

Deloitte.

Zero Trust solutions
AI-Native Security
Operations Platform™



Introduction

Transform security operations through AI, automation, and orchestration

Integrating Artificial Intelligence (AI), automation, and orchestration in a Security Operations Center (SOC) can radically transform security operations, evolving them into a more proactive, efficient, and effective entity capable of addressing the modern threat landscape with agility and precision. This approach not only fortifies the security posture of organizations but also requires resource utilization and enhances strategic decision-making.

Through the strategic application of AI, analytics, behavioral analysis, and threat detection are enhanced, enabling the system to carefully anticipate and identify potential cybersecurity threats and breaches. Automation fortifies security operations, facilitating automated responses to threats, meticulous alert management, and proficient patch management, determining that low-risk incidents are managed without unnecessary manual intervention. Additionally, the orchestration of processes enables the synchronization of various security tools, a coordinated incident response, and the streamlined sharing of threat intelligence among different teams and technologies. This proactive defense mechanism, supported by AI and automation, encompasses vigilant threat hunting, meticulous vulnerability management, and broad data protection, which safeguard networks from undetected threats and vulnerabilities. The technology not only enhances decision-making and strategic planning through actionable insights and AI-driven analytics but also mitigates alert fatigue and augments analyst productivity by smartly consolidating alerts, providing enriched contexts, and automating mundane tasks.

Ultimately, an AI-native SOC benefits organizations by significantly increasing operational efficiency, enhancing the security posture, requiring resource utilization, and improving accuracy by minimizing human error and enabling quicker and more adept responses to security incidents.

An enterprise-scale security operations capability is a critical function in modern, digital organizations. That said, the cost and resourcing challenges for this critical function are very real. From a technology perspective: collecting and centralizing logs, developing technical policies and threat identification use cases, funding the many necessary security tools such as Security Information and Event Management (SIEM), User and Entity Behavior Analytics (UEBA), Data Loss Prevention (DLP), Network Analytics (NA), etc. comes with significant overhead. From a resourcing perspective: sourcing, hiring, and retaining highly trained cybersecurity specialists is increasingly difficult in a world where we have a supply and demand issue and companies competing for limited resources. Add to this the rapidly shifting threat landscape and business model changes post-pandemic, and many organizations are challenged to evolve their security operations capability to meet the increasingly sophisticated threats they face.

Additionally, due to multi-dimensional threats such as ransomware, many organizations are struggling to find the applicable balance (and associated investment) between prevention, detection, and response security capabilities. Our view is that this balance is no longer relevant. To use a historical parallel, we can compare the inception and evolution of security operations to the evolution of aircraft safety. Initially, hobbyists made and launched aircraft with limited consideration of the long-term implications. Commercialization soon followed along with expansion driven by repeatable design and build capabilities. Over time the increased demand for airplanes and pilots led the airline industry to realize they had a critical need to establish rigorous and repeatable processes as well as a need for regulation to establish the operational boundaries of the industry.

The Federal Aviation Administration (FAA) was born (as well as other regulatory agencies in other countries) to provide mandatory rules that many industry players comply with. Thereafter, the processes by which planes take off and operate while airborne became highly controlled using automated computer systems. In general, the methods by which pilots interacted with air traffic control systems became rigid in order to prioritize the safety of the passengers.

However, even today humans are still responsible for creatively navigating unforeseen events. For example, pilots leverage autopilot systems to handle routine tasks -- maintaining the aircraft's course, altitude, speed, and heading. However, for critical events or in an emergency, the pilot will typically take control of the plane when automation alone is not enough ... such as takeoff, landing, and addressing unforeseen circumstances in flight.

This example underscores the point of this paper; highly standardized processes, AI decisions and automation provide an effective baseline for operations, but human innovation and creativity are still required to achieve the expected/desired result in extraordinary or dynamic situations.

The important point is that there should be a distinction between where security operations should have a tight, repeatable, and measurable process and where creative cybersecurity specialists may perform maneuvers to mitigate potential issues. This paper highlights ways to create a consistent set of core processes and allow room for creativity within the process set for organizational security operations.

Security operations basics

The foundational security operation model was built 20 years ago, and many SOCs follow a few core security operations processes:

- Vendor tool-driven alert triage using a "follow the sun" model (a type of global workflow in which issues are handled by and passed between offices in different time zones, increasing responsiveness and reducing delays) with defined handoffs.
- Traditional security operations, characterized by their organized and systematic methodologies, provide a stable and consistent framework for incident response processes. However, they often face challenges in adaptability, scalability, and responsiveness when confronting progressively sophisticated threats. The intrinsic constraints of such traditional frameworks have prompted a shift towards more cohesive, forward-looking, and flexible security models in numerous modern security operations.
- Security operations use case design and development processes resulting in analyst playbook development. The traditional playbook is typically a more general document, outlining the organization's approach and worker responsibilities when responding to a limited number of scenarios.
- While traditional security operations yield regular and systematic insights via static reporting, they frequently find it hard to keep pace with the rapid advancements and complexities in the realm of cyber threats. The reports were characterized by their inflexibility, extensive generation time, and absence of responsiveness to the continually changing threat environment.
- Intelligence is typically consumed from vendor-provided feeds and focuses on known indicators of compromise (IOCs) -- which may or may not be address the threats facing a specific organization.
- For security operations' environmental awareness, activities such as Configuration Management Database (CMDB)



verification, asset enumeration, and data flow analysis are managed, however, they are typically conducted manually and are point-in-time assessments.

- Traditionally, organizations outsourced security operations to leverage the experience, technology, and cost benefits of external providers, enabling cost reduction and continuous, required security monitoring through a "follow the sun" model. However, these providers often lack motivation to advance underlying technologies or introduce new capabilities because it may require further training, updating their team's experience, or even recruiting new personnel, many of which could negatively impact their profitability.
- Broad security operations, while sometimes seen as tedious due to their meticulous procedures, are crucial for safeguarding organizations. A broad approach to security requires people, processes, and technology to function in harmony, forming a 'triage', rather than just leaning on cutting-edge tools. An organization's maturity is reflected in its capacity to effectively blend these components. Foundational security operations are pivotal, setting the groundwork for future security strategies, emphasizing the significance of resource integration for bolstering organizational security.

Security operations complexity

Cyberattacks continue to increase, and stopping an attack is as difficult as ever. When an effective attack occurs, the target's

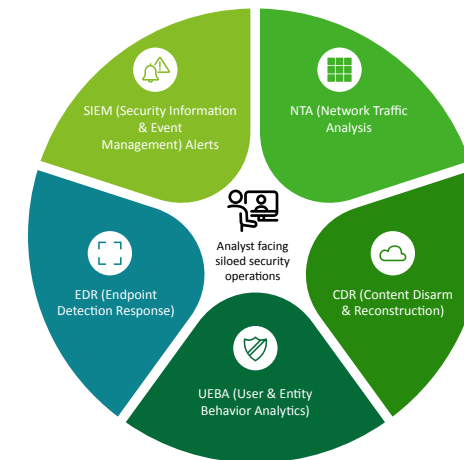
security team will typically perform forensics to figure out what happened, how the attackers got in, which systems were affected, and what data was compromised. While often effective, this is a reactive exercise, stitching together data from many different systems and conducting human led analysis across large data sets looking for patterns indicative of the phases in the cyber kill chain lifecycle. So, the question is: if the organization has the information needed to solve the puzzle after the fact, why can they not prevent the attack in the first place? The answer is the difference between manual data analysis and AI and Machine Learning (ML) driven data analytics. Many security organizations simply have too much data to manage in too many silos – coming from SIEM alerts, Network Traffic Analyzers (NTA), Endpoint Detection and Response (EDR), User and Entity Behavior Analytics (UEBA), and Cloud Detection and Response (CDR).

Analysts may gather this data post-event for retrospective analysis, but isolated data is less effective in real time. AI-native security solutions are adept at real-time inspection and dynamic adjustment of models and algorithms, enabling them to preemptively identify and address potential threats. Furthermore, these AI-driven solutions enhance the accuracy and efficiency of detecting and responding to security breaches and vulnerabilities, adjusting to emerging patterns, and continuously learning from the evolving threat landscape, thus bolstering organizational protection. The challenge lies not in the driver but in the vehicle.

It is also worth noting that many security operation playbooks and procedures were established years ago and were largely built for a different era.

Today's expanded enterprise attack surface generates much more security data from a volume perspective and that data is both more complex and more siloed than only a few years ago. Today, network, endpoint, identity, and cloud data often remain in separate systems.

Given the traditional level of effort required to build connectors and manually map data elements into a larger data model, typically only a subset of logs are going into the SIEM but a flood of alerts is coming out. As a result, security operations analysts are overloaded with alerts, many requiring manual analysis in order to triage, and often multiple levels of escalation and further analysis are needed before taking effective action. Alert overload increases the risk of missed threats and longer adversary dwells time. At the same time, security engineers struggle to integrate new data streams and create new detection rules and playbooks, while IT architects continue to integrate new technologies creating further events and alerts. The results are predictable: 'alert fatigue', slow investigation progress, and increased risk to the enterprise.



While the technology supporting cyber defense has continued to progress and security requirements are changing, the design and use of traditional SIEM has not kept pace. Many other components of the security architecture have been modernized. In Figure 1, the endpoint moved from antivirus (AV) to EDR and to Extended Detection and Response (XDR); the network moved from a "hard shell" perimeter to Zero Trust and Secure Access Service Edge (SASE); the applications moved from the data center to the cloud and microservices.

In contrast, the security operations model designed 20 years ago is still heavily dependent on a traditional SIEM.

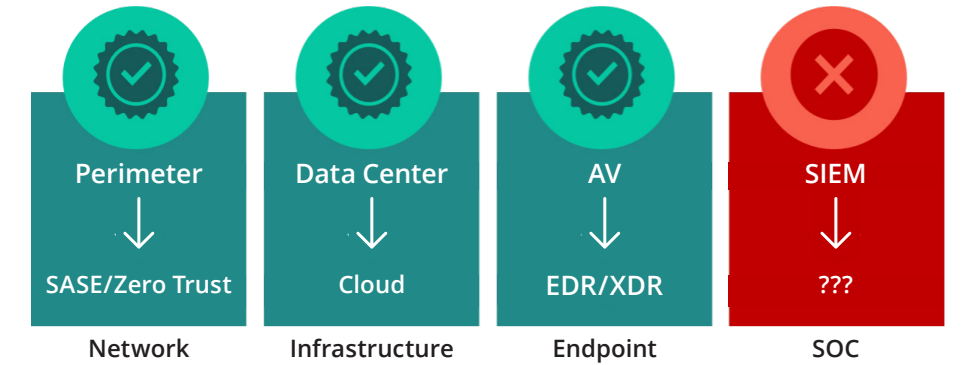


Figure 1: The Evolution for IT & Security

Currently, our clients operate under a security model primarily centered around human analysts, regardless of whether the solution is hosted on-premises or migrated to the cloud. Security operations analysts in this model meticulously scrutinize hundreds of alerts daily and manually gather contextual data for triage, dedicating a substantial amount of time to addressing false positives and exerting manual effort. However, as the volume of alerts has escalated and integrating data from an increasing number of systems has become more challenging, the efficacy of this human-led approach is compromised, showing significant strains and limitations. Instead, the modern way to scale effective security operations is with automation as the foundation and with analysts working on a small set of high-risk incidents. Just as flying a commercial airplane no longer requires constant, hands-on control by the pilot, an automation-led security operation handles the bulk of low-risk, repeated alerts, analysis tasks, and mitigations. This frees the analysts to work on urgent, high-impact incidents while the underlying platform autopilots the security to safe outcomes, learning from each activity and offering information and effective recommendations to the central control. This is our vision for the AI native security operations platform.

Transformation to the AI-Native Security Operations Platform™

The AI-Native Security Operations Platform™ should advance faster. Security operations should be agile to support dynamic software development and operations.

Security should be continuously integrated throughout the life cycle of operational processes. The value of an AI-Native Security Operations Platform™ should provide automation and efficiency with resource and time savings for the development of net new playbooks.

Ultimately, better data modeling and integration combined with automated analytics and detection, ease the burden on security engineers, who no longer need to build custom correlation rules to integrate data and detect threats. Unlike legacy security operations, the AI-Native Security Operations Platform™ leads with data science applied to massive data sets rather than human judgment and rules designed to catch yesterday's threats.

The AI-Native Security Operations Platform™ should be built on a new architecture with:

- Modern technologies utilizing AI-enabled features for comprehensive and automated integration, analysis, and triage of logs and data, facilitating rapid ingestion and correlation of new log sources. This stands in stark contrast to the slower, antiquated methods of manual data ingestion and log source integration that were once prevalent.
- Unified workflows that allow analysts to be productive.
- Embedded intelligence and automated response that can help block attacks with minimal analyst input.

- The following components are core to an AI-Native Security Operations Platform™: Artificial Intelligence (AI) powered security operations, intelligent threat modeling, agile security operations for the secure software development lifecycle (SSDL), agile data sciences and analytics, and intelligent threat hunting.

Artificial intelligence (AI) powered security operation center (SOC)

The goal of an AI-powered SOC is to proactively collect, process, and analyze internal and external intelligence data to lead intelligent threat analysis. Proactive threat intelligence should be built from various sources and AI-powered without pre-populated biases. The internal and external intelligence data should be able to correlate what may negatively impact an organization. When a breach occurs, the security operations team should be able to figure out very efficiently what happened, what systems were touched, and what data was compromised. Intelligence data should provide analysis so that the organization may promptly figure out before the damage occurs and prevent it from happening. The AI-powered SOC should recognize the need for formalizing intelligence collection using a framework and automatically take actions leveraging collected observation. Here are steps to get to the next generation security operations: a powerful way to revamp the way security data are pulled via siloed sources such as network, endpoint, identity, and cloud data.

Here are the AI-powered SOC core solutions using modern technology and creativity:

- Automated alert triage and queue management capabilities that logically process and prioritize high incidents.
- Alert volumes decrease, reducing the necessity for manual intervention by security operations personnel.
- Security operations personnel are cross-trained to take IR actions and automate remediation for low-severity incidents.

- Security operations teams detect and tune their environment on the fly via the continuous development of automated playbooks.
- Access to a dashboard empowers security operations by offering real-time visibility, facilitating proactive threat detection and efficient incident response, enabling performance monitoring, supporting data analysis, fostering collaboration, and enhancing situational awareness in the dynamic security environment.
- Investigation responses, encompassing alert triage, incident analysis, and response planning, make use of automated scoring and machine learning capabilities to forecast the accuracy of incident verdicts.
- Data collections are made continuously from network devices, firewalls, and endpoint tools to keep asset inventory management current.
- Security operations platform vendors provide intelligent and proactive security capabilities that extend service beyond simple alert triage, escalation to response, and remediation activities.

Intelligent threat modeling

One core capability for modern security operations that goes hand in hand with next-generation security operationalization is intelligent threat modeling. In fact, the value of modeling as a process extends beyond its own boundaries and should shape security operations, from detection development to threat hunting, to metrics design and implementation. Security operations have a defined process for how threat modeling could be conducted and who the expected consumers are for various models. Threat modeling frameworks are typically aligned to MITRE ATT&CK® framework as the next-generation security operationalization. The MITRE ATT&CK® framework is a broad matrix of tactics and techniques designed for threat hunters, defenders, and red teams to classify attacks, analyze attack attribution and objectives, and analyze an organization's risk.

Organizations may use the framework to analyze security gaps and prioritize mitigations based on risk.

This framework values consistent and continuous collection of integrated business contextual information from business owners themselves including impending mergers and acquisitions, external business relationships, and foreign entity relationships. Personnel assigned to threat modeling activities should be granted able focused time to make and foster these external connections. The contextual and organizational data should be paired with intelligence and threat research to make educated decisions around security posture and prioritization. Ultimately, security operations investigations and their subsequent conclusions should represent a framework language such as MITRE.

This helps align the security operation monitor and respond to the applicable threats across the phases of the cyber-attack lifecycle. These insights may analyze potential coverage gaps and be used to augment existing security operations metrics. Finally, those performing modeling should be well positioned in the security operations to have operational authority to shape detection and code development efforts, tool purchases, and operational playbooks and job aids. It is a security operations manager's decision whether to designate dedicated time and resources for these efforts or protect portions of the team's time for the effort, but in each case, the practice should be acknowledged and protected and the impact on operations should be documented and measured through formalized procedural rigor.

Agile security operations leveraging secure software development lifecycle (SSDL)

Security operations of the future need to embody the agility intrinsic to a modernized software development environment. Utilizing secure software lifecycle methodologies enables organizations to consistently monitor their security postures with near-real-time speed and automation.

An AI-Native Security Operations Platform™ aligns robust security protocols effectively into the SOC process, maintaining the integrity and safety of organizational assets while efficiently identifying genuine incidents with agile speed. This comprehensive strategy embraces the principle of continuous security, integrating the swift, adaptive nature of agile software development methodologies with ongoing visibility and cutting-edge security measures.

Utilizing the automation and orchestration features of Deloitte's Secure Software Development Lifecycle™ (SSDL) services, organizations can compile and streamline alerts from diverse sources such as endpoints, networks, firewalls, cloud flow logs, and third-party security instruments with agile efficiency. Deloitte's experience aids in aggregating a vast volume of alerts into actionable incidents, effectively integrating automated case management within the client's established ecosystems.

The integrated workflows enable the flow of information in the form of alert data through the multi-connected technology stack across the client organizations. Integrating various toolsets and their streamlined workflows greatly improves the operational team's efficiency and effectiveness. This leads to better outcomes as the operations team does not have to perform initial triage and correlation functions to identify the source of the alert. The operations team is presented with correlated and chained incident information where the attributes such as the source of the alert, its validity and existence, and its impact are represented in line with the organization-defined metrics. This enables the underlying goal of less human analysis and broader automated and AI-driven engineering to solve security challenges. An example of doing this is covered in the [SSDL white paper](#).

This approach serves as an early warning system, providing near-real-time notifications about issues, anomalies, and potential risks within the organization's environment. These alerts enable various business unit development and operation

teams to swiftly identify and address problems such as code vulnerabilities, build failures, performance bottlenecks, and security breaches. By promptly alerting operation teams to issues, these integrated workflows help reduce downtime, improve code and operations quality, and enhance overall productivity. Additionally, alerts facilitate collaboration and communication among team members, determining that problems are addressed collaboratively, and stakeholders stay informed.

The AI Native Security Operations Center platform leverages artificial intelligence and machine learning technologies to enhance cybersecurity operations' detection, response, and overall effectiveness.

The critical components of an AI-driven SOC include:

- **Data Collection and Aggregation:** Data should be collected and aggregated from various sources, i.e., implement a technology stack within the environment to lay the foundation of AI-driven analysis.
- **Data Normalization and Enrichment:** Data normalization helps standardize data, making it consistent and easier to work with. Consistent data allows security analysts to apply the same analysis and detection techniques to many data sources. Enriching security data with additional information provides context for careful threat detection and response. Contextual data can include threat intelligence feeds, known indicators of compromise (IOCs), and information about the organization's assets and network topology. This additional context helps security analysts understand the relevance and severity of security alerts.
- **Machine Learning Models:** Implementing machine learning algorithms to analyze and model security data. These models can detect anomalies, identify patterns, and predict potential threats based on historical and real-time data.
- **Alert Generation and Prioritization:** Automatically generate alerts based on machine learning models and behavioral analysis. Prioritize alerts based on risk and impact to focus resources on the many

critical threats based on severity.

- **Threat Intelligence Integration:** Integrate threat intelligence feeds to stay updated on the latest threat indicators, malware signatures, and known attack patterns. This helps in proactive threat hunting and identifying emerging threats across the organization's landscape.
- **Automation and Orchestration:** Enable automation and orchestration workflows to respond to alerts. Simple and repetitive tasks can be automated, reducing response time and allowing analysts to focus on more complex issues.

Agile data science and analytics

Unfortunately, even with world-class intelligent threat modeling, and continuous SSDL initiatives, the team may still encounter unforeseen threats that are using advanced techniques and methods leveraging AI. However, there is hope because the combination of statistical and analytical modeling coupled with the next-generation security operation processes provides fighting chances. Modern computing and cloud-enabled services unlock the power to inspect large amounts of data at high speed. Beyond faster processing of known data analysis methods, it is now financially and computationally accessible to use trained and untrained data science techniques to query massive amounts of data for faint signals that may be early warning signs of a breach. Many advanced threats are stealthy and quiet as they collect data and credentials until the time to attack.

Traditional security operations have the capability to detect early signs of cyber-attacks. They might be using various tools and techniques to monitor systems and networks for signs of threats or vulnerabilities. However, these operations often encounter challenges because their methods for detecting threats aren't as structured or less prescriptive as they could be. Traditional methods might involve monitoring a variety of signals and using more general or broad criteria to detect threats.

Scenario-driven alerting, on the other hand, involves defining specific scenarios

or patterns that signal a threat and setting up alerts to be triggered when those scenarios occur. This approach is generally more precise and can be more effective in efficiently identifying and responding to specific types of threats.

In addition, proactive threat detection requires more skilled investigation and advanced tools to act effectively. Modern computing and enabled services are powerful capabilities and security operations should rally around. Such a program is among the ways security operations may classify patterns and anomalies previously undetected by legacy security operations methods.

Security operations professionals are no longer required to look for the needle in a haystack by continuously defining what a new needle looks like. With the AI-Native Security Operations Platform™'s agile data science and analytics, security operations may now query the haystack to tell them what unusual hay looks like. For example, looking at user behavior, anomalies may reveal compromised accounts and systems faster than rule-based approaches because of how compromised accounts may stand out against a baseline of typical healthy behavior.

Next-generation security operations can gain success through adopting data science and analytics-focused processes and technologies. Determining continual learning and intelligent tech adoption, along with cross-training teams in data science and analytics, will enhance their future efficacy and efficiency in security operations.

Intelligent threat hunting

Even with thorough preparations and precautionary steps taken by security operations, unexpected and new threats are normally likely to occur. This highlights the importance of normally being prepared for unforeseen cybersecurity challenges and continuously updating and adapting security strategies to manage new kinds of threats.

For such events, intelligent threat hunting plays an increasingly important role. As a concept, threat hunting is not new; however, it is rarely viewed as a formalized process in security operations.

The threat-hunting team should be intimately familiar with threat trends, patterns, and behaviors from the next-generation threat intelligence.

Gathering and understanding information related to potential cyber threats is important. Threat intelligence data might include information about different types of cyber threats, their indicators, mechanisms, and strategies used by adversaries. The methods and rules for collecting and absorbing the data should be systematically and automatically defined within a given process. This means using software and automation to determine consistency, scalability, and efficiency in gathering and utilizing threat intelligence. Organizations should determine that capable and reliable sources of information are well-integrated into the data collection and absorption process. It means choosing and managing the sources in a way that they harmonize with the established process and with each other, to provide comprehensive

and effective threat intelligence. Lastly, the gathered threat intelligence should not just be accumulated but also analyzed and converted into actionable insights. Actionable observations mean that the data should inform decisions, guide responses, and/or proactive measures to manage and mitigate cyber threats effectively.

From there, the threat hunter should either work closely with the threat modeling team or use their knowledge of organizational technology and business functions to threat model how relevant attacks may manifest in the environment.

Using intelligent threat models as a guide, threat hunters can proactively review the data from the intelligent AI-Native Security Operations Platform™ systems in the environment focusing on severe malicious or anomalous behavior. Finally, it is critical that the threat hunter leverage the intelligent systems to automate document findings and remediation actions to promote lessons learned and create internal intelligence.

The threat-hunting cycle requires high-quality skills, creativity, and focus by leveraging intelligent threat-hunting capabilities and finding threat events faster with accuracy. For leveraging AI-Native Security Operations in intelligent threat hunting, focus on the following three core components:

- **Data Collection and Integration:** This involves the identification, gathering, and integration of previously unknown cybersecurity methods or techniques. By continuously updating and enhancing threat intelligence, security operations

can stay ahead of evolving cyber threats, determining a more adaptive and robust defense mechanism.

- **Threat Data Modeling:** Utilizing threat data to develop computational or statistical models aids in understanding, predicting, and mitigating cyber threats. By processing and analyzing data, selecting relevant features, and deploying trained models, AI can enhance threat anticipation, risk management, and resource allocation.
- **Continuous Network Monitoring:** Observing and analyzing daily data flow and activities across a network is essential for detecting anomalies and determining proactive security. Continuous monitoring supports the identification of deviations from established baselines, enabling swift incident responses, policy enforcement, and ongoing security posture requirement.

Deloitte and Palo Alto Networks AI-Native Security Operations Platform™ framework

This is the new vision for autonomous AI-Native Security Operations Platform™. Deloitte associates with Palo Alto Networks' Cortex® XSIAM™ to provide extended security intelligence and automation management. Cortex XSIAM is a cloud-delivered, integrated, AI-native security

operations feature that unifies critical functions, including AI-powered CDL, TIP, XDR, EPP, ASM, UEBA, SOAR, CDR, and Management, Reporting, and Compliance.

Clients may consolidate multiple products into a single, integrated platform, cutting costs, improving operations, and increasing analyst productivity. Deloitte and Palo Alto Networks may provide an intelligent data foundation that may easily integrate telemetry from the source, providing unified security operations across hybrid IT architecture. Below is the overview of the Deloitte and Palo Alto Networks' approach.

- **AI-Powered Cloud Data Lake (CDL)** Delivers log management, correlation and alerting, reporting, and long-term data retention.
- **Threat Intelligence Platform (TIP):** Aggregates, scores, and distributes threat intelligence data, including the industry-leading Deloitte and Unit 42™ threat feed, to third-party tools and enriches alerts for context and attribution.
- **Extended Detection and Response (XDR):** Gathers telemetry from sources for unrivaled detection coverage and accuracy, with a high-quality number of technique-level detections in the 2022 MITRE ATT&CK® evaluations.

- **Endpoint Protection Platform (EPP):** Prevents endpoint attacks with a demonstrated endpoint agent that blocks exploits, malware, and file-less attacks and collects full telemetry for detection and response.
- **Attack Surface Management (ASM):** Provides embedded attack surface management (ASM) capabilities for an attacker's view of your organization, with asset discovery, vulnerability assessment, and risk management.
- **User and Entity Behavior Analytics (UEBA):** Uses machine learning and behavioral analysis to profile users and entities and alert them on behaviors that may indicate a compromised account or malicious insider.
- **Security Orchestration, Automation, and Response (SOAR):** Automates many use cases with hundreds of built-in playbooks and offers customization with a visual drag-and-drop playbook editor.
- **Cloud Detection and Response (CDR):** Analyzes cloud inspect, flow, and container host logs together with data from other sources for broad detection and response across your hybrid enterprise.
- **Management, Reporting, and Compliance:** Simplifies operations, centralizing many configurations, monitoring, and reporting functions, including endpoint policy management, orchestration, and response.

A streamlined data onboarding process allows security operation teams to easily add new data sources while an extended data model normalizes and correlates data for schema on-read data access. Cortex XSIAM also automatically stitches together endpoint, network, cloud, and identity data, so it may detect advanced threats with precision and simplify investigations with cross-data insights. Security analysts may swiftly investigate incidents by providing an integrated picture of the attack, including intelligent alert grouping and collected information about the root cause. Embedded automation may enrich alerts, respond to malicious activity, and close low-risk alerts before they reach the queue

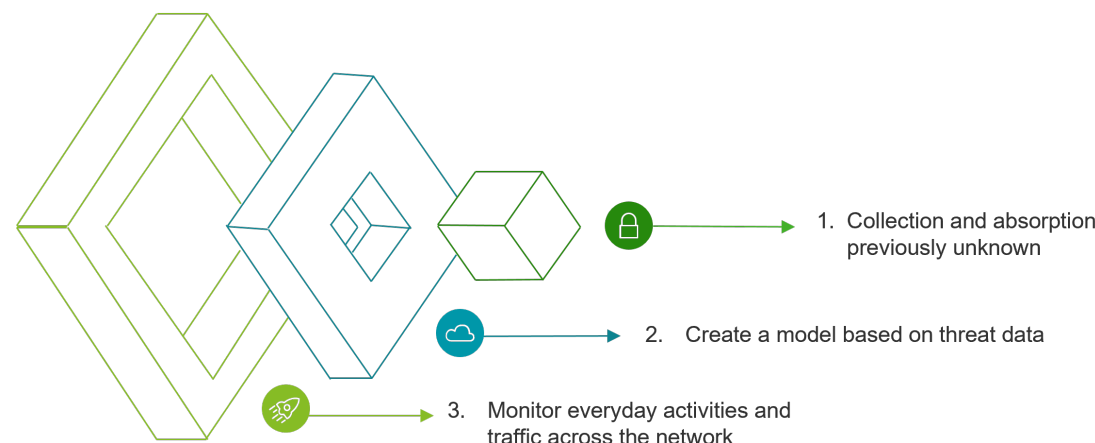


Figure: Intelligent threat hunting

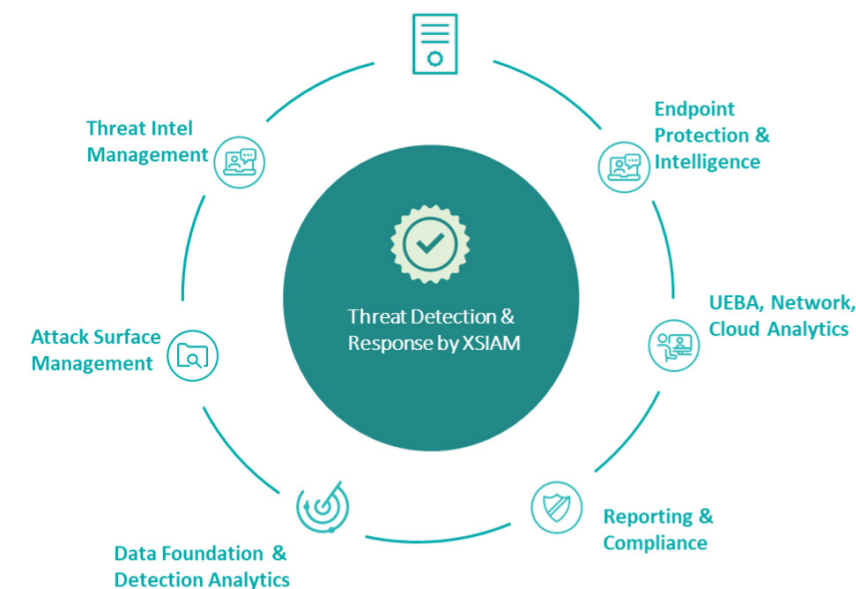


Figure: AI-native security orchestration and automation

allowing analysts to focus on the few threats that require human intervention.

The AI-powered SOC solution, XSIAM, is in production at Palo Alto Networks security operations center and is reducing over one trillion events per month to a handful of analyst incidents per day.

The legacy security operations capabilities required clients to operationalize and enhance the product. XSIAM provides continuous threat intel updates from more than 85,000 clients, updates machine learning (ML) detection models, and automatically distributes the latest protections to deployments to safeguard clients from advanced and fast-moving threats. In addition, using AI-Native Security Operations Platform™ accelerators, Deloitte and Palo Alto Networks share the responsibility of helping protect our client's ongoing operations.

The AI-Native Security Operations Platform™ service is structured in two main phases: Design and Develop, and Build, Test, and Deploy.

In the **Design and Development** phase, initial planning encompasses project requirements, timelines, team members, action items, and the finalization of the project plan. Discovery involves analyzing deployment scenarios and information gathering for targeted design workshops. This leads to the Design and Develop Workshop.

The **Build, Test, and Deploy** phase starts with onboarding, where configurations are set to launch the Cortex XSIAM framework. This includes setting up the Cortex Data Lake, Application, access validations, add-on applications deployment, third-party log collection, agent deployment, and recommendations for integrating firewalls or third-party logs. Endpoint Policy Tuning is then initiated, adjusting event profiles and creating required exclusions. Analytics are crafted once a significant portion of endpoints are tuned. Correlation rules are established post-onboarding and mapped as per the set use cases.

Knowledge transfer is central, covering areas such as the incident response

process, elements of security operations, and incident management methodologies. Deloitte may provide insights on incident management, focusing on identification, investigation, mitigation, and improvement strategies.

Other areas of focus could be an assessment of XSIAM capabilities, system usage, custom correlations, alert management, incident handling, playbook triggers, alert tuning, dashboard reporting, threat intelligence management, and attack surface management. Custom or standard playbooks might be deployed. Documentation, including configuration, operations, maintenance, and recommendations, will be provided by Palo Alto Networks.

Monitor and Operate

Deloitte, being a globally recognized consultancy and managed service provider, offers a plethora of managed services for an AI-Native Security Operations Center (SOC) to clients. Here are managed services offerings, particularly focusing on the integration and management of AI-Native SOC:

• Security Monitoring and Management:

- **24/7 Monitoring:** Continuous oversight of network traffic, applications, and infrastructure.
- **Incident Response:** Efficiently responding to and managing cybersecurity incidents and breaches.
- **Threat Hunting:** Proactively seeking signs of malicious activity within your network.

• AI and Automation Integration:

- **AI Implementation:** Deploying AI algorithms to enhance analytics, behavioral analysis, and threat detection
- **Automation Setup:** Implementing automated workflows for incident response, patch management, and other SOC activities.

• Strategy and Governance:

- **Compliance Management:** Determining that the AI-Native SOC adheres to relevant regulatory and compliance standards.

- **Security Strategy:** Developing and refining cybersecurity strategy in line with organizational objectives.

- **Policy Management:** Creating, implementing, and managing security policies.

• Technology Management:

- **Technology Requirement:** Determining that many SOC technologies are functioning optimally and updating them as required.
- **Integration Services:** Combining AI-Native SOC technologies with existing IT systems and processes.

• Endpoint Protection Management:

- **Device Management:** Overseeing many endpoint devices to determine they are secure and compliant.
- **Vulnerability Management:** Regularly identifying, assessing, and mitigating vulnerabilities.

• Threat Intelligence Services:

- **Intelligence Gathering:** Collecting data on emerging threats and cyberattack strategies.
- **Intelligence Analysis:** Analyzing data to identify potential threats to the organization.

• Cloud Security Services:

- **Cloud Access Security:** Managing and monitoring access to cloud environments.
- **Cloud Compliance:** Determining cloud configurations comply with regulatory requirements.

• Identity and Access Management:

- **User Management:** Administering user accounts, privileges, and credentials.
- **Access Reviews:** Regularly reviewing and checking user access to resources.

• Cyber Risk Management:

- **Risk Assessment:** Evaluating the cybersecurity risks associated with organizational assets and data.
- **Risk Mitigation Planning:** Developing strategies to mitigate identified cybersecurity risks.

• Training and Awareness Programs:

- **Employee Training:** Developing and delivering cybersecurity training for staff.
- **Phishing Simulations:** Conducting simulated phishing attacks to improve employee awareness and response.

• Customized Reporting and Analytics:

- **Security Reporting:** Creating customized reports on the cybersecurity posture and incidents.
- **Advanced Analytics:** Utilizing AI to derive actionable insights from data across the security landscape.

• Advisory and Consulting Services:

- **Cybersecurity Consulting:** Providing experience on enhancing the cybersecurity posture.
- **SOC Advisory:** Offering insights and advice on requiring and enhancing SOC operations.

• Managed Detection and Response (MDR):

- **Advanced Threat Detection:** Utilizing AI to detect sophisticated threats.
- **Coordinated Response:** Managing and coordinating the response to detected threats.

• Incident Response and Forensics:

- **Forensic Analysis:** Investigating incidents to determine their cause and impact.
- **Incident Mitigation:** Implementing strategies to contain and mitigate incidents.

This broad suite of services aligns with various aspects of managing and requiring an AI-Native SOC, determining organizations can protect their assets while adhering to compliance standards and leading practices in cybersecurity. This also enables them to leverage Deloitte's experience to enhance their cybersecurity posture amidst evolving threat landscapes. Remember that specific offerings may vary, and it may likely be leading to discuss directly with Deloitte for precise services tailored to specific circumstances.

If you start with a Greenfield environment, detection and response of incidents are no longer limited to building your own

detection stack. This offers the benefit of offloading the integration design complexity to the vendor, correlating alerts with data infused from individual alerts and other stack components, and spending more time triaging the alert. Otherwise, if your organization has an existing set of security tools, build a framework using a fabric of interconnection that steps back from integrated point technologies and their silos, which is infamously referred to as Analyze your organization's maturity and the build priority requirement matrix.

In this matrix, document the security areas organizations are doing well and begin ideating the tactical areas that may potentially impact the businesses. Perform this activity with cross-functional teams and recognize the benefit, expected time to build/develop/integrate, the effort required, and obstacles. There should be a greater emphasis placed on strategic items that may effectively measure the impact of business risk profiles.

When selecting an approach, it is required to evaluate the efficacy outcomes of a detection stack in comparison with the types of threats it detects. Detection of common types of threats is tactical; however, selecting the leading tool that may detect uncommon threats is strategic. The management of security risks should align with organizational priorities and address the threat landscape specific to your organization.

To build an AI-Native SOC, focus on automating what is already pre-defined, mundane, and tedious by providing SOAR orchestration playbooks that may take what you do and make it agile. The challenge lies not only in identifying what action is required to be taken but also in taking them. Organizations should review and understand the below points before making the decision.

- Think about how efficiently and at what expense the measures should be taken.
- Who should be involved in their coordination or agreement? Should further steps be taken in addition to those needed for the attack's containment and recovery?

- Could the expense of the legal defense and consumer notice outweigh the actual loss of the server, user, and data?

There may be classic organizational incident response management documents in place, but the specific lies in categorizing and identifying the common pattern of alerts with respect to time, the many common types of alerts by considering the quantity, quality, and quickness of each alert, as it is related to the amount of operational effort required to take an action. Build SOAR operational processes and task-level playbooks only after observing daily operations for an acceptable period. Consider the cross-functional teams' inputs required to build the playbook and align on the approach that require the need to be included for automating a workflow.

Lastly, automate the operational approach to reduce human effort and alert fatigue. Concentrate not only on initial deployment but also on how to continuously build new use cases that improve the efficiency of your security operations program with SOAR. The high-level goal is to set up the approach with the applicable functionalities with proper action to reduce the impact of human effort. Not to mention automation normally comes with its own catch if the human approval element is neglected. It is highly important that critical automation playbooks add the essence of human approval so that we can fix the applicable problem with the applicable solution. Eventually, the result should be directly related to how well it could mitigate the consequences of the discovered alert. To enhance security operations processes and alerting incorporate improved event enrichment, event correlation, and aggregation. Building a playbook may be time-consuming but once done, it may add a lot of value to the operations.

Operating Procedures for the categorized alerts and slowly build upon them by observing the operational cost benefit of introducing automation.

The strength of the Deloitte and Palo Alto Networks relationship

The strength of the Deloitte and Palo Alto Networks relationship Deloitte's award-winning Cyber and strategic Risk consultants have joined forces with Palo Alto Networks and its security capabilities platform. Together, we're working to provide a broad range of capabilities that simplify the complex software development lifecycle, while increasing speed, agility, and enablement so that organizations like yours may better protect their infrastructure and workloads at many stages of the development lifecycle. Our joint solution may aid you in creating a cyber-minded culture for your organization so that it can move forward faster and stronger, fuel more innovation, and stay more resilient in the face of persistent and ever-changing threats—while accelerating time to market and reducing costs.

Authors

Palo Alto Networks Alliance Leaders



Kieran Norton

Principal
US Cyber & Strategic Risk
Deloitte & Touche LLP
kinorton@deloitte.com



Jane Chung

Managing Director
US Cyber & Strategic Risk
Deloitte & Touche LLP
jachung@deloitte.com



Siddharth Kantroo

Advisory Senior Manager
US Cyber & Strategic Risk
Deloitte & Touche LLP
skantroo@deloitte.com



Anthony Polzine

Senior Manager
Global Partner Solution Architect Palo Alto Networks
apolzine@paloaltonetworks.com

Deloitte.

About this publication

This publication contains general information only and Deloitte and Palo Alto Networks are not, by means of this publication, rendering accounting, business, financial, investment, legal, tax, or other professional advice or services. This publication is not a substitute for such professional advice or services, nor should it be used as a basis for any decision or action that may affect your business. Before making any decision or taking any action that may affect your business, you should consult a qualified professional adviser. Deloitte and Palo Alto Networks shall not be responsible for any loss sustained by any person who relies on this publication.

As used in this document, "Deloitte" means Deloitte & Touche LLP, a subsidiary of Deloitte LLP. Please see www.deloitte.com/us/about for a detailed description of our legal structure. Certain services may not be available to attest clients under the rules and regulations of public accounting.