A Practical Guide for Federal Agencies Adopting Zero Trust in the SOC

Mature Your Zero Trust Strategy with Visibility into Critical Assets and Continuous Monitoring

paloalto NETWORKS | CORTEX BY PALO ALTO NETWORKS

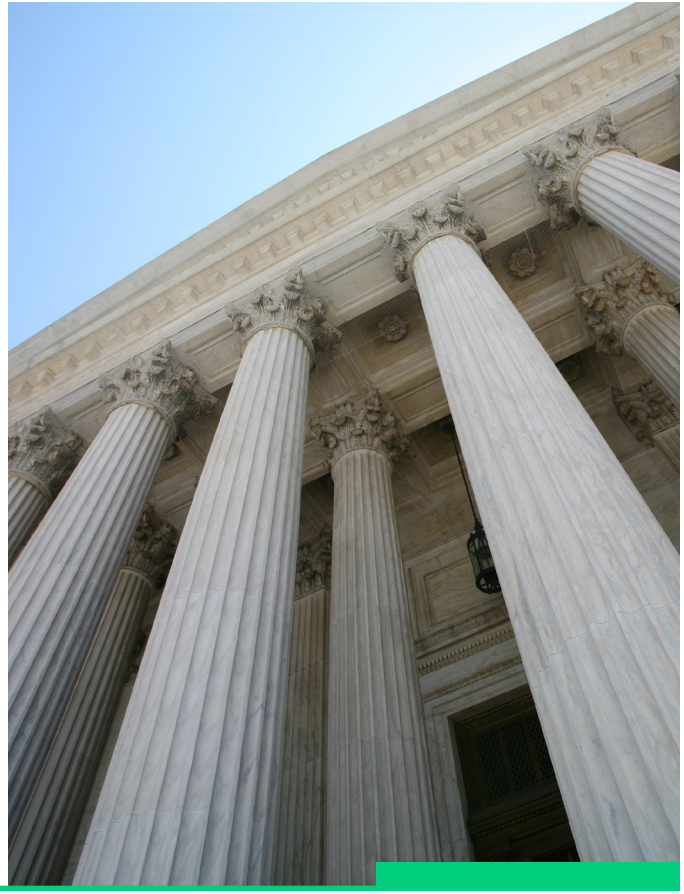

Table of Contents

Introduction

The purview of the SOC has traditionally been focused on the perimeter, yet perimeter-centric strategies for security are no longer sufficient. With the advent of embedded systems, IoT, and (nearly) ubiquitous wireless connectivity, our collective attack surface has grown beyond traditional defense techniques. The location of infrastructure and systems extends beyond the traditional network perimeter to the supply chain, cloud, remote workers, every connected device or endpoint.

Each of these requires visibility and control over activity and behavior to prevent compromises and breaches. As this highly distributed working environment becomes the new normal, we must adapt the controls employed to ensure a consistent and reliable cybersecurity posture. As such, our trust decisions need to be reevaluated to secure modern enterprise ecosystems.

The concept of Zero Trust has been around for a while and was introduced by then Forrester Research Analyst John Kindervag as a way of addressing threats that were circumventing traditional security models. This new model completely changed the way we think about IT security by assuming previously "trusted" infrastructure was in fact compromised and potentially hostile. Organizations quickly learned that they must shift from "Trust, but verify" to a "Assume breach, always verify" model.

Generally speaking, Zero Trust relies on strict "inline" continuous verification and validation of entities and processes attempting to access network resources and the removal of implicit permit in all aspects of the enterprise. Organizations can no longer assume users will be the primary attack vector and must shift this definition of what constitutes an entity to include endpoints, servers, cloud-based resources, containerized systems, and applications or services. If the entity has the means to interact with the network or gain access to a resource, it should be monitored and secured. The primary goal is to prevent successful breaches or corruption of data, applications, and business-critical systems.

USERS ARE EVERYWHERE

76% of employees want to be hybrid, even after the pandemic.[1]

Government systems have significant challenges in this defense paradigm due to the highly federated nature of identity in government enterprise. Visibility across multiple entities improves efficacy and efficiency in detection due to the additional telemetry sources and correlation thereof. Silos between agencies limit

1. *The State of Hybrid Workforce Security*, Palo Alto Networks, August 25, 2021.

that "crowdsourcing" of information—even to the detriment of the entire ecosystem—and movement of users between collaborating agencies complicates forensics and attribution.

The principles in Zero Trust are designed to reduce exposure and unauthorized access from across the threat landscape. They've been thoughtfully developed to address the security of critical applications and sensitive data across an organization. These principles can easily become a part of any security strategy and serve as continuously evolving building blocks of a security posture. Some of them include:

- **Multifactor authentication (MFA)**: A security protocol that requires individuals to be authenticated with more than one required security procedure. Typically, this is some combination of things one knows (e.g., passwords or a PIN), things one has (e.g., fob, badge), and/or physical markers (e.g., biometrics, voice recognition, or fingerprints).
- **Policy of least privilege**: A policy in which entities are given the minimum amount of access they need to carry out their jobs. This helps reduce pathways and exposure to malware, attackers, and the chances of data exfiltration. For example, an application should only be able to perform certain system functions that are consistent with its established baseline of operation.
- **Microsegmentation**: A network is divided into separate segments or "secure zones" in data centers or in cloud deployments that require different access credentials to help isolate users, devices, and even workloads. This also helps limit lateral (or east-west) movement in internal networks if breached. Entity behavior should be baselined to determine what normal looks like so deviations can be detected, alerted on, and responded to.

Security Challenges for Government Organizations

Renewed Focus on Cyber Hygiene and Implementation of Cybersecurity Technologies

These many challenges require new techniques and technologies for government SOCs to stay ahead of the adversary. The May 2021 Executive Order on Improving the Nation's Cybersecurity highlights many of these, including implementing Zero Trust architectures (ZTAs),

What is Zero Trust? It's a strategic approach to cybersecurity that secures an organization by eliminating implicit trust and continuously validating every stage of digital interaction.

requiring endpoint detection and response (EDR) across the entire federal government, requiring secure cloud capabilities (including ZTA in the cloud), and calling for enhanced visibility and monitoring of all public-facing, internet-connected government assets. Continuous verification that services and tools are deployed safely fulfills an important aspect of this executive order, and attack surface management enables that mission.

Recent Binding Operational Directive Requires Enhanced Visibility and Monitoring

Observers will also note the federal government has become increasingly proactive in eliminating attack vectors before an attack can be executed. In November 2021, the Department of Homeland Security/Cybersecurity and Infrastructure Security Agency (DHS/CISA) published Binding Operational Directive (BOD) 22-01: Reducing

the Significant Risk of Known Exploited Vulnerabilities, establishing requirements for United States government agencies to remediate a broad range of known, critical vulnerabilities.

As of late January 2022, the Known Exploited Vulnerabilities Catalog accompanying BOD 22-01 contained over 500 individual Common Vulnerabilities and Exposures (CVEs), at least 165 of which involve internet-facing products and services. These numbers continue to increase since BOD 22-01's initial publication, as CISA adds new CVEs to the KEV catalog.

Novel Threats Posed by New Vulnerabilities and Bad Actors Continue to Proliferate

Protecting government networks is of the utmost importance to reduce the risk of data breaches that could compromise taxpayer and government employee personally identifiable information (PII) records or sensitive national security information.

However, recent high-profile CVEs, from SolarWinds to Microsoft Exchange and Apache Log4j, have demonstrated that attackers are always searching for ways to target federal networks, and the frequency at which new CVEs are released is ever-increasing. SOCs are overwhelmed with events, alerts, and incidents to investigate, not to mention remaining on call for emergencies and incident response. Ensuring system availability and reliability is also crucial to the function of government, to provide and protect services, and further US interests abroad. Federal SOCs must continue to meet this challenge head-on, leveraging new capabilities to remain ahead of bad actors looking to damage networks, steal information, and dent the credibility of the United States.

COVID-19 Pandemic Accelerates Shift to Telework

The onset of the COVID-19 pandemic in 2020 dramatically accelerated the plans of many agencies to accommodate additional telework opportunities for the federal workforce. Adding to an already complicated, geographically distributed footprint, expanding telework also expands network boundaries to include employees' personal or government-furnished devices, connecting via potentially insecure networks.

One of the results has been a dramatic increase in the number of remote access exposures (such as Remote Desktop Protocol or telnet) that could be exploited by a bad actor to gain a foothold on a government network. Palo Alto Networks 2022 Cortex Xpanse Attack Surface Threat Report found that Remote Desktop Protocol (RDP) accounted for one-third of all security issues exposed on the public internet. SOC teams will need to ensure they maintain visibility and management of their network perimeter to successfully detect and defend against attacks utilizing this vector.

Agencies Must Transform to Keep Pace

In response to these many threats and challenges, government SOCs can leverage key technologies to inform their security operations strategy, improve government cybersecurity, supply chain, and critical infrastructure to establish centralized visibility and operational control over federal information technology These technologies include:

- **Integrated XDR, logging, SOAR, and hunting** to prevent anomalous behavior at the endpoint, correlate data across the enterprise, automate response, and support government-wide hunt and incident response teams
- **Internet Operations Management and national vulnerability and incident remediation** to identify and monitor the government's entire internet-facing attack surface and remediate vulnerable software running in the infrastructure

What Is a Zero Trust Approach?

Zero Trust is a strategic approach to cybersecurity that works to eliminate implicit trust and verify every stage of a digital transaction. This means controlling and monitoring all entities (e.g., users, managed devices, IoT devices, workloads) and their access to all resources, especially resources related to critical infrastructure and sensitive data. It is important to remember that an entity is not restricted to the items above but includes any resource that may attempt to access another resource.

Building the Zero Trust Enterprise

Users, Applications, and Infrastructure

Zero Trust for users starts with verifying and validating identity, followed by:

- Verifying the integrity of the user's device.
- Blocking or allowing access based on the verified identity.
- Ensuring verified identity everywhere in your security architecture where you block and allow access.

When applying Zero Trust to applications, it requires a similar approach:

- Address the identity of the entity trying to access an organizational resource.
- Validate access between applications and workloads, and validate the transaction to ensure the content related to applications is not malicious, regardless of where these applications are in private cloud, public cloud, etc.
- Likewise, verify the flow of an application transaction to ensure rogue steps or processes have not been injected that could alter the associated outputs.

	Identity	Device/Workload	Access	Transaction
Zero Trust for Users	Validate users with strong authentication	Verify user device integrity	Enforce least-privileged user access to data and applications	Scan all content for malicious activity and data theft
Zero Trust for Applications	Validate developers, DevOps, and admins with strong authentication	Verify workload integrity	Enforce least-privileged access for workloads accessing other workloads	Scan all content for malicious activity and data theft
Zero Trust for Infrastructure	Validate all users with access to infrastructure	Identify all devices, including IoT	Least-privileged access segmentation for native and third-party infrastructure	Scan all content within the infrastructure for malicious activity and data theft

Figure 1: Building a comprehensive Zero Trust approach: Users, Applications, and Infrastructure

Apply this same rigor when looking at infrastructure by validating the identity of entities connecting to the infrastructure:

- IoT presents a significant security risk due to a lack of built-in security, so validating identity can secure all devices including IoT.
- Continuously discover and monitor all internet-connected assets and changes to existing assets to ensure full visibility of exposure risks.
- Apply least-privileged access and segmentation, and monitor transactions of native and third-party supply chain infrastructure to include third-party software libraries tied into custom business applications.
- Scan all content within the infrastructure for malicious activity and data theft.

A Holistic Approach to Zero Trust: The Role of the SOC

Zero Trust is an ongoing process that requires continuous evolution and refinement as each organization's requirements change and subsequent technology shifts occur—especially during digital transformation initiatives. As such, continuous monitoring of the threat landscape should be a core requirement in any Zero Trust journey. Monitoring needs to go beyond any single security tool to ensure broadened visibility across the entire attack surface. Routine gap analysis must become part of security methodologies to ensure holes not previously identified can be brought to light and remediated, as required. This makes the role of the SOC critical in the continued audit and maintenance of any Zero Trust security posture.

Specifically, the SOC plays a key role in reassessing trust:

- Continuously verifying Zero Trust policies
- Identifying gaps in your Zero Trust posture
- Limiting attack impact through automated enforcement
- Speeding up investigation with automated threat data collection
- Continuously discovering critical assets

APPS ARE EVERYWHERE

80% of organizations have a hybrid cloud strategy,[2] and the average organization uses 110 SaaS apps.[3]

2. *2021 State of the Cloud Report*, Flexera, March 2021.
3. "Average number of SaaS apps used by organizations worldwide 2015-2020," Statista, February 16, 2022.

Case in point, an organization might implement MFA to correctly identify entities and grant access to applications. The SecOps team can analyze an entity's activity with machine learning, behavioral analytics, and human insights to detect insider abuse and disable a rogue entity's access to mitigate damage. Even with a mature Zero Trust implementation that secures entity access and permissions, organizations still need a SOC to reassess trusts through threat detection, response, automation, and risk management.

As a first step, organizations should establish controls that address Zero Trust concepts across all entities to include users, applications, and infrastructure. The SOC has the best vantage point to ingest a broad set of security telemetry, perform continuous monitoring, and to validate and verify Zero Trust controls.

SOC Transformation: A Critical Step in Modern Zero Trust

In order to implement Zero Trust, SOC teams have to be aggressive in tuning their detection alert settings as failure to do so results in higher alert volumes. This tuning process must be continuous, and it has to be informed by ever-changing organizational priorities not only driven by leadership but informed by the analysts subsequently tasked to carry out these actions. Multiply this across the average 30 or more tools that the SOC uses, and it becomes inevitable there will be a deluge of low-fidelity alerts and noise cluttering analysts' dashboards while providing minimal operational value.

For instance, while we have tools that can create alerts, they require analysts to either confirm if the alert is legitimate or to close it as a false positive. As a result, SOC analysts can spend inordinate amounts of time investigating and validating a *single* alert. Plus, they may end up using numerous other, and often siloed, tools just to gather enough information to decide if an alert should be escalated. Continuously tuning alerting and monitoring processes ensures SOC analysts can focus on what matters.

An over-abundance of alerts that lack useful context can result in even more time to chase down related data and sift through logs, while hoping to correlate alerts which may or may not be related. Combine scattered context data, tool sprawl,

Figure 2: Traditional approaches to security aren't working

swivel-chair operation across various networks, and even gaps in security staffing, and it becomes clear that organizations need a better way forward to defend and protect their critical infrastructures.

Modern security threats are also evolving at a faster pace than security technologies. While well-funded threat actors are investing in new tools like machine learning, automation, and artificial intelligence, SOCs built around legacy security information and event management (SIEM) fail to provide a flexible and scalable solution that keeps pace with digital transformation, cloud initiatives, *and* advanced attack campaigns. Security analysts struggle to identify, manage, and remediate critical threats when faced with overwhelming challenges such as noisy false positives, event storage (volume and cost), poor investigation workflows, the adoption of hybrid and multicloud architectures, and the proliferation of devices and endpoints.

Another challenge impacting SOCs is the use of playbooks for incident investigation, triage, response, and remediation that are not standardized across departments or teams. This results in unpredictable documentation, unsustainable response actions, and processes or procedures that aren't trainable or scalable. This can introduce additional challenges related to timely and repeatable response activities, which can increase the mean time to remediate (MTTR) for a given incident or incident type.

Issues from legacy SOC environments can include:

- Lack of visibility and context
- Increased complexity of investigations
- Alert fatigue and noise from a high volume of low-fidelity alerts generated by security controls
- Lack of interoperability of systems
- Lack of automation and orchestration
- Inability to collect, process, and contextualize threat intelligence data

A Way Forward: Embracing AI, Automation, and Orchestration

When embarking on a Zero Trust journey, an organization first needs to define a unified security policy. This typically starts with identifying critical assets and deploying a Zero Trust architecture with strict, least-access policies across users, applications, and infrastructure. After all, you can't secure what you don't know you have, much like you can't secure processes that aren't documented and understood across work centers.

Automate Workflows

Security leaders must consider whether a tool requires a human to configure or run. Is an expert needed to interpret or triage the outputs? Are people needed to test things? Security leaders can identify repeatable, low-level tasks that can work with human decision-making to help accelerate incident investigations.

While advancements in machine learning and artificial intelligence hold great promise, retaining the human element for knowledge transfer in either direction is imperative to achieve optimal outcomes for a smooth SOC transformation. As automation capabilities begin to mature, humans can and should own a smaller and smaller piece of workflows.

With too many manual processes involved in security operations and incident response (IR), including numerous threat feeds to monitor, investing in automation capabilities such as those in a security orchestration, automation, and response (SOAR) solution can help orchestrate actions across the product stack for faster and more scalable IR.

One-to-Five-Year Prediction on Automation Takeaways

New SOC operations can start using automation from day one, while more established organizations will have to retool and figure out where the move to automation can begin. This is a good three-year goal for an established organization: to move 50% of SOC work into the hands of automation. By year five, most SOC teams can be closer to 75% of activities automated yet continue to rely on human engineers for other activities like threat hunting.

Augment People with Machine Learning-Driven Intelligence

A key component to rearchitecting your SOC for Zero Trust is to ensure that security teams are leveraging machine learning to its full potential to augment and complement humans in security. Advanced analytics and AI can significantly reduce the amount of time that teams spend processing massive amounts of data in the enterprise to come up with critical security insights. By automatically detecting anomalous patterns across multiple data sources and automatically providing alerts with context, machine learning today can deliver on its promise of speeding up investigations and removing blind spots in the enterprise.

This works by training machine learning models, using them to detect patterns among and across the data, and then testing and refining the processes. ML techniques can gather, integrate, analyze, and interrogate the data to reduce the amount of time and knowledge needed for a human to perform these tasks. This also minimizes the challenge for a SOC team trying to find threat context and evidence embedded in data across multiple layers of security.

At a high level, machine learning techniques can:

- **Integrate**: Enable the data to tell a story about what is happening.
- **Analyze**: Extract insights about the problem space and make predictions.
- **Automate**: Accelerate human decision-making, enrich incident data, and automate system-level action, workflows, and decision-making.

"We treat all the use cases the same in the most extreme way that we can. We don't give anyone any discounts just because we can assume something about them based on where they are, who they are, what they're trying to do, and so on. It turns out that that approach leads to a much simpler infrastructure because all of a sudden we don't need to buy different equipment or different solutions, different technology for securing users, depending on the situation or securing applications based on the situation.

We can use one architecture, one system, one solution, one technology to secure all users all the time, wherever they are, whatever it is they're trying to do because we're going to run them through the same security checks and the same is true for securing applications and so on."

—Nir Zuk, Co-Founder & CTO, Palo Alto Networks

Achieve Comprehensive Zero Trust Faster with the Cortex Suite of Products

You can begin or accelerate your SOC journey by deploying the Cortex suite of products: Cortex XDR, Cortex XSOAR, and Cortex Xpanse, which seamlessly work together as a force multiplier across your security operations.

Better together, SOC teams can experience immediate high-level advantages:

Cortex XDR: Helps keep your organization safe from attack by delivering leading endpoint protection and enterprise-wide threat detection and response across network, cloud, endpoint, and virtually any data source. Patented behavioral and machine learning-based analytics pinpoint evasive threats and provide the intelligence you need to respond before a breach can occur.

Cortex XSOAR: Provide a single platform for SOC teams to manage all their incidents and threat intelligence feeds. With 900+ prebuilt integrations for security tools used in the SOC and thousands of automated workflow scripts or playbooks, XSOAR enables SOC teams to be

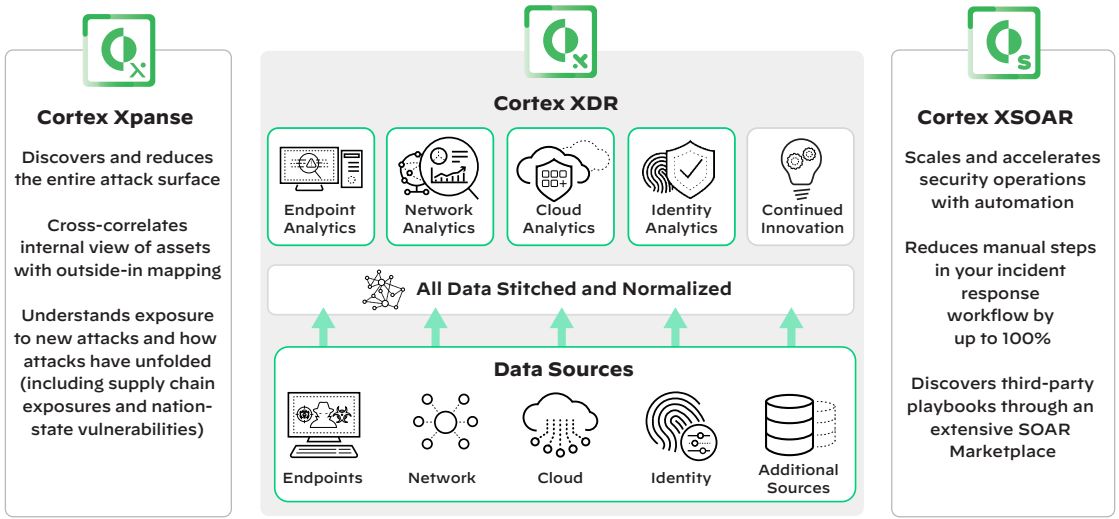

Figure 3: The Cortex suite of products is a force multiplier

Cortex XDR is part of Palo Alto Networks Government Cloud Services. As a cloud-delivered service that integrates data from your existing security sensors, Cortex XDR has achieved a Federal Risk and Authorization Management Program (FedRAMP) Moderate Authorization, meeting federal requirements for additional levels of assurance. Learn more.

as aggressive as they need to be in their alert settings to implement Zero Trust without worrying about the flow of alerts impacting analyst workload. As a result of using XSOAR to automate their SOC operational processes, a US electric utility company was able to reduce the number of cases by 30% within the first month of operations.

Cortex Xpanse: A complete, accurate, and up-to-date inventory of an organization's global, internet-facing cloud assets and misconfigurations to continuously discover, evaluate, and mitigate risks on an external attack surface and evaluate supplier risk or assess the security of M&A targets.

While each standalone product brings its own unique features and benefits, when combined, the positive results increase exponentially. These three products help lower the risk and impact from breaches with a comprehensive product suite for security operations, empowering enterprises with best-in-class detection, investigation, automation, and response capabilities, bar none.

With end-to-end native integration and interoperability, SOC teams can close the loop on threats with continual synergies across the Cortex ecosystem. All three products work in concert to monitor the threat landscape and provide the most robust detection, response, and investigation capabilities:

- Cortex XDR and Cortex Xpanse provide ultimate visibility and detections across the internet attack surface, endpoints, cloud, and network, including remote workers.
- Cortex XDR can leverage Cortex XSOAR to automate malware investigation and response.
- Cortex Xpanse works together with Cortex XSOAR to automatically enrich incidents using Xpanse asset information and to automate the remediation of newly discovered assets.
- Cortex XSOAR leverages Cortex XDR and Cortex Xpanse to provide high-fidelity detections and alerts to drive automated incident response workflows.

What's Next? Future-Forward with XSIAM

While Cortex products address key SOC requirements for visibility, protection, and automation, most organizations still depend on SIEM as a core component of SecOps. But SIEM products have failed to deliver on the promise of effective, centralized threat detection and response, burdening analysts with endless alerts and manual processes. Security teams need a central platform that incorporates and automates multiple security functions into a single foundational solution with visibility into enterprise-wide security data.

Extended security intelligence and automation management (XSIAM) is purpose-built to address this need, harnessing the power of AI-driven automation to radically improve security outcomes and transform the manual SecOps model. By building an intelligent data foundation and automating unified SOC functions, XSIAM accelerates response, outpaces threats, and dramatically streamlines analyst activities.

XSIAM is designed to be the center of SOC activity, replacing SIEM and specialty products by unifying broad functionality into a holistic and automated solution. XSIAM is revolutionary in the way it operates, using intelligent automation to transform the analyst-driven model of today's security products. With XSIAM, organizations can consolidate security data and tools, automate activities, and eliminate security gaps, delivering dramatically better protection and streamlined operations.

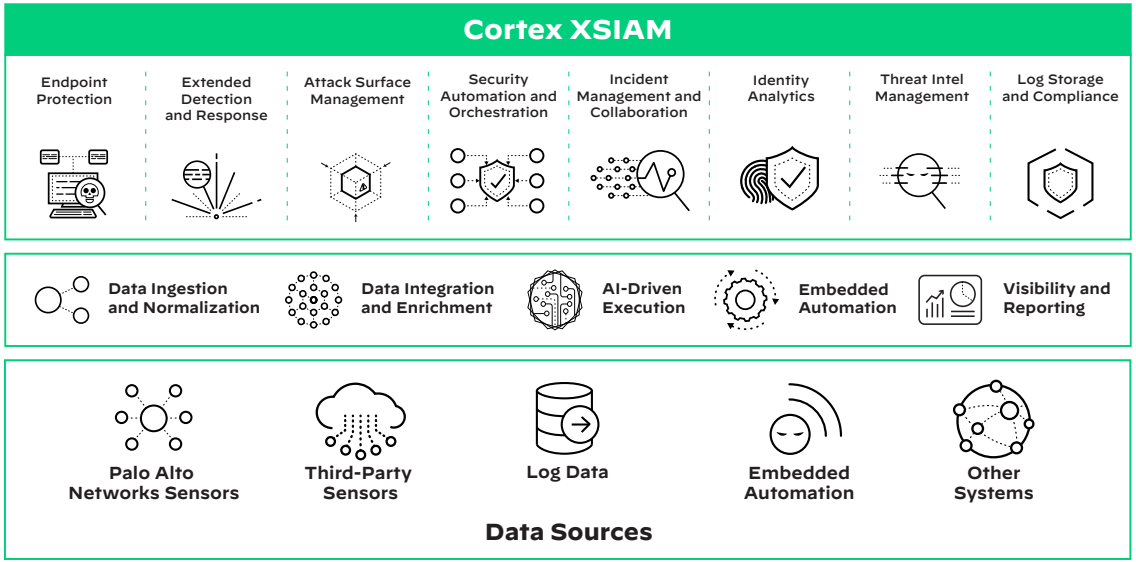

Figure 4: XSIAM is the AI-driven platform for the modern SOC

Powered and Protected by Cortex

Government agencies face unique security challenges, including the need for tools that provide visibility and enhanced operational efficiency over large, federated, geographically distributed networks. Furthermore, effective command and control (C2) is difficult to achieve for organizational structures in which multiple teams share responsibility to secure and defend agency infrastructure. Government agencies are also a frequent target of APT groups, nation-states, and other bad actors, given these agencies' vast stores of sensitive national security information and employees' and citizens' data. A breach may impact national security, leak PII, or negatively affect the availability of crucial government services.

The Cortex portfolio offers an end-to-end security solution that ensures every step of the security process is covered.

Palo Alto Networks is committed to bringing the newest and most advanced and integrated security solutions to market. We invite you to take a look at our solutions, reach out, and talk to us. We're here to help you learn more, do more, and secure more.

Get to Know Cortex

Visit our product pages for more information:

- Cortex Xpanse
- Cortex XSOAR
- Cortex XDR
- Cortex XSIAM
- Cortex portfolio page

More Zero Trust Resources

Digital transformation is accelerating with key shifts such as the expanding hybrid workforce and continued migration of applications and data to the cloud. As we make this transformation, information security teams have the opportunity to adopt a modern Zero Trust approach that fits these significant shifts. Enjoy our resources to learn more:

- Download our Architecting the Zero Trust Enterprise whitepaper
- Read our blog, Building the Zero Trust Enterprise: The Role of the SOC.
- Contact the Palo Alto Networks federal team for additional information, or visit the Palo Alto Networks federal website.

The full Cortex product suite is available on the DHS CISA CDM program Approved Products List (APL). Learn how Cortex XDR, XSOAR, and Xpanse can improve outcomes and efficiencies for your SOC. Download the DHS Approved Products List.

paloalto NETWORKS

3000 Tannery Way
Santa Clara, CA 95054

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

www.paloaltonetworks.com

cortex_ebook_practical-guide-to-adopting-zero-trust_121922