Best Practices Implementing Zero Trust with Palo Alto Networks

Version 10.0

paloaltonetworks.com/documentation

techDOCS paloalto NETWORKS

Contact Information

Corporate Headquarters:
Palo Alto Networks
3000 Tannery Way
Santa Clara, CA 95054
www.paloaltonetworks.com/company/contact-support

About the Documentation

- For the most recent version of this guide or for access to related documentation, visit the Technical Documentation portal www.paloaltonetworks.com/documentation.
- To search for a specific topic, go to our search page www.paloaltonetworks.com/documentation/document-search.html.
- Have feedback or questions for us? Leave a comment on any page in the portal, or write to us at documentation@paloaltonetworks.com.

Copyright

Last Revised

October 28, 2020

Table of Contents

Zero Trust Best Practices....5
What Is Zero Trust and Why Do I Need It?....7
Zero Trust Viewpoint....8
Zero Trust High-Level Best Practices....8
How Do I Start My Zero Trust Implementation?....9
The Five-Step Methodology....10
Step 1: Define Your Protect Surface....10
Step 2: Map the Protect Surface Transaction Flows....11
Step 3: Architect a Zero Trust Network....12
Step 4: Create the Zero Trust Policy....13
Step 5: Monitor and Maintain the Network....16
Zero Trust Resources....17

Zero Trust Best Practices

This document describes what a Zero Trust strategy is and how to implement it in your network using a five-step methodology that guides you through best practices for identifying your critical protect surfaces, mapping your critical transaction flows, architecting your Zero Trust network, creating Zero Trust policy, and maintaining the deployment. Sections include links to detailed information from Palo Alto Networks, including how to configure next-generation firewalls (physical and virtual) and security capabilities from Palo Alto Networks to prevent data breaches.

> What Is Zero Trust and Why Do I Need It?
> Zero Trust Viewpoint
> The Five-Step Methodology
> Zero Trust Resources

What Is Zero Trust and Why Do I Need It?

Zero Trust is a business-driven, strategic approach to securing your most critical data, applications, assets, and services (DAAS) as well as your users based on what is important to your particular business, in a *protect surface*. Zero Trust strategy is infrastructure-neutral, so you can apply it all physical and virtual locations—network, public cloud, private cloud, and endpoint. The concept behind Zero Trust is simple: trust is a vulnerability. Trust nothing in the digital environment—packets, identities, devices, or services—and verify everything. There is no such thing as default trust.

Implementing the strategy is not something you do once and cookie-cutter copy from network to network because each environment and protect surface is different; and as businesses change over time, the goal and DAAS elements also change. Strategy is business-specific and security strategy is specific to protecting what's important to your particular business.

The goal of Zero Trust strategy is to eliminate trust from the network. Eliminating trust helps prevent successful data breaches, simplifies operations through automation and a reduced rulebase, and simplifies regulatory compliance and audits because Zero Trust environments are designed for compliance and easy auditing.

Zero Trust Viewpoint

When you understand Zero Trust, you see trust for what it is—a vulnerability that attackers exploit. Attackers can steal credentials, spoof information in packet headers, and even be "trusted" employees or partners. Edward Snowden was a trusted user who had the right antivirus software and the right patch level on his workstation. He also used Multi-Factor Authentication. But nobody cared about where he went on the network or the packets he generated because he was a trusted user, so he could explore the network and find and exfiltrate sensitive data. The lesson is that outcome of digital trust is digital betrayal; don't trust identities, applications, or data. When you take a Zero Trust viewpoint, you:

- Align security with business functions because business functions determine what you need to protect.
- Inspect and log all packets at Layer 7 when they access a resource.
- Access all resources in a secure manner regardless of location.
- Apply consistent security policy in all locations.
- Manage security and segmentation policy centrally.
- Accommodate changes as your business changes.

Trust is a failure point you avoid by implementing a Zero Trust strategy.

- Zero Trust High-Level Best Practices
- How Do I Start My Zero Trust Implementation?

Zero Trust High-Level Best Practices

The following best practices prepare for and help you transition your network to a Zero Trust architecture:

- Define your desired business outcomes before architecting your Zero Trust environment. The Zero Trust model supports and enables secure business functions.
- Design from the inside-out instead of from the outside-in to protect what's most valuable to your business first. Your most valuable assets are more likely to be in your data center than at your perimeter.
- Use an integrated, centrally managed platform that reduces the total cost of ownership, rather than a collection of point products that don't work well together. Palo Alto Networks shares information among platform elements and enables centralized management and simplified operation using Panorama, GlobalProtect, and Prisma Access to provide consistent policy, prevention, and protection across all locations.
- Use Palo Alto Networks Next-Generation Firewalls as segmentation gateways to consolidate security technologies on one platform and to apply consistent security policy in all locations natively at Layer 7 using App-ID, User-ID, and Content-ID. A segmentation gateway segments and controls the network based on applications, users, and data, and should provide granular access control and secure all traffic as it crosses microperimeters and gains access to a protect surface.

  You don't need to change your infrastructure to create microperimeters because you create microperimeters in Layer 7 policy by allowing only authorized users to access only the protect surfaces they need to access for business purposes.
- Segment your network based on what's valuable to your business to prevent unauthorized lateral movement.
- Apply the principle of least-privileged access to your protect surfaces. Determine who needs access to what resources, how they need access, and when they need access. Allow only the exact level of access required for each user and device, assert identity (including proper authorization), and then map Layer 7 policy to identity.
- Decrypt, inspect, and log every packet through Layer 7 that regulations, compliance, and your business practices allow you to inspect. You must inspect and log Layer 7 traffic. Remember, every attacker knows how to bypass security controls at Layer 3 and Layer 4.

- Create a strategy for tagging workloads to group objects and registering tags dynamically to help automate security policy.
- Develop processes to operate, maintain, and continually update prevention controls as you develop your strategy and design the network. Document processes, educate and train personnel, set baselines, and measure progress against the baselines.
- Transition to a Zero Trust environment gradually, one segment at a time, beginning with one or more non-critical segments from which you learn and gain experience. Zero Trust segments coexist with legacy segments, so you can use a safe, iterative approach instead of a risky rip-and-replace approach.

As the importance of applications diminishes, you can be less aggressive with protection. For example, you don't need to apply the same protection to a chat app that you need to apply to business-critical apps. Collaboration with business leaders helps to determine which applications are the most critical to protect.

How Do I Start My Zero Trust Implementation?

Education and collaboration begin the journey to Zero Trust security. You and other stakeholders who will identify what's valuable to your business and how to protect it need to understand Zero Trust concepts, principles, and goals.

1. Create a Zero Trust Center for Excellence. This is a cross-functional team of business leaders (business and technical decision makers), IT, information security, infrastructure, application developers, and other stakeholders. The team defines and identifies each protect surface and the data, applications, assets, and services (DAAS elements) that make up each protect surface. They prioritize the most valuable protect surfaces for your business and plan and implement the Zero Trust strategy. The team remains involved in maintaining the deployment as the business changes. Business leaders can speak to desired business outcomes, compliance requirements, and the value of business assets.
2. Attend a Zero Trust workshop to prepare everyone and get everyone on the same page. Contact your Palo Alto Networks sales representative for more information and to schedule a workshop.
3. Follow The Five-Step Methodology to map the segmented network you want to build.
4. Start the transition with one or more small, well understood, low-risk (not critical to business operation) segments to learn from the experience. Don't start with critical assets. Next, test your learning on one or more practice segments. When you feel ready, place your most business-critical protect surfaces (the DAAS elements that make up the protect surface) in Zero Trust microperimeters, one microperimeter per protect surface. After that, convert the next most valuable set of protect surfaces to Zero Trust, etc.

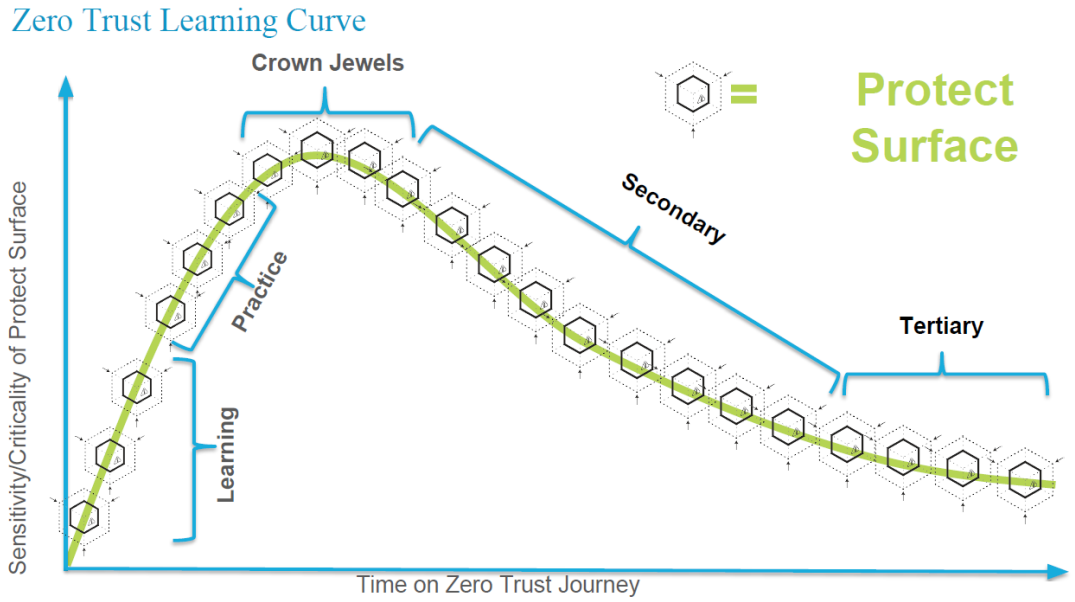

The Five-Step Methodology

The five-step methodology for implementing a Zero Trust strategy presents a logical, clear path to protecting your environment, data, applications, assets, services, and users. The way you apply the methodology depends on what you're protecting and your business requirements—what's critical to your business—but the outcomes you're working toward are the same:

- Segment the network effectively and efficiently to prevent lateral movement.
- Protect business-critical data and systems from unauthorized applications and users.
- Protect business-critical applications from unauthorized access and usage.
- Enforce policy seamlessly across network, cloud, and endpoints to simplify management and apply consistent policy everywhere.

The five-step methodology works whether you're implementing a Zero Trust strategy in the cloud, on a private network, or on endpoints, regardless of infrastructure.

- Step 1: Define Your Protect Surface
- Step 2: Map the Protect Surface Transaction Flows
- Step 3: Architect a Zero Trust Network
- Step 4: Create the Zero Trust Policy
- Step 5: Monitor and Maintain the Network

Step 1: Define Your Protect Surface

A protect surface is what's valuable to your business: the data, applications, assets, and services (DAAS) you need to protect to ensure normal business operation. Defining your protect surface enables you to focus on defending what really matters to your business instead of trying to identify and protect the entire attack surface or focusing on just the perimeter. The protect surface is also much smaller than the attack surface or the perimeter, so it's easier to protect.

Define your protect surface based on the most crucial DAAS elements for your business:

- **D**ata. What data needs to be protected? Think about intellectual property such as proprietary code or processes, personally identifiable information (PII), payment card information (PCI), and personal health information (PHI) such as Health Insurance Portability and Accountability Act (HIPAA) information.
- **A**pplications. Which applications consume sensitive information? Which applications are critical for your business functions?
- **A**ssets. Which assets are the most sensitive? Depending on your business, that could be SCADA controls, POS terminals, medical equipment, manufacturing equipment, and groups of critical servers.
- **S**ervices. Which services can attackers exploit to disrupt IT operations and negatively impact the business, such as DNS, DHCP, and Active Directory?

Each critical DAAS element is part of a protect surface (or in some cases is a protect surface). For example, if your business provides health care, then personal health information (PHI) is critical to your business. The *Data* is the patient information. The *Applications* are the applications used to access PHI data—for example, EPIC. The *Assets* are servers that store the data and equipment that generates PHI, such as medical scanners or physicians' workstations. The *Services* are services used to access the data, such as single sign-on and Active Directory.

As you follow the five-step methodology, you'll place each protect surface in its own microperimeter (segmented by a Palo Alto Networks physical or virtual next-generation firewall, which acts as a segmentation gateway) so that you control exactly who accesses the element, how they access it, and when they access it. Secure each protect surface in manner that is appropriate for that protect surface. A microperimeter is easier to manage and defend than a broad perimeter that encompasses DAAS elements

that users with different access requirements need to reach. It also moves protections closer to the critical data.

Prioritize what to protect first based on what's critical to running your business. Your most valuable assets are often in your data center or in the cloud. After you implement Zero Trust on one or more non-critical protect surfaces to gain experience, defend your most critical protect surfaces. You may not know all of the applications in your data center when you start, but you know your most critical applications. Afterward, move on to the next set of protect surfaces on the priority list and keep going through the list until you reach your security goals.

Use the following tools to gain visibility into your network traffic and help identify the DAAS elements that make up your most critical protect surfaces:

- The team's knowledge of the business. For example, business leaders can speak to the strategic value of applications.
- Insert one or more next-generation firewalls transparently into your network in virtual wire (vwire) mode, which is a passthrough mode that requires no topology changes because vwire interfaces don't have IP or MAC addresses, to gain visibility into traffic. Check the Traffic logs to view and analyze network traffic. If you already have managed firewalls in your network, use Panorama logs.
- View logs in the Cortex Data Lake and use third-party asset discovery tools that work with Cortex from one of Palo Alto Networks' integrated partners.
- Use Prisma SaaS to discover users, assets, and data for SaaS applications and gain visibility into those applications.
- If you run PAN-OS 9.0 or later on the next-generation firewall or on the Panorama that manages your firewalls, use Policy Optimizer to help identify key applications on existing Security policy rules. (Policy Optimizer even shows you all of the applications on port-based rules.) If you can't use Policy Optimizer, use Expedition to gain visibility into applications.
- Application Dependency Mapping tools to discover application dependencies (the resources an application uses, such as databases, load balancers, servers, etc.) automatically.

Step 2: Map the Protect Surface Transaction Flows

Map the transaction flows (interactions) between your critical DAAS elements and users to understand their interdependencies—who has business reasons to access each element, in what manner, and at what time. Map the transaction flows to understand and architect the network. Mapping helps you understand how to create security policy that allows only authorized users access to specific data and assets using the specified applications (principle of least-privileged access).

There are many ways to map transaction flows, and some techniques for defining your protect surface also apply to mapping its transaction flows:

- Leverage existing flow diagrams if you have them (compliance and auditing sometimes require businesses to create flow diagrams).
- Work with application, network, and enterprise architects, and business representatives to understand the purpose of applications and the transaction flow the architects and business representatives envision.
- Insert one or more next-generation firewalls transparently into your network in virtual wire (vwire) mode to gain visibility into traffic. Check Traffic logs to view and analyze traffic.
- Use third-party tools from Palo Alto Networks' integrated partners.
- Use log information from the Cortex Data Lake to gain visibility into and map transaction flows. The Cortex Data Lake aggregates logs from the next-generation firewall, VM-Series firewalls, Prisma Access, and Cortex XDR.
- For applications, map the workflows, including the flow of application data across the network, the computing objects required for each application, and who uses each application.
- For data, find out who uses the data, where you collect, store, use and transfer the data, and how the data is stored, encrypted, archived, or destroyed after use.

- For assets, find out the asset's location, who uses the asset, when they use the asset, and where the asset fits into workflows.
- For services, map the service workflows across the environment.

In addition to revealing who uses what applications where and when, mapping transaction flows provides granular visibility that aids with disaster recovery planning and compliance. It also gives you an opportunity to optimize workflows and examine who has legitimate business reasons to access the DAAS elements in each protect surface.

When you understand transaction flows through your network, you'll know how to segment the network and where to insert controls because you'll understand who uses each protect surface, how they use it, where it's located, and which elements interact to enable each critical application.

Step 3: Architect a Zero Trust Network

Armed with an understanding of your protect surface and transaction flows, begin architecting your Zero Trust network based on what's valuable to your business. Architect the business-critical protect surfaces you identified in Step 1: Define Your Protect Surface from the inside-out. As you develop the architecture, keep in mind ease of operation and maintenance, and flexibility to accommodate protect surface and business changes. Run the Best Practice Assessment tool to set a best practice configuration baseline and measure progress toward your Zero Trust goals.

The cornerstone of the architecture is segmentation gateways—physical or virtual Palo Alto Networks Next-Generation Firewalls that connect your network segments and enforce Layer 7 policy. Run all traffic through a segmentation gateway, place segmentation gateways as close as possible to the resources they protect, and use them in conjunction with other Palo Alto Networks capabilities to automate as much as possible. Next-generation firewalls:

- Create a microperimeter in Layer 7 policy around each protect surface. This prevents lateral movement because the microperimeter provides granular policy controls for who (User-ID) accesses what applications (App-ID) and resources in what manner (Content-ID) and at what time through the segmentation gateway. Segment based on how transactions flow across your network and how your users and applications access data and services.
- Aggregate security capabilities into a single control point for all traffic entering and exiting the protect surface. The segmentation gateway should enforce policy, decrypt encrypted traffic, and apply protections such as:
  - DNS Security (use the DNS Security service, which provides multiple real-time threat intelligence sources, infinitely scalable real-time analysis of DNS requests, and advanced DNS signatures).
  - Intrusion prevention (Vulnerability Protection, Anti-Spyware, and Antivirus profiles).
  - Blocking potentially dangerous file types.
  - Preventing unknown and Day 1 threats (WildFire).
  - URL Filtering.
  - Data Loss Prevention (DLP).
- Decrypt and inspect traffic at Layer 7 in real-time.
- Log every packet from Layer 2 through Layer 7. Send logs to the Cortex Data Lake from Panorama for managed firewalls, from individual firewalls (firewalls not managed by Panorama), from Prisma Access (formerly GlobalProtect™ cloud service), and from Cortex XDR to centralize and aggregate your on-premise and virtual (private and public cloud) log storage for physical and VM-Series firewalls.
- Use APIs for tight integration with third-party defense tools from partners.
- Automate feedback loops that detect events and automate responses.
  - Tag workloads and use tags as filtering criteria to determine the members of dynamic address groups in security policy. This enables you to automate actions based on log forwarding events to an HTTP(S) server. The log forwarding event triggers the action by dynamically adding or removing members of a dynamic address group used in security policy in real-time. The security policy

determines if the members of the dynamic address group are allowed or denied access and the firewall enforces the action. For example, set up a DNS sinkhole in an Anti-Spyware security profile to automatically quarantine potentially compromised systems that attempt to access the sinkhole. Use tags and log forwarding to add and remove those systems dynamically from a dynamic address group that is attached to a policy rule which blocks and logs all traffic to the sinkhole address. You can then investigate potentially compromised systems when notified by log alerts.
    - Use Cortex XDR to automate analyzing your network, discovering anomalous behavior that indicates a potential intrusion, and alerting on that behavior so you can investigate and remediate the issue. Cortex XDR provides visibility into network traffic, simplifies threat investigation by correlating logs, and enables you to identify the root cause of alerts and respond immediately. Use Cortext XDR APIs to integrate with Cortex XSOAR and automate responses using Desmisto response playbooks that are tailored to your business workflows, which can reduce response time from days to minutes.
    - Use WildFire to automate discovery of new malware. When WildFire discovers malware anywhere in the world, it takes at the most five minutes before WildFire updates your security profiles to protect you against the new malware.
- Use templates and template stacks in Panorama to automate policy deployment.
- Use tools such as Ansible, Terraform, and Python to automate, orchestrate, and accelerate protecting Prisma Cloud deployments.

Palo Alto Networks enables you to architect your Zero Trust environment and apply consistent security across all locations:

- Panroma centralizes management policy control for multiple next-generation firewalls and increases operational efficiency compared to managing firewalls individually.
- Corporate network and data center: Use next-generation firewalls to segment the network into microperimeters for your protect surfaces.
- Public cloud: Use Prisma Access, which uses on-premise or VM-Series next-generation firewalls, and Prisma Cloud (an API-based cloud infrastructure security solution), to implement Zero Trust policy in cloud environments. Virtual private clouds (VPCs) define protection boundaries to segment workloads.
- Private cloud: Use VM-Series firewalls to implement Zero Trust policy.
- Branch office and mobile users: Use Prisma Access to provide cloud-based security and to avoid round-trips to corporate network resources. Configure Prisma Access for users and also Prisma Access for networks to secure branches.

  Alternatively, use an on-premises next-generation firewall with the GlobalProtect subscription service to extend security policy and enforcement to remote users and branch offices.
- Endpoints: Layer protection using the next-generation firewall for segmentation and the first layer of protection and using Cortex XDR agent for the second layer of protection. Enforce consistent policy using GlobalProtect (on-premise installation) or Prisma Access (installed using Panorama and managed for you in the cloud) VPNs to extend policy to remote endpoints and enable policy to move with the user. Prisma Access requires the GlobalProtect app on mobile user endpoints. In all cases, install the GlobalProtect app on managed endpoints and use GlobalProtect Clientless VPN on unmanaged endpoints (endpoints on which you can't or don't want to place an agent, such as partner systems or personal devices). Apply Multi-Factor Authentication when appropriate to protect high-value assets.
- SaaS applications: Use Prisma SaaS to scan, analyze, classify, and help protect SaaS applications. Redirect SaaS application traffic for unmanaged devices through your next-generation firewall (traffic from managed devices goes through Prisma Access, GlobalProtect, or a next-generation firewall).

Step 4: Create the Zero Trust Policy

Zero Trust policy consists of allow rules—rules that allow only authorized users to access specific resources using the specified applications at the right time in the right places. If traffic doesn't match a rule, the firewall automatically blocks the traffic. This is important because:

- It's much easier to know the applications you want to allow to support your business than to take on the never-ending task of identifying and blocking all the applications you don't want to allow.
- All breaches and malicious activity happen on allow rules. Focus security on traffic you allow and allow only the traffic required for business.

Zero Trust policy is based on the Kipling Method. Answering Rudyard Kipling's 6-tuple of questions, "who, what, when, where, why, and how," shows you how to decide whether to allow or block traffic and how to create security policy that safeguards each protect surface. Palo Alto Networks provides the capabilities to implement the Kipling Method in security policy:

- **Who** should access a resource?
  - User-ID identifies users and enables you to control who accesses a resource in policy. Through a lens of least-privileged access (who needs to know?), allow access only to individuals, groups, and devices that have legitimate business reasons to access a resource.
  - Create Authentication policy to verify the identities of users when they attempt to access resources. Authentication policy also determines whether to require Multi-Factor Authentication (MFA).
  - Use MFA to protect sensitive services and applications by requiring at least one more authentication factor in addition to entering a password in Authentication Portal, such as a one-time-use code delivered to a cell phone or email, before the firewall allows access to sensitive services, applications, and resources. For remote users, configure GlobalProtect to facilitate MFA notifications (you must also configure MFA on the firewall).
  - For devices that use GlobalProtect, configure Host Information Profiles (HIPs) to define access policy for hosts, enforce policy on those hosts, and prevent devices that don't meet your security and maintenance standards from accessing resources. For example, you can use a HIP to ensure that endpoints have encryption enabled, the host's antivirus signatures are up-to-date, etc. If a host doesn't meet the HIP requirements, the security policy blocks access.
- **What** application is used to access the resource?
  - Create application-based Layer 7 policy using App-ID, which identifies applications regardless of port, protocol, or evasive tactics so that you allow only the right applications on your network. Policy based on Layer 3 and Layer 4 relies on IP addresses an attacker can spoof and leaves ports open to evasive applications.
  - Set the Service to application-default to safely enable applications on their default ports and prevent evasive applications from accessing your network on non-standard ports.
  - If the firewall runs PAN-OS 9.0 or later or a Panorama appliance running PAN-OS 9.0 or later manages firewalls running PAN-OS 8.1 or later, use Policy Optimizer to examine existing policy rules (both application-based rules and legacy port-based rules), identify unused rules, and identify rules with unused applications. For firewalls that run older versions of PAN-OS, use Expedition to examine policy rules. (If you need to migrate a legacy configuration to a PAN-OS device, follow the Best Practices for Migrating to Application-Based Policy.
- **When** do users access the resource?

  For applications users access only during certain hours, apply a schedule (**Objects** > **Schedules** on Panorama appliances and firewalls) to the policy rule to prevent suspicious access during off-hours. Adversaries often attack and attempt to exfiltrate data outside of normal business hours to reduce the chance of discovery.
- **Where** is the resource located?

  Add the location of the destination resource to the policy. When appropriate, also restrict the source (zone and IP address) of the traffic.
- **Why** is the data accessed—what is the data's value if lost (toxicity)?

  Classify data to understand its toxicity—why is the data worth protecting? Would you have to disclose the loss if an attacker exfiltrated the data? Set up Data Filtering to prevent sensitive information from leaving your network and use data classification tools to provide metadata about the data.

Understanding the toxicity of data helps you determine how to protect data, what to do with data after using it, and how to tag it for use in policy.

- **How** should you allow access to the resource?

  Apply Content-ID and best practices to protect against threats in application traffic:

  - Apply the philosophy of least-privileged access to security policy. Allow only users with legitimate business reasons to access only the applications they need to access for business purposes at only the proper times and only in the proper way.
  - Log all internal and external traffic through Layer 7. The firewall policy rules enable logging by default. Forward logs to the Cortex Data Lake (or to Panorama or to Log Collectors) to consolidate logs for easier and more thorough analysis.
  - Apply policy and threat prevention consistently across all locations (network, cloud, endpoints), for all local and remote users so the policy follows the user wherever the user goes, for all applications, and for all resources. Inconsistent policy increases vulnerabilities, is difficult to understand and maintain, and may negatively affect compliance requirements and audits. Use physical next-generation firewalls and virtual VM-Series firewalls as segmentation gateways to apply consistent Zero Trust, Layer 7, Kipling Method policy in the network and the cloud. Use Prisma Access (cloud) and GlobalProtect (on-premise installation and with Prisma Access) to extend consistent Zero Trust policy to endpoints. For unmanaged endpoints (endpoints on which you don't want to or can't place an agent), use GlobalProtect Clientless VPN to apply consistent policy. Create and reuse Panorama templates and stacks to apply consistent policy across similar locations, such as your data centers or your perimeters.
  - Configure security profiles (Vulnerability Protection profiles for IPS, Antivirus and WildFire profiles to protect against malware including day-one malware, Anti-Spyware profiles to prevent command-and-control threats, File Blocking profiles to block or alert on risky file types, and URL Filtering to control website access, help prevent phishing attacks, and enforce safe search for search engines) and apply them to all allowed traffic. Follow best practices for data center firewall and perimeter firewall security profiles.
  - Use WildFire best practices to detect and prevent zero-day malware.
  - Use decryption best practices to decrypt as much traffic as regulations and business requirements enable you to decrypt so you can inspect as much traffic as possible. You can't protect your network against threats you can't see.
  - Use the DNS Security service to provide infinitely scalable real-time access to DNS signatures, real-time analysis of DNS requests, and advanced DNS signatures generated using machine learning and predictive analysis.
  - How also includes determining what to do with sensitive data after you use it—abstract it using encryption, tokenization, or masking, or dispose of it by archiving or deleting it. Archive stale data (approximately 80% of data on most systems hasn't been accessed for two or more years).
  - Use Cortex XDR to refine and improve policy.

The Kipling Method enables you to create security policy that defends each protect surface appropriately because it leads you to understand who should have access, how they should access it, when they should access it, and the protections to apply. You develop policy rules by developing business statements based on the Kipling Method. For example:

	Who	What	When	Where	Why	How
Method	User-ID	App-ID	Time limits	System object	Classification	Content-ID
On-Premise	Epic_Users	Epic	Any	Epic_Srvr	Toxic (data has high value)	Decrypt, inspect (security profiles), log traffic

	Who	What	When	Where	Why	How
Cloud	Sales	Salesforce	Working hours	USA	Toxic (data has high value)	Decrypt, inspect (security profiles), log traffic

In both cases, the firewall allows only traffic that satisfies all of the conditions in the Kipling tuple and passes inspection. The firewall automatically denies all traffic that doesn't match an allow rule.

In addition to security, authentication, and decryption policy, use DoS and Zone protection best practices to protect vital servers from denial-of-service (DoS) attacks.

For firewalls that you haven't configured yet, use IronSkillet Day 1 configuration templates to implement a Day 1 best practice policy, then tune the policy to best suit your protect surfaces.

Step 5: Monitor and Maintain the Network

Security is an iterative process because logging and monitoring reveal improvements to make and because your business and network change over time. Follow the operational processes you developed when architecting the network to maintain and continually update prevention controls.

- Decrypt, inspect, and log all traffic (internal and external) through Layer 7.
- Forward logs to the Cortex Data Lake from Panorama for managed firewalls, from individual firewalls (firewalls not managed by Panorama), from Prisma Access, and from Cortex XDR to centralize and aggregate your on-premise and virtual (private and public cloud) log storage. This provides visibility into your network traffic and protect surfaces.
- Update policy and potentially add new protect surfaces based on intelligence from Cortex XDR, which uses Cortex Data Lake data and machine learning to automate analyzing your network based on your network's normal behavior and identifying anomalous behavior that may indicate an intrusion or other threat. Threat activity that targets DAAS elements which aren't in a protect surface can highlight protect surfaces you didn't consider when you originally defined your protect surfaces.
- Use Cortex XDR to gain visibility into your network traffic, simplify threat investigation by correlating logs, and enable you to identify the root cause of alerts and respond immediately.
- Use Cortext XDR APIs to integrate with Cortex XSOAR and automate responses using Desmisto response playbooks that are tailored to your business workflows, which can reduce response time from days to minutes.
- Use Prisma Cloud to aggregate and provide visibility into configuration data, user activity information, and network traffic information. Prisma Cloud analyzes data and delivers concise and actionable insights.
- Follow Best Practices for Applications and Threats Content Updates to get new and modified App-IDs and to keep your threat signatures up-to-date.
- Use the Best Practice Assessment tool to measure progress toward a best-practice configuration and to help you transition to a best practice security posture.
- Monitor network activity, use predefined reports, and generate custom reports to gain visibility into your environment.
- Keep the cross-functional team together to help maintain your Zero Trust deployment as the network and the business evolve, and create education and training to ensure that new members of the team understand the strategy and the implementation.
- Continue to automate actions and responses as automation capabilities advance.

Zero Trust Resources

The following technical documentation, white papers, webcasts, videos, and other resources provide more information and context for your Zero Trust strategy. In addition to the information in this document and the listed resources, you can engage Palo Alto Networks Professional Services consulting team of experts to help you design and implement your Zero Trust strategy.

- How to Build a Zero Trust Network (on-demand webcast)
- Debunk the Myth Around Implementing Zero Trust (on-demand webcast)
- Zero Trust Overview
- Zero Trust (Palo Alto Networks Zero Trust web page)
- Best Practices for Executing on Zero Trust (transformation roadmap)
- Simplify Zero Trust Implementation Using a Five-Step Methodology (white paper)
- Secure the Cloud: Zero Trust Cloud Security
- Zero Trust Cloud Security (video)
- The Truth About Zero Trust (infographic)

Palo Alto Networks Technical Documentation

Transition to Best Practices:

- Getting Started with the BPA
- How to Run a BPA (video)
- Understanding BPA Results (video)
- Live community Best Practice Assessment page

Best Practices Documentation Portal:

- Getting Started with Best Practices
- Internet Gateway Best Practice Security Policy
- Data Center Best Practice Security Policy
- Best Practices for Migrating to Application-Based Policy
- Best Practices for Securing Administrative Access
- Best Practices for Applications and Threats Content Updates
- Decryption Best Practices
- DoS and Zone Protection Best Practices
- WildFire Deployment Best Practices

Expedition

IronSkillet (Day 1 configuration templates)

Customer Support

Prevention Posture Assessment (complimentary consultative assessment of your prevention capabilities)

Palo Alto Networks NextWave Technology Partners