

SIMPLIFY ZERO TRUST IMPLEMENTATION WITH A FIVE-STEP METHODOLOGY

Companies are often reluctant to begin the Zero Trust journey because they believe it is difficult, costly, and disruptive. 20th-century design paradigms can create problems when designing a 21st-century Zero Trust network. However, building Zero Trust networks is actually much simpler than building legacy 20th-century hierarchical networks. Because most of us learned to design networks from the outside in, based on classifying users as “trusted” and “untrusted”—an approach that has since proven unsecure—we struggle to adapt our design thinking to the Zero Trust methodology.

It's not necessary to rip and replace your existing network to deploy Zero Trust. Zero Trust augments your existing network, with each Zero Trust network designed for a specific protect surface. The Zero Trust network is interconnected with your existing network to take advantage of the technology you already have. Then, over time, you iteratively move your additional data sets, applications, assets, or services over from your legacy network to your Zero Trust network. This helps make deploying Zero Trust networks manageable, cost-effective, and nondisruptive.

Following a five-step methodology supports deployment of your Zero Trust network with ease. This white paper provides an in-depth explanation of these steps and highlights how Palo Alto Networks provides a tightly integrated platform that aligns to each step, simplifying the process of protecting your critical assets.

The Five Steps to a Zero Trust Network




Figure 1: The five-step methodology

1. Define Your Protect Surface

The original five-step methodology talked about defining the sensitive data you need to protect, which makes up the protect surface. The protect surface is orders of magnitude smaller than the attack surface and is always knowable. Over time, it became apparent that Zero Trust protections should expand beyond data to include other elements of the network.

When defining the protect surface, you need to consider all critical data, application, assets, or services (DAAS). This could include:

- **Data**—payment card information (PCI), protected health information (PHI), personally identifiable information (PII), and intellectual property (IP)
- **Applications**—off-the-shelf or custom software
- **Assets**—SCADA controls, point-of-sale terminals, medical equipment, manufacturing assets, and internet of things (IoT) devices
- **Services**—DNS, DHCP, and Active Directory®



Palo Alto Networks Next-Generation Firewalls, in physical or virtualized form, provide comprehensive Layer 7 visibility to help you determine your DAAS profile. Palo Alto Networks also has extensive partnerships with leading third-party companies to help with additional data and asset discovery. Cortex XDR™ detection and response by Palo Alto Networks utilizes network, cloud, and endpoint products as sensors, feeding data into Cortex Data Lake to provide visibility into the activity of users, devices, applications, and services for greater insight into the individual protect surfaces across your enterprise environment.

Define Your Protect Surface by Deploying a Next-Generation Firewall Transparently in Your Network

Palo Alto Networks offers the unique ability to place a Next-Generation Firewall in line within your network without requiring any network changes. Known as a virtual wire deployment, this allows the Next-Generation Firewall to function as a “bump in the wire” on your network. This simplifies firewall installation and configuration, letting you insert the Next-Generation Firewall into existing topology without assigning MAC or IP addresses to the interfaces, redesigning the network, or reconfiguring surrounding network devices.

The device is invisible to the traffic, transparent to the network, and doesn't require network topology changes to actively see and block threats. With the Next-Generation Firewall in the path of your traffic, it can log for protect surface discovery purposes without any disruption to your network.

2. Map the Transaction Flows

To properly design a network, it's critical to understand how systems should work. The way traffic moves across the network, specific to the data in the protect surface, determines how it should be protected. This understanding comes from scanning and mapping the transaction flows inside your network in order to determine how various DAAS components interact with other resources on your network.


It's common to approximate flows by documenting what you know about how specific resources interact, even without a complete picture. This information still provides valuable data so that you don't arbitrarily implement controls with zero insight.

Zero Trust is a flow-based architecture. Once you understand how your systems are designed to work, the flow maps will tell you where you need to insert controls.

Remember that Zero Trust is an iterative process. Start with what you know. As you move through the steps in this methodology, you'll gather more information that will enable more granularity in your design. You shouldn't delay your Zero Trust initiative just because you don't have perfect knowledge.

Palo Alto Networks Next-Generation Firewalls deliver deep, application-layer visibility with granular insight into traffic flows. Policy Optimizer, a feature on our Next-Generation Firewalls available as part of the PAN-OS® 9.0 software release, gives deep visibility into applications to help you prioritize rule migration, identify rules that allow unused or overprovisioned applications, and analyze rule usage characteristics.

Additionally, Cortex Data Lake collects telemetry from the network via Next-Generation Firewall appliances; the cloud via VM-Series Virtualized Next-Generation Firewalls; and endpoints via Traps™, an endpoint protection and response agent. With this data centralized, Cortex XDR taps into Cortex Data Lake to validate established interaction and provide details around that interaction to help refine the use of communication and understanding of the flow.



3. Architect a Zero Trust Network


Traditionally, the first step of any network design is to architect it. Individuals get “reference architectures” for the network and must work to make them usable for their business. In the Zero Trust journey, architecting the network is the third step. Further, Zero Trust networks are bespoke, not some universal design. With your protect surface defined and flows mapped, the Zero Trust architecture will become apparent.

The architectural elements begin with deploying a Next-Generation Firewall as a [segmentation gateway](#) to enforce granular Layer 7 access as a microperimeter around the protect surface. With this architecture, each packet that accesses a resource inside the protect surface will pass through a Next-Generation Firewall so Layer 7 policy can be enforced, simultaneously controlling and inspecting access.

There is a significant misunderstanding that Zero Trust is only about access control, but least-privileged access control is only one facet of it. Another is the inspection and logging of every single packet, all the way through Layer 7, to determine if packets are clean. This is done by inspecting all network traffic for malicious content with multiple integrated security services, including intrusion prevention systems (IPS), sandboxing, URL filtering, DNS security, and data loss prevention (DLP) capabilities.

A Tailored Suit

Think about how custom clothing is made. The designer first measures you, then creates a pattern, and only begins sewing after that. Zero Trust mirrors this process. You can't architect an effective, secure network without first understanding what needs to be protected and how those systems work.



Palo Alto Networks Next-Generation Firewalls take advantage of our powerful, exclusive App-ID™, User-ID™, and Content-ID™ technologies to define authoritative Layer 7 policy controls and prevent compromise of protect surfaces. Because these segmentation gateways are offered in both physical and virtual form factors, this architectural model can work everywhere you may have a protect surface, whether in on- or off-premises physical data centers, or private, public, or hybrid cloud environments.

Endpoint security, such as Traps, can prevent compromise of the protect surface by known and unknown threats, whether from malware, fileless attacks, or exploits. Secure access offerings, such as Prisma™ Access by Palo Alto Networks, cloud service, extend the policy of each microperimeter down to the endpoints attempting to access protect surface resources.

The Security Operating Platform® delivers telemetry from all core Palo Alto Networks technologies to Cortex Data Lake, enabling machine learning and automation via Cortex XDR for improvement of policy in later stages of your deployment.

The architecture would still be incomplete without important third-party offerings. Palo Alto Networks integrates with multiple multi-factor authentication (MFA) providers to add fidelity to User-ID. To round out and simplify Zero Trust architectures, our powerful API provides deep integrations with more than 250 third-party partners, including anti-spam/anti-phishing technologies, DLP systems, software-defined wide area networks (SD-WAN), and wireless offerings.

4. Create the Zero Trust Policy


Once you've architected your Zero Trust network, you need to [create the supporting Zero Trust policies following the Kipling Method](#), answering the who, what, when, where, why, and how of your network and policies. For one resource to talk to another, a specific rule must whitelist that traffic. The Kipling Method of creating policy enables Layer 7 policy for granular enforcement so that only known allowed traffic or legitimate application communication is allowed in your network. This process significantly reduces the attack surface while also reducing the number of port-based firewall rules enforced by traditional network firewalls. With the Kipling Method, you can easily write policies by answering:

- **Who** should be accessing a resource? This defines the "asserted identity."
- **What** application is the asserted identity of the packet using to access a resource inside the protect surface?
- **When** is the asserted identity trying to access the resource?
- **Where** is the packet destination? A packet's destination is often automatically pulled from other systems that manage assets in an environment, such as from a load-balanced server via a virtual IP.
- **Why** is this packet trying to access this resource within the protect surface? This relates to data classification, where metadata automatically ingested from data classification tools helps make your policy more granular.
- **How** is the asserted identity of a packet accessing the protect surface via a specific application?

No Traffic Should Be "Unknown"

In Zero Trust, there is no "unknown traffic." If you don't know what the traffic is, it shouldn't be allowed to access the protect surface. Unknown traffic validates that the trust model is broken and needs to be repaired. Such traffic must be made known by defining a Kipling Method tuple to determine whether the traffic should be allowed or not.

Palo Alto Networks Next-Generation Firewalls deliver a default deny out of the box, and with a Zero Trust architecture, you'll have very few deny rules. For example, you might have an Active Directory group embedded into a custom User-ID for a single individual who should not have access the protect surface. Instead of making changes to the Active Directory, you can simply create a deny rule to that protect surface using the individual's specific domain credential with a User-ID.



Who	What	When	Where	Why	How	Action
User-ID	App-ID	Time	System Object	Classification	Content-ID	—
Sales	Salesforce	Working hours	US	Toxic	SFDC_CID	Allow
Epic_Users	Epic	Any	Epic_Svr	Toxic	Epic_CID	Allow

To simplify the process, you should create policies primarily on your segmentation gateways' centralized management tool. Palo Alto Networks Panorama™ provides this functionality and is where the Kipling Method is applied.

Palo Alto Networks powerful Next-Generation Firewall technology and unique features let you write policies that are easy to understand and maintain while providing maximum security transparent to your end users. User-ID technology helps define the who, Application-ID helps define the what, and Content-ID helps define the how, all of which is enforced throughout your deployment, including by WildFire® malware prevention service, as well as Threat Prevention, URL Filtering, and DNS Security services. PAN-OS 9.0 delivers enhanced policy creation capability, notably through Policy Optimizer, which continuously helps you understand how to increase the fidelity of your Zero Trust policy. Additionally, you can create policies for Prisma SaaS based on how software-as-a-service applications are accessed.


Focus on the Right Traffic

It's common for companies to look at the wrong things. For example, many focus on assessing traffic that the system has denied. Although this might provide some insight, you should focus on the allowed traffic, instead. Data breaches happen inside Allow rules. Attackers search for poorly configured technologies and misconfigured resources that let them pass malicious traffic through wide-open rules.

5. Monitor and Maintain the Network

The last step in this iterative process is to monitor and maintain your network. This means continuously looking at all internal and external logs through Layer 7 and focusing on the operational aspects of Zero Trust. Inspecting and logging all traffic on your network is a pivotal facet of Zero Trust.

It's important to send the system as much telemetry as possible about your environment. This data will give you new insights into how to improve your Zero Trust network over time. The more your network is attacked, the stronger it will become, with greater insight into making policies more secure. Additional data gives you insight into the protect surface, such as what you should include in it and the interdependencies of data within it, that can inform architectural tweaks to further enhance your security.



All telemetry generated by Palo Alto Networks endpoint, network, and cloud security technologies is sent to Cortex Data Lake, where the data is stitched together to enable machine learning and analytics.

Next-Generation Firewall and VM-Series data is consolidated into a singular view under Panorama, which raises an alert when a malicious or suspicious occurrence should be investigated. AutoFocus™ contextual threat intelligence service, enables this investigation with a combination of machine intelligence from WildFire and human intelligence provided by Palo Alto Networks Unit 42 threat research team, resulting in policy improvement and a more refined protect surface. The MineMeld™ engine within AutoFocus can aggregate, enforce, and share threat intelligence from third-party sources, providing further context for improved Zero Trust policy. MineMeld can seamlessly integrate with your Next-Generation Firewall inside or outside your Palo Alto Networks deployment.

Prisma Cloud provides public cloud security and compliance monitoring, scanning all audit and flow logs across multi-cloud environments for root user and overly permissive administrator activities. Prisma Cloud builds deep contextual understanding of your cloud environment, allowing detection of user anomalies—based on activity and location—that could signal compromised credentials, brute force attacks, and other suspicious activities. Prisma Cloud also correlates threat intelligence data to provide visibility into suspicious IPs and host vulnerabilities across your resources, which can quickly be isolated to avoid additional exposure. This data provides insight that allows you to fine-tune Zero Trust privileges.

Cortex XDR takes advantage of Cortex Data Lake to create profiles of users and devices, acting as a baseline of normal use. This allows the behavioral analytics engine to detect threats based on anomalies targeting your protect surface.

In evaluating current or additional protect surface policies, Cortex XDR allows you to search the telemetry within Cortex Data Lake for communication and interactions between entities. You can also analyze the telemetry to prove the condition or get valuable insight into how your policy should be modified. In rare instances, the search can identify an unknown threat vector not factored into the protect surface. Cortex XDR will then facilitate a deep investigation of the newfound threat so you can uncover what occurred and react accordingly.


Get Started on Your Zero Trust Journey

You should do three things to get started on your Zero Trust journey:

- **Build a Zero Trust center for excellence.** It's important to create a cross-functional working group of all of the teams that will be working on your Zero Trust architecture as you start your journey. This includes IT and cybersecurity teams that will be doing the actual deployment as well as business leaders who can help define the business objectives necessary to create a successful, powerful architecture.
- **Do a Zero Trust workshop.** To get everyone on the same page, your Zero Trust center of excellence should engage in a workshop to ensure everyone understands the basic concepts of Zero Trust, the business objectives of your environment, and the starting protect surface. Plan the prototype Zero Trust network during this workshop so your architects and engineers can move into a more formal design phase.
- **Start with something low-risk.** Rather than the “crown jewels” of your organization, start by deploying Zero Trust in a low-risk environment where your implementation teams can get hands-on experience and confidence in building out your Zero Trust network.

Zero Trust is a powerful prevention strategy when implemented across your entire environment—the network, endpoint, and cloud. This five-step methodology relies on using a Next-Generation Firewall as a segmentation gateway, as part of a tightly integrated platform, to simplify protection of your sensitive data and critical assets. Using every part of the Security Operating Platform, and taking advantage of the platform's deep integrations with technology partners, you can define a significant portion of your Zero Trust network for a seamless Zero Trust environment.

Palo Alto Networks offers Professional Services to help you execute on your Zero Trust strategy. Our expert consultants will work with your team to step through the five steps of Zero Trust implementation to maximize protection for your most valuable assets. With our comprehensive approach, Zero Trust becomes actionable, simple to deploy, and a powerful business enabler.



Build a Zero Trust center of excellence



Get engaged with a Zero Trust workshop



Start low-risk vs. the crown jewels

Figure 2: Starting your Zero Trust journey



3000 Tannery Way
Santa Clara, CA 95054
Main: +1.408.753.4000
Sales: +1.866.320.4788
Support: +1.866.898.9087
www.paloaltonetworks.com

© 2019 Palo Alto Networks, Inc. Palo Alto Networks is a registered trademark of Palo Alto Networks. A list of our trademarks can be found at <https://www.paloaltonetworks.com/company/trademarks.html>. All other marks mentioned herein may be trademarks of their respective companies. simplify-zero-trust-implementation-with-a-five-step-methodology-wp-082919