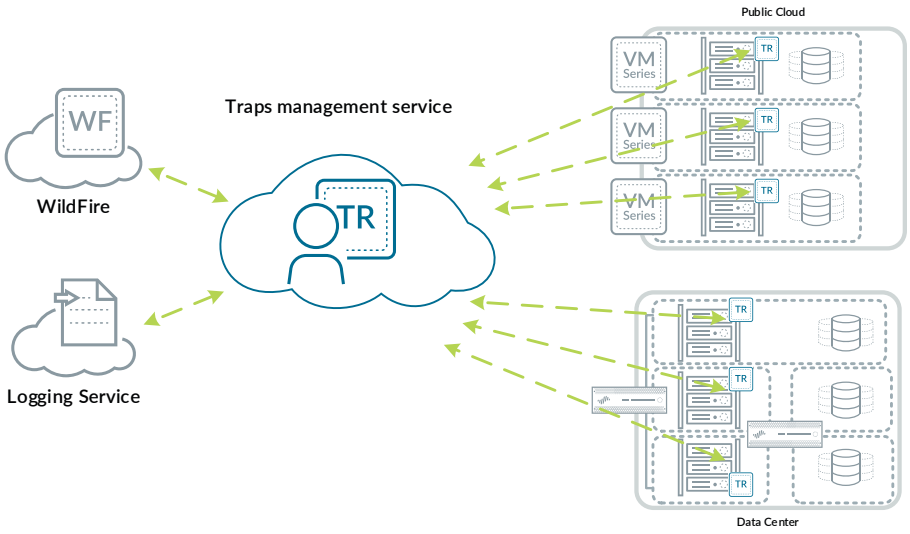

Because attackers most often target application vulnerabilities when attempting to compromise servers, the Traps exploit-prevention profile is key to extending the Zero Trust security model to servers and blocking the core techniques used by zero-day exploits. In Traps Management Service, enable the following exploit protection capabilities on the Windows and Linux server agents:

- **Known Vulnerable Processes Protection**—Common applications in the operating system can contain bugs and vulnerabilities that an attacker can exploit. By enabling this capability, Traps protects these processes from attacks that try to exploit known process vulnerabilities.
- **Exploit Protection for Additional Processes**—Extend protection from exploitation attempts to the processes that the known-vulnerable process-protection doesn't cover. Add processes that interact with the sensitive data or end-users to this list.

- **Operating System Exploits Protection**—Attackers leverage the operating system itself to accomplish malicious actions. Enable this capability so that Traps protects operating system mechanisms, such as privilege escalation, and prevents attackers from using them for malicious purposes.

Like traffic flows, the applications that run on the servers in your private data center and public cloud environment should be well defined. Use restriction profiles to reduce the avenues from which an attacker can compromise your servers. For your highest sensitivity Windows servers, define restriction policy that restricts:

- The locations from which executable files can run.
- Access to all network locations except for those that the application explicitly requires.
- The executable files that users can launch from external drives.
- The executable files that users can launch from optical disc drives.

Windows servers should also have the following malware protection capabilities enabled:

- **Ransomware protection**— Targets encryption-based activity associated with ransomware in order to analyze and halt ransomware before any data loss occurs.
- **Prevent malicious child process execution**—Prevents script-based attacks used to deliver malware by blocking known targeted processes from launching child processes that attackers commonly use to bypass traditional security approaches.
- **Portable executables and DLLs examination**—Analyze and prevent malicious executable and DLL files from running. Whitelist the files and executable signers that your application requires.

The Traps Management Service provides granular control over service-protection settings and security of the Traps agent running on the servers. This default protection prevents attempts to disable or make changes to Traps processes, services, registry keys and values, and files.

Traps Management Service sends all of its logs to the Logging Service for storage and to be used as data for security analytics running on the Application Framework. Traps logs provide detail on the processes and activities that generate and receive traffic that was logged by the next-generation firewall.

Protecting the Public Cloud with Evident

The policies and capabilities used to secure data and applications in the public cloud are effectively the same as those used in the private data center. One major area of difference is that you do not have complete control over the infrastructure. This is troublesome because a single mistake in the configuration of the cloud environment can leave it open to attack and loss of sensitive data. Another area of difference is that the data stored in the public cloud environment is at greater risk of exposure than data in the private data center. So, when high and moderate sensitivity data and applications live in the public cloud, you must secure the cloud infrastructure and govern the data stored in the cloud environment as part of Zero Trust. Evident helps secure your public cloud environments and provide data governance to public cloud storage.

Figure 14 Evident integration into the public cloud environment

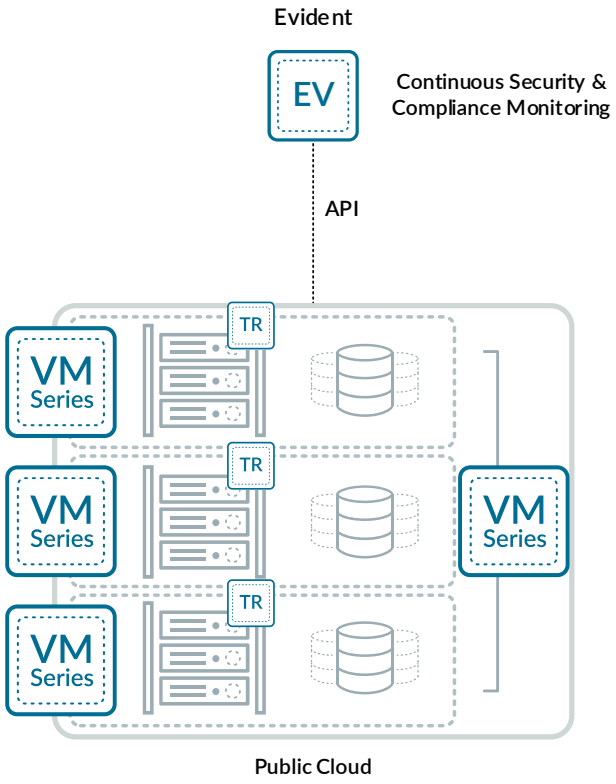

EVIDENT CONTINUOUS MONITORING AND COMPLIANCE REPORTING

Evident automatically validates best security practices against your public cloud resources in AWS and Azure. Evident continuously monitors your public cloud environments and alerts you when resources fail to pass validation. For all public cloud resources regardless of data sensitivity, enable Evident to monitor all security signatures including:

- **Default VPC NACL**—Monitors for the default NACL that AWS supplies with a new VPC, which is insecure by default.
- **SSL Certificate Expiration**—Monitors for SSL certificates that are or are about to expire.
- **S3 Bucket with Global Upload and Delete ACL Permissions Enabled**—Monitors for S3 permissions that permit anyone to add, update, or remove the content in your S3 bucket.

Alerts with:

- High severity are a great potential risk to your business assets and you should examine them as soon as possible.
- Medium severity should be scheduled to be fixed and tracked.

For the quickest remediation of alerts, configure Evident integrations with the tools your organization uses to track work requests. Evident integrates into ServiceNow, Splunk, Slack, Jira, and others.

Evident also can monitor the cloud environment for compliance with industry and regulatory standards. Configure Evident to monitor your cloud environment for compliance with the standards that are relevant to your organization. Supported Compliance benchmarks include HIPAA, ISO 27001, NIST 800-53/FedRAMP, NIST 800-171, PCI, and SOC 2.

EVIDENT PUBLIC CLOUD ENVIRONMENT STORAGE SECURITY

After your data leaves your network and is stored in the public cloud environment's storage, in-line network security devices can't see access and changes to the data. The Zero Trust security model requires visibility and control over sensitive data regardless of the data's location. Building on the Zero Trust protections at the internet perimeter, Evident secures the data stored in public cloud buckets and containers.

Evident connects directly to storage service's API. This connection provides visibility and control over the data. Evident automatically assesses risk through content, activity, and security control policies. Content policies scan the content of data for information that is critical, sensitive, or subject to compliance and assign a risk value. You can mitigate the risk automatically or manually by quarantining, changing share access, or alerting the owner or an administrator.

Evident can manage AWS S3, Google Cloud Storage, and Microsoft Azure, giving you consistent visibility and control across each.

Summary

Breaches and data loss have serious consequences for organizations and their customers. Zero Trust is a security model developed specifically to address the security of sensitive data and critical applications in an enterprise organization. Zero Trust policy leverages the Palo Alto Networks Security Operating Platform capabilities and functionality, including:

- Next-generation and VM-Series firewalls provide the inline protection required to segment data from other applications and endpoints in the network.
- Traps advanced endpoint protection inspects system process execution and file system behavior through local and dynamic analysis.
- Evident inspects asset accessibility and risk through API integrations into the public cloud storage services.
- Logging Service serves as the central cloud-based repository for all security platform data and logs.

Related Guides

Zero Trust Overview—Introduces the concepts of the Zero Trust security model and how to implement it using the Palo Alto Networks Security Operating Platform.

Reference Architecture Guide for Cisco ACI—Presents a detailed discussion of the available design considerations and options for the next-generation firewall in a Cisco ACI-based private data center deployment.

Reference Architecture Guide for VMware NSX—Presents a detailed discussion of the available design considerations and options for the next-generation firewall in a VMware NSX-based software defined data center deployment.

You can use the feedback form to send comments about this guide.

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

B-000230P-1 08/18