Security Analytics Trends 2026

Tackling Blind Spots, Insider Threats, and Compliance Challenges

exeon

Executive Summary

Rising Threat Complexity

Cyberattacks are surging in frequency and sophistication. In early 2025, organizations faced a **47% increase in weekly cyber attacks on average** [1]. This uptick underscores the need for advanced security analytics like User and Entity Behavior Analytics (UEBA), Network Detection and Response (NDR), and Security Information and Event Management (SIEM) solutions that can detect subtle threats amid the noise. Organisations across banking and financial services (BFSI), manufacturing, the public sector, and critical infrastructure like energy, utility and healthcare are reassessing their security monitoring as threat actors grow more adept at evading traditional defenses.

47% increase of cyberattacks per organization weekly

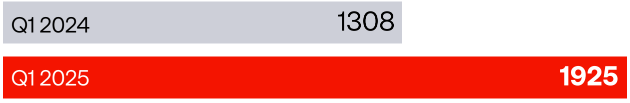

Shadow IT and Visibility Gaps

go to page 5 ⟶

Many organizations lack visibility into significant portions of their IT environment. IT departments are often unaware of roughly two-thirds of the SaaS applications in use [2], and as much as 30–40% of IT spending occurs in shadow IT outside official oversight [3]. Custom business-critical applications, unmanaged APIs, and cloud services frequently operate as blind spots – their logs may sit idle or never even reach the SIEM due to complex formatting requirements. The result is fragmented visibility, leaving security teams uncertain if they're truly catching all threats. Indeed, **71% of SOC analysts believe their organization may already be compromised without their knowledge** [4].

71% of SOC analysts believe their org is already compromised)

Insider Threats Undetected

go to page 9 ⟶

83% of organizations experienced an insider attack in the past year

Insider threats – whether negligent users, malicious insiders, or stolen credentials – have become a top concern. A staggering **83% of organizations experienced an insider attack in the past year** [5], and nearly half say such incidents are increasing in frequency [6]. Yet these threats often blend in with normal activity and go unnoticed: the average time to contain an insider incident is over 80 days [7]. Traditional tools often miss the subtle warning signs in user or entity behavior. Without AI-driven baselining and custom behavioral analytics, credential abuse and policy violations can occur for months undetected – a particularly acute risk in data-rich industries like finance and healthcare.

SOC Alert Overload and Inefficiency

go to page 14 ⟶

Security Operations Centers (SOCs) are inundated with data and alerts, leading to alert fatigue and slow response times. SOC teams handle 4,484 alerts per day on average, and **67% of those alerts are ignored due to volume and false positives** [8]. Over half of security teams feel overwhelmed by the volume of alerts, and 55% admit they lack confidence in prioritizing and responding to them [9]. Analysts waste nearly a quarter of their time (27%) chasing false positives [9], delaying real incident response. Tool sprawl exacerbates the problem – enterprises juggle dozens of disparate security tools, with complexity itself now a major impediment (52% of executives say security solution fragmentation limits their ability to combat threats [10]). The outcome is an overtaxed SOC that struggles to react quickly (often too late in the attack kill-chain) and precisely (high false alarm rates).

67% of 4,484 alerts per day are ignored due to volume and false positives.

Data Privacy and Regulatory Pressure

go to page 19 ⟶

24 hours for companies to report major incidents according to NIS2 directive

Data sovereignty and privacy regulations impose constraints on how and where security data gets analyzed, even as new laws demand more robust monitoring. Many security analytics solutions – especially UEBA/SIEM technology – rely on cloud analytics, but organizations in regions like DACH (Germany, Austria, Switzerland) and other jurisdictions face GDPR and data residency requirements that limit sending sensitive user logs to foreign or public clouds. This can constrain analytics and reduce operational control if suitable EU-based or on-premises options aren't available. At the same time, regulations are raising the bar: the EU's Digital Operational Resilience Act (DORA) (effective January 2025) requires financial institutions to continuously monitor ICT systems and report major incidents within hours [11], and the NIS2 directive expands to new sectors (manufacturing, healthcare, etc.) **with a mandate to report incidents within 24 hours of detection** [12]. These rules make timely detection and reporting of anomalous activities (including those spotted by UEBA) a legal obligation. Organizations without automated, real-time behavioral analytics and reporting face a heavy manual workload to meet compliance, tying up resources.

Let's dive deeper into these challenges

In the sections that follow, we dive deeper into these trends and challenges – from the hidden risks of shadow IT to the importance of detecting insider threats – and examine how companies in BFSI, manufacturing, public sector, healthcare and other critical industries are adapting their security strategies going into 2026.

Each section provides data-driven insights and relevant statistics, offering a comprehensive outlook for security leaders – from CISOs and CIOs to SOC managers as well as risk and compliance heads – to understand the evolving landscape.

Shadow IT and Blind Spots in Monitoring

Despite heavy investments in cybersecurity, organizations continue to struggle with **visibility gaps** in their IT environment. A major cause is **shadow IT** – systems and applications that operate outside the purview of central IT or security teams. Studies show that IT departments are unaware of about two-thirds of the SaaS apps running in their organization [2]. In fact, the number of cloud services in use can be

3× higher than what IT knows about [2]

This unchecked adoption of tools (often by business units seeking convenience) results in silos of data and logs that may never make it into centralized security monitoring.

Blind spots arise when critical systems – from legacy ERP databases to new cloud workloads or APIs – are not integrated with the SOC's monitoring capabilities. Many firms still focus their defenses on endpoints or network perimeter traffic, while neglecting application-layer logs, user activity feeds, or IoT/OT systems. Even when logs from these sources are collected, they often "sit in the SIEM" without effective analysis, due to formatting complexities or lack of correlation rules.

For example, custom application logs might be ingested as raw data but not translated into useful behavioral insights or alerts because building the necessary parsers and use-case rules is time-consuming. As a result, executives and security teams might have a false sense of coverage, when in reality large portions of their digital infrastructure are dark to them.

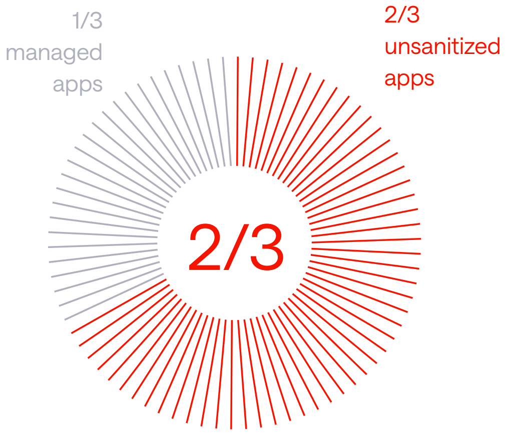

IT departments are unaware of nearly two-thirds of SaaS apps in use.

71%

Over 71% of SOC analysts fear their organization may have been compromised without them knowing [4].

30 - 40%

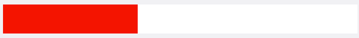

Shadow IT accounts for 30–40% of all IT spending [3].

68%

Up to 68% of organizations have undocumented or “shadow” APIs [14].

Fragmented Visibility:

Over **71% of SOC analysts fear their organization may have been compromised without them knowing** [4], a telling indicator of the visibility void. Similarly, security surveys reveal that 76% of SMBs perceive shadow IT as a moderate or severe threat [13] – unknown assets are essentially unmonitored assets.

And it’s not just small businesses; Gartner estimates that in large enterprises, shadow IT accounts for **30–40% of all IT spending** [3], which underscores how much technology runs outside official oversight. Beyond SaaS applications, undiscovered APIs and service accounts present another blind spot.

Up to **68% of organizations have undocumented or “shadow” APIs** [14], which may expose data or allow access without the security team’s awareness. In one study, nearly 31% of observed malicious requests targeted unknown or unprotected APIs [15], exploiting these gaps. Each unseen application or interface is a potential foothold for attackers and an entry point for insider abuse that won’t trigger any alarm if logs aren’t analyzed.

> “Compounding the issue, organizations often face tough choices about log retention and SIEM ingestion due to cost or performance constraints”

Traditional SIEM licensing is typically based on volume of data ingested, leading many companies to filter out “less important” logs or shorten retention to control expenses. Ironically, some of the logs deemed high-volume (and thus expensive) are exactly those that could provide early warning of an attack – identity and application logs, for instance, where 90–95% of events may be routine but the few anomalies are critical [16]. According to industry analysts, **nearly 80% of SIEM logs are “noise” with little analytical value** [17], and security teams spend considerable effort storing and processing this noise.

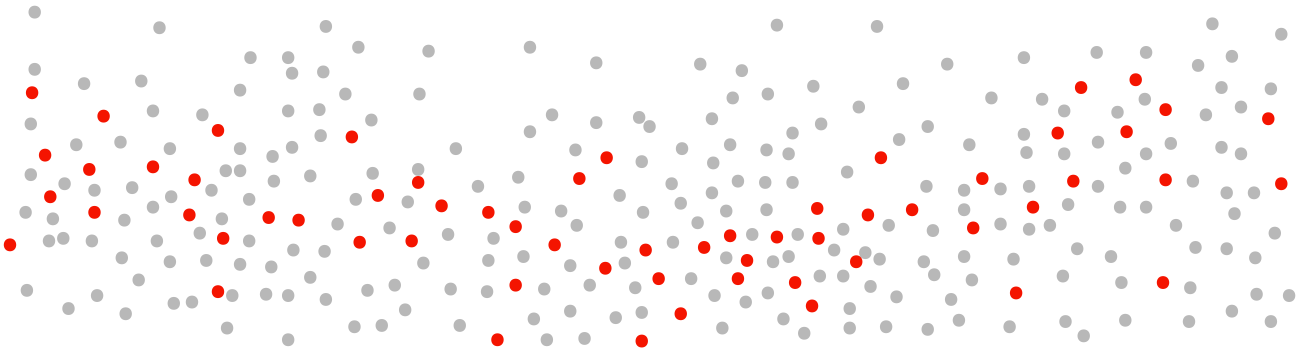

Nearly 80% of SIEM logs are “noise” with little analytical value [17].

The hidden cost is that to stay within budget, teams sometimes drop high-value data sources entirely, inadvertently creating dangerous blind spots [18][16]. In other words, cost-driven data filtering means attackers can hide in the gaps – if logs from an ERP system or an Okta identity service aren’t ingested, an insider abusing those platforms could fly under the radar.

The net effect of shadow IT and partial monitoring is a **fragmented security picture.** Different tools and silos each provide a narrow view (endpoint, network, cloud, etc.) with no single vantage point. Under these conditions, subtle indicators of compromise can go unnoticed because they only become evident when data is correlated across systems. For instance, a suspicious use of a privileged account in a core banking application might only stand out if seen alongside unusual network access patterns – something a siloed approach would miss. Multiple surveys confirm this challenge: **76% of security professionals say too many point tools create blind spots in their defenses** [19]. The move to cloud and hybrid IT has only increased this complexity, as visibility into cloud workloads and SaaS usage lags behind on-premise capabilities [20].

76% of security professionals say too many point tools create blind spots in their defenses [19].

To address these blind spots, organizations are gradually shifting toward more unified monitoring and data integration. Practices like centralized log management and security data lakes are gaining traction, allowing firms to ingest diverse data sources without the high cost of keeping everything in a traditional SIEM.

More importantly, there’s a growing recognition that **behavioral analytics** must extend beyond endpoints to users, applications, and even service accounts. The use of advanced security analytics technologies like UEBA is one response to this need – by baselining normal behavior for each entity (user or machine), these security analytics tools can turn raw log data from previously under-monitored systems into meaningful anomalies and alerts.

The 2026 Outlook

The trend 2026 is toward converged platforms that combine network, cloud, endpoint and identity telemetry to close the visibility gaps. As noted in one industry report, modern SIEM/analytics solutions are evolving into the “central nervous system” of the SOC, correlating across all domains and applying AI to highlight events that would be invisible in isolated silos [21][22]. Organizations that succeed in this integration will significantly reduce their shadow IT risks and uncover threats lurking in formerly blind areas.

Missed Insider Threats and Undetected Attacks

76%

of organizations identify growing business and technological complexity, including factors like cloud adoption, increased user access, and IT sprawl, as key drivers of rising insider risk.

Insider threats – whether malicious employees, compromised insiders (stolen credentials), or unintentional mistakes – have emerged as one of the most pernicious security challenges. In high-control industries like BFSI or pharmaceuticals, insiders have privileged access to sensitive data and systems, making them uniquely positioned to cause harm. Unfortunately, many traditional security tools are ill-suited to detect insider behavior that appears legitimate. Unlike malware or external attacks that might trigger antivirus or firewall alerts, insider attacks often involve an authorized user doing allowed things, but in anomalous ways or inappropriate contexts.

The prevalence of insider incidents has climbed steeply.

A 2024 global survey found **83% of organizations experienced at least one insider security breach in the prior year** [5], up from just over half of organizations a few years earlier. Moreover, 74% of organizations believe insider attacks have become more frequent in recent times [23], fueled by factors like the rise of remote work, cloud collaboration, and the sheer complexity of modern IT (which provides more hiding places).

Notably, **76% of organizations identify growing business and technological complexity, including factors like cloud adoption, increased user access, and IT sprawl, as key drivers of rising insider risk** [24]. With more employees accessing data from more places, it becomes harder to baseline “normal” behavior without advanced analytics.

Certain industries face particular insider threat risks. In financial services, there is the danger of internal fraud or data theft (e.g. an employee with access to transaction systems abusing credentials). In healthcare, insiders might snoop electronic health records or misuse patient data. Manufacturing and pharma worry about supply chain attacks and intellectual property theft by employees or contractors. Public sector and critical infrastructure organizations face potential sabotage or espionage from insiders. In all cases, the impact of an insider incident can be immense – a recent study pegged the average annual cost of insider threats at $17.4 million per organization in 2025 (up from $16.2M in 2023) [7]. These costs accrue from investigations, system downtime, remediation, regulatory fines, and loss of trust. Particularly costly are **incidents involving compromised credentials, which average $779k per incident** [25] because they often allow attackers to roam undetected as legitimate users.

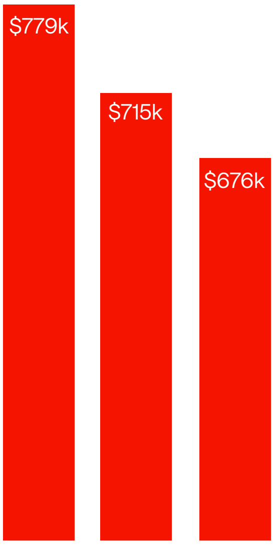

Credential theft is the most expensive at **$779k**, followed by malicious acts at **$715k**, and negligence at **$676k**.

> "On average, companies take over 80 days to contain an insider incident."

The detection problem is at the heart of why insider threats are missed. On average, companies take over 80 days to contain an insider incident [7] – meaning an insider may operate for nearly a full quarter before being stopped. Part of this delay is that initial detection is slow.

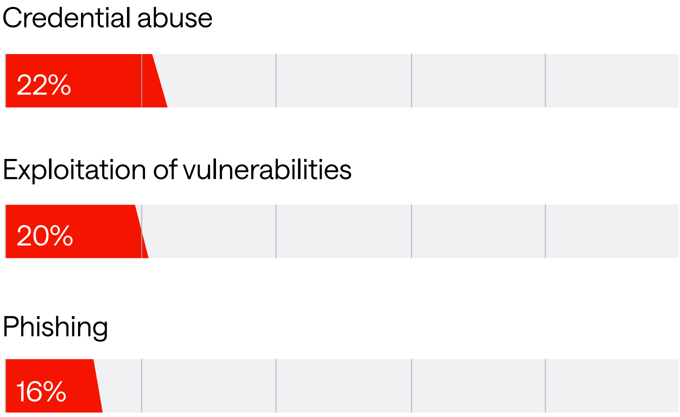

Known initial access vectors in non-Error, non-Misuse breaches (n=9,891) [26]

Many insider actions do not trigger clear-cut alerts: for example, an employee gradually downloading more data than usual, or accessing systems they rarely use, might not set off any static rules. Without behavioral baselining, such patterns blend into the background. **Verizon's data breach research highlights that the human element is a factor in the majority of breaches, and that stolen credentials are leveraged in over 22% of breaches** [26] – essentially turning an external attack into an "insider" incident by using legitimate access. Once an attacker is logging in with a valid user's credentials, traditional perimeter defenses (firewalls, etc.) offer no help [27]. The attacker's actions must be caught through recognizing abnormal usage by that user account, which is exactly what UEBA is designed to do.

5x increase

Organizations experiencing 11-20 insider attacks annually saw a fivefold increase YOY

> “An AI assistant embedded in a CRM system could be manipulated to export sensitive customer data to an external location by a cleverly crafted user input.”

Examples of missed insider threat scenarios abound.

In many cases, employees had been violating policies or exhibiting suspicious behavior for weeks but were only caught after a major incident or whistleblower report.

Common red flags that are often overlooked include:

- an employee accessing systems they never have before,
- logins at odd hours or from unusual locations,
- sudden spikes in database queries or file downloads,
- or multiple authentication failures on sensitive systems.

If these events are not correlated and analyzed in context, they might be written off as normal IT issues or remain buried in log files. The 2024 Insider Threat Report noted **a fivefold increase in organizations experiencing a high volume (10+ per year) of insider incidents**, indicating that **many smaller incidents likely fly under the radar until they accumulate** [6].

Importantly, the notion of who, or what, can be an insider is evolving. As organizations deploy **agentic AI systems** (AI-powered agents capable of acting autonomously across email, file systems, identity platforms, and business applications), these tools are gaining system-level access on par with or greater than human users. If compromised – whether through prompt injection, supply chain tampering, or malicious retraining – an agent could misuse privileges under the guise of “helping users,” while quietly exfiltrating data or modifying records [60]. For example, an AI assistant embedded in a CRM system could be manipulated to export sensitive customer data to an external location by a cleverly crafted user input. Because these agents typically fall outside conventional identity and behavior monitoring frameworks, their activity may be neither logged properly nor baselined for anomalies. The result is a **new class of non-human insiders** with broad access but limited oversight.

Security teams will increasingly need to monitor **machine and agent behavior** in parallel with human users. This raises important architectural questions: Should each AI agent be treated as its own entity within security analytics models? How do we baseline “normal” behavior for a non-human actor that executes varied tasks for many users across systems? These are urgent topics for SOCs and risk leaders as adoption of agentic AI accelerates in enterprise workflows. Security analytics systems that do not account for these entities risk blind spots where attackers can pivot through seemingly benign machine behavior that is, in fact, malicious or compromised.

Insider threats are also not limited to direct employees.

Contractors, third-party partners, or service accounts can act as insiders. For example, a support contractor with VPN access might exfiltrate data, or an attacker might compromise a third-party software account to get inside the network. These scenarios stress the need for entity behavior analytics (the “E” in UEBA) – monitoring not just human users but also service accounts, APIs, and machines for anomalies.

Modern security analytics solutions use machine learning to establish a baseline of normal behavior for each user and entity (what resources they access, typical work hours, network activity patterns, etc.). They can then flag subtle deviations: e.g., if a user who normally accesses HR records on weekdays suddenly bulk-downloads a customer database on a weekend, or if a privileged account starts running queries it never ran before. AI-driven UEBA is increasingly essential because insider attacks evolve and “learn” from detection methods. Static rules alone (like fixed thresholds for downloads) often produce either too many false positives or miss the context that makes an action truly suspicious. Security analytics tools can adapt as behavior changes over time and as new types of insider tactics emerge.

In fact, security surveys show 64% of cybersecurity professionals consider malicious insiders or compromised accounts a greater danger than external attackers [28], reflecting a consensus that advanced behavioral monitoring is a top priority.

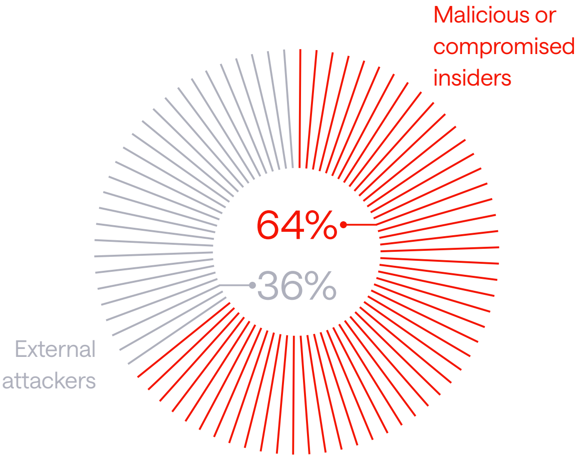

Greatest cybersecurity threat according to professionals.

The 2026 Outlook

In summary, missed insider threats remain a critical vulnerability across industries. The combination of more complex IT environments, more determined adversaries (including nation-state actors who recruit insiders or use stolen creds), and higher stakes (both financial and regulatory) means that organizations must strengthen their insider threat programs.

This includes not only deploying advanced security analytics technologies, but also establishing insider response playbooks, conducting regular user access reviews, and fostering a culture where employees are aware of insider risks (e.g. know how to spot social engineering attempts aimed at stealing credentials). Going forward, security strategies must explicitly include AI agents as monitored identities. By 2026, we can expect insider threat detection to be a standard component of cybersecurity strategies – not an afterthought – especially in sectors like BFSI, pharma, and critical infrastructure where the “insider damage” can be catastrophic.

SOC Overload and Operational Inefficiencies

Security teams today face an operational crisis: the sheer volume of alerts and data to process has overwhelmed human analysts, leading to slower responses and important signals being missed. Nowhere is this more evident than in the SOC (Security Operations Center), the nerve center where threats are supposed to be detected and neutralized. Across industries, SOC analysts consistently report high levels of stress and burnout, directly attributable to alert overload and inefficient processes.

One global study of 2,000 SOC analysts painted a stark picture: an average enterprise SOC receives about **4,484 security alerts per day, of which 67% are simply ignored due to the unmanageable volume and frequent false alarms** [8]. This means only roughly one-third of potential issues even get a look from humans, a gap that adversaries can exploit. Another industry survey found 51% of SecOps teams feel overwhelmed by the volume of alerts they must handle, and 55% are not confident in their ability to prioritize and respond effectively [9]. In practice, many SOCs can only triage a fraction of incoming alerts; the rest backlog or auto-close, creating a dangerous “alert fatigue” feedback loop.

False positives are a huge contributor to the noise. With legacy SIEM rules or signature-based detection, it’s common to get alert storms for benign events. Analysts end up spending about 1/4 of their time (27%) investigating false positives [9], essentially chasing ghosts. This is wasted effort that delays the investigation of real threats. It also saps morale – constantly responding to false alarms can lead teams to become jaded, or worse, start ignoring alerts altogether. In fact, approximately **70% of SOC staff report that the relentless alert overload negatively impacts their well-being and work/life balance** [29].

Cases are documented where analysts intentionally turn off noisy alert rules or tune them to be less sensitive, just to reduce the torrent of notifications [29]. The risk, of course, is that actual attacks might then slip through unnoticed.

4,484

daily alerts on average are handled by SOC teams.

67%

are simply ignored due to the unmanageable volume and frequent false alarms.

70%

of SOC staff report that the relentless alert overload negatively impacts their well-being and work/life balance. [29].

Several root causes contribute to this operational inefficiency:

Exponential Growth of Log Data:

As organizations enable logging on more systems and expand into cloud, IoT, and remote work, the amount of security telemetry is exploding. Yet not all logs are intelligent signals – as noted earlier, up to **80% of ingested logs have little security value** [17], but they still consume analyst attention and SIEM resources.

80% of ingested logs have little security value.

Tool Sprawl and Siloed Systems:

Many enterprises have accumulated a patchwork of point security tools (for endpoint, network, cloud, identity, etc.) that are not well-integrated. **The average organization uses 83 different security solutions from 29 vendors** [10], creating a fragmented workflow for SOC analysts. They must swivel-chair between multiple consoles and correlating data manually. Over half of security executives (52%) say this complexity itself is the biggest impediment to effective operations [10]. When an alert comes in, gathering context might require querying three or four separate systems – a slow and error-prone process.

83 The average organization uses 83 different security solutions from 29 vendors.

Shortage of Skilled Analysts:

The cybersecurity skills shortage compounds the issue – there is an estimated **3.4 million person deficit in the cyber workforce** [30], meaning many SOCs are understaffed. Those on the front lines juggle too many alerts per person. Critical investigation and hunting tasks get deferred because the team is stuck in reactive mode.

3.4 million

person deficit in the cyber workforce.

Lack of Context/ Enrichment:

Alerts that lack context (the who/what/where of an event) force analysts to do extra legwork. For example, a raw alert like “excessive failed logins on Server X” might be a minor IT issue or a brute-force attack – to know which, an analyst has to dig for related data (which user? where from? what else did they do?). Without automation to enrich alerts with context, triage is slow and many alerts remain unaddressed.

The consequences of SOC overload are significant delays in detection and response. While some organizations reduce complexity by adopting unified security platforms, others successfully integrate Best of Breed solutions through lightweight, API-first architectures. Industry research shows that reducing tool fragmentation – whether via consolidation or seamless interoperability – can dramatically improve response times.

One study found that

Organizations with streamlined security operations detected incidents 72 days faster and contained them 84 days faster on average than those managing siloed tools.[31] [32]

This illustrates how much inefficiency can be hiding in a disjointed SOC workflow. When an attack like ransomware can encrypt an environment in hours, taking days or weeks to respond is obviously unacceptable.

Alert fatigue can also lead to missed attacks.

Real incidents often generate early warning alerts that may be low-severity or buried among false positives. If those alerts are ignored, the attack proceeds until a much more obvious symptom (like a server outage or big data exfiltration) is noticed – by then, damage is done. According to one survey, **71% of analysts believe there may already be compromises in their environment that they haven’t caught** [4], which aligns with the idea that overload is causing a confidence crisis in detection capabilities.

71%

of analysts believe there may already be compromises in their environment that they haven’t caught [4].

To combat these operational challenges, organizations are pursuing several strategies going into 2026:

Investing in Automation and AI: Automating level-1 triage tasks can dramatically cut down the workload. Many SOCs are deploying Security Orchestration, Automation and Response (SOAR) tools or automated playbooks to handle routine alerts (like auto-quarantining malware, gathering basic incident data, etc.). AI and machine learning are also being leveraged to reduce false positives – for example, AI-driven SIEM/UEBA/NDR can learn which alerts are truly meaningful by cross-referencing behavior patterns, thus filtering out noise. There is also interest in Generative AI assistants to help analysts summarize incidents or even suggest remediation steps, acting as force-multipliers for junior staff.

Alert Prioritization and Risk Scoring: Modern detection platforms now often apply a risk score to entities or incidents, helping analysts focus on the most critical issues first. By correlating multiple low-level events into one high-priority incident (for example, combining an alert of unusual login, plus a data access anomaly, plus a phishing email click for the same user), the systems can present a coherent story that this chain of events likely indicates a breach. This reduces the number of separate alerts and provides needed context in one package.

Interoperability Over Platform Lock-In: As security teams face mounting workload and alert fatigue, organizations are seeking ways to streamline operations without sacrificing depth or flexibility. While some pursue full-stack platforms, many are now favoring modular, interoperable solutions that integrate easily into existing ecosystems. Instead of monolithic “all-in-one” platforms, forward-leaning teams prioritize tools that are easy to deploy (e.g. log forwarding, out-of-the-box normalization), standards-based (API-ready), and analytics-rich, enabling fast time-to-value without vendor lock-in. This “best-of-breed, well-integrated” approach supports unified visibility while retaining control over architecture and data flow. The pressure to reduce tool fragmentation is real: **74% of security executives say their teams’ workload is excessive** [33], and 52% cite disconnected tools as a key barrier to effective threat response [34]. By 2026, many organizations aim to consolidate operational workflows – not necessarily into a single vendor stack, but into a cohesive, interoperable detection and response layer that connects endpoint, network, identity, and application telemetry in a manageable, analyst-friendly way.

Contextual and Early Alerting: Instead of tens of thousands of raw alerts, the goal is to generate fewer, smarter alerts. Behavior analytics plays a role here by focusing on anomalies that truly deviate from normal baselines (often indicating an advanced threat or insider). Additionally, continuous threat intelligence integration can help the SOC know which alerts map to known threat campaigns, lending insight into which alerts cannot be safely ignored. Early detection is key – catching an attacker in the reconnaissance or lateral movement stage [35] (when they test the waters or move between systems) prevents later-stage damage. This requires real-time analysis of patterns that span multiple systems (for example, a user account trying to access several servers it never did before might indicate lateral movement).

Ultimately, addressing SOC overload is as much about process and people as it is about technology.

Many organizations are revising their incident response workflows, implementing tighter SLAs for alert handling and better handoff procedures (so critical alerts aren’t dropped between shifts). Training and retaining skilled analysts is also crucial – some companies are rotating staff to prevent burnout and providing on-call support for major incidents to avoid fatigue. **KPIs like mean time to detect (MTTD) and mean time to respond (MTTR) are being closely watched by CISOs as key metrics of SOC efficiency.**

The 2026 Outlook

For 2026, an emerging best practice is to treat the SOC like a continuous improvement environment: regularly tune detection rules to eliminate noise, decommission tools that don't provide value, and periodically simulate high-alert scenarios to ensure the team can cope without collapsing under pressure.

In summary, SOCs in BFSI, manufacturing, government, and other sectors are all grappling with these operational issues – a bank's SOC might drown in fraud alert data, while a utility company's SOC has to monitor both IT and legacy (OT) apps without losing track. The solutions lie in smarter tech (behavior analytics, AI/ML, integrated platforms) and smarter operations (workflow automation, tool rationalization). The organizations that master these will gain a significant advantage: faster and more precise incident response, which in turn can mean the difference between quickly foiling an attack or suffering costly damage.

Data Privacy and Regulatory Challenges for Analytics

While organizations rush to improve their threat detection and response, they must do so under the constraints of data privacy, sovereignty, and regulatory compliance. These factors are especially salient in regions like Europe and Middle East as well as in industries like finance, critical infrastructure, and government, where data is highly sensitive and heavily regulated.

Two intertwined – and potentially contradicting – challenges stand out:

1. Limitations on data hosting/processing (which affect how security analytics can be deployed),

2. New regulations mandating stronger operational resilience and reporting.

Data Hosting & Sovereignty Limitations:

Many security analytics solutions leverage cloud infrastructure for scalability and advanced analytics. Cloud-based analytics can offer advantages in processing big data (applying AI across massive log volumes, etc.). However, for companies in jurisdictions with strict privacy laws, sending internal user behavior data or system logs to a public cloud (especially one hosted in another country) raises red flags.

> "Organizations must improve their threat detection and response while adhering to data privacy, sovereignty, and regulatory compliance constraints."

European data protection laws (GDPR), for instance, require that personal data (which could include user IDs, IP addresses, authentication logs, etc.) be handled with care and not transferred to regions without equivalent protections. Organizations in DACH, other European countries, and in the Middle East often have policies or regulations demanding that sensitive log data (like employee or customer information in security logs) remain within national or EU borders. As a result,

the “cloud-only” delivery model of some security analytics providers doesn’t meet these needs, forcing security teams to either seek on-premises or EU-local solutions, or accept reduced functionality.

Additionally, certain sectors have explicit sovereignty rules: government agencies and critical infrastructure operators may be required to use domestic clouds or on-prem systems for security monitoring to avoid espionage risks. In the financial sector, regulations like the ECB cloud outsourcing guidance and others mandate oversight and auditability of any cloud service – pushing some banks to prefer self-hosted analytics. If a security analytics tool requires uploading all Active Directory logs to a US-based cloud for analysis, many European banks simply won’t consider it, no matter how effective the analytics might be.

This creates a trade-off:

some of the most advanced analytic tools might be off-limits due to data privacy concerns. Organizations end up with constrained analytics if they can’t leverage cloud AI/ML at scale. Even when cloud is used, privacy concerns drive demand for data anonymization or pseudonymization in logs (which can complicate analysis by removing user identifiers). Vendors are responding by offering EU-based cloud instances or hybrid deployments (where sensitive data stays on-prem and only anonymized features go to the cloud). Nonetheless, it remains a limitation – the flexibility to host the solution where the customer needs it is now a competitive factor. Those providers that support on-prem or private cloud deployments are seeing uptake in highly regulated markets, whereas cloud-only solutions face hesitancy. In short, **“analytics sovereignty”** is becoming important:

organizations want powerful analytics and control over where the data lives.

At the same time, **privacy-preserving analytics techniques** are emerging. There’s interest in methods like federated learning (where models train on local data and only share model updates, not raw logs) or advanced encryption that allows cloud analysis without exposing plaintext data. These are still developing, but the key point is that user behavior data is sensitive, and analytics platforms are adapting to comply with privacy laws. As one market analysis noted, most **security analytics vendors are incorporating privacy measures to comply with global data protection laws while still analyzing large volumes of behavior data** [36]. This can include features like role-based access controls (so security analysts only see pseudonyms for users), data minimization (only storing security-relevant fields), and clear data retention policies aligned with regulations.

Regulatory Demands – DORA, NIS2, and Beyond

On the flip side of restrictions, regulators are also driving organizations to enhance their security monitoring and incident reporting capabilities. Notably, the EU has introduced new rules that explicitly or implicitly require continuous threat detection, including the use of advanced security analytics:

DORA (Digital Operational Resilience Act):

Focused on EU financial entities (banks, insurers, investment firms), DORA came into force with full applicability in January 2025 [37]. It mandates that firms maintain robust security monitoring and incident handling as part of operational resilience.

One of the stringent requirements is that significant cyber incidents must be **reported to regulators within hours** of being classified as major [11]. This means banks and financial institutions need the ability to detect anomalies rapidly and assess their impact. Without automated monitoring picking up suspicious activity, an organization might not even realize it's under attack until it's too late to meet the reporting deadline. DORA essentially forces BFSI firms to invest in real-time detection (the kind security analytics tools provide) because manual or ad-hoc monitoring will not catch incidents fast enough to comply [11]. Moreover, DORA expects firms to have **comprehensive ICT risk management frameworks**, implying that logs from all critical systems should be watched. The era of ignoring certain legacy system logs is over – if those systems are essential to operations, they fall under the resilience framework and need monitoring for anomalies.

DORA also brings scrutiny to third-party providers (including cloud services) [38]. Financial firms must ensure their security data in the cloud is protected and that they have continuity plans if a provider fails.

This dovetails with the privacy issue: regulators would not look kindly on a setup where a bank cannot account for or control security data sent to an external analytics cloud, especially if that poses concentration risk or oversight difficulties[39]. So DORA pushes for both improved analytics and careful vendor management – a tricky balance.

5 Pillars of DORA regluation

01 ICT risk managemt

02 Incident reporting

03 Digital resilience testing

04 ICT third-party risk management

05 Information sharing

NIS2 (EU Network and Information Security Directive 2):

24h

Within 24 hours, organizations must provide an “early warning” that a significant cybersecurity incident is suspected.

72h

Within 72 hours, organizations must provide a detailed report about the incident.

1 month

Within 1 month, organizations must provide a final report describing the incident, cause and mitigation efforts.

NIS2, adopted in 2022 and to be transposed by member states by 2025, broadens the scope of mandatory cybersecurity practices and incident reporting across many sectors. It now covers medium/large organizations in sectors like manufacturing, healthcare, energy, transport, financial market infrastructure, public administration, and more [40][41]. A key requirement is the **24-hour initial notification** for significant incidents [12]. This is an even tighter window than many were used to (the original NIS was 72 hours). To report an incident within 24 hours of awareness, organizations need to be aware almost immediately when something is wrong – which again underscores continuous monitoring.

NIS2 explicitly calls out the need for appropriate **risk management measures, including detection capabilities and fast response** [42]. It even suggests that entities should consider SIEM solutions to aggregate and analyze data across the network for effective threat detection [43] – essentially encouraging the uptake of advanced security analytics for compliance.

For industries like healthcare and manufacturing that may not have had the same level of SOC maturity as banks, NIS2 is a wake-up call. Hospitals, for example, now must monitor for suspicious network activity that could indicate ransomware, and report it quickly. This can be challenging given resource constraints, so automation is key. In general, NIS2 pushes companies to implement “state-of-the-art” cybersecurity measures, a term that is interpreted as current best practices – certainly, leveraging AI-driven security analytics fits that description [44][45].

Organizations are looking to automation as a compliance lifeline[46][47]: using tools that can detect and even report on incidents with minimal human intervention, to meet the tight timelines sustainably.

Other Regulations and Standards:

Beyond DORA and NIS2, there are numerous sector-specific or national rules heightening security requirements. For instance, in Germany the IT Security Act and BSI guidelines for KRITIS (critical infrastructure) operators mandate continuous monitoring and incident reporting. The US has directives for pipeline and utility security that require anomaly detection. The banking sector globally, via Basel/IOSCO, is emphasizing operational resilience which includes cyber-attack detection. Even privacy regulations like GDPR indirectly require good security monitoring – a breach of personal data must be reported within 72 hours under GDPR, which again implies you need to know you had a breach within that time frame. All this creates a regulatory environment in 2026 where not having advanced threat detection is not just a security risk, but a compliance risk.

One particular challenge is the resource burden of compliance reporting. Regulations like DORA and NIS2 require detailed reporting and sometimes continuous updates. If organizations lack tooling to automatically compile the relevant data (e.g. logs, indicators, impacted systems) for an incident report, they might have to pull this together manually under tight deadlines. This can be extremely labor-intensive – pulling log files, correlating events, writing up incident descriptions – potentially diverting the SOC from actually responding to the threat. For example, NIS2 requires a follow-up report within 72 hours and a final report within one month [48][49], meaning multiple stages of documentation. Companies are therefore increasingly interested in solutions that can generate compliance-ready reports or audit logs as a by-product of their monitoring. Indeed, some vendors now tout features like “one-click regulatory reporting” for incidents [50], which can output the details needed to inform authorities. Automation in this space is crucial – otherwise, the effort to manually comply could “bind” a lot of human resources, as the user prompt indicated. Security teams may find themselves writing reports more than hunting threats, if not careful.

> Automation in this space is crucial – otherwise, the effort to manually comply could “bind” a lot of human resources.

The 2026 Outlook

In summary, regulatory and privacy factors are a double-edged sword: they compel organizations to improve security operations (great for motivating investment), but also constrain how those operations can be implemented and add overhead. The leading organizations are approaching this proactively – for instance, many European banks spent the last 2 years upgrading their SIEM and analytics, knowing DORA was coming into force. Likewise, companies under NIS2 are performing gap assessments to ensure they have the monitoring tools and incident response plans needed to fulfill the 24-hour rule [42]. We also see cross-functional collaboration increasing: compliance officers, data privacy officers, and CISOs are working together to select solutions that satisfy both security objectives and legal requirements (for example, choosing a security analytics platform that can be hosted in-country to satisfy privacy, while still delivering on detection needs).

One positive outcome of these pressures is that they encourage standardization and documentation. Firms are cataloguing their critical systems and data flows (often required by regulation), which in turn helps identify what needs monitoring (closing those blind spots mentioned earlier). They are also defining what constitutes an “incident” that must be reported, which helps clarify detection rules and priorities. As we move into 2026, we can expect regulatory scrutiny on cybersecurity only to rise – but organizations that turn compliance into an opportunity to build a strong, privacy-conscious security analytics capability will benefit not just by avoiding fines but by better protecting their stakeholders’ data and trust.

Conclusion and outlook for 2026

As enterprises across sectors brace for the evolving cyber threat landscape, it's clear that traditional security approaches must give way to more intelligent, integrated strategies. The trends discussed – from shadow IT and insider threats to SOC overload and new regulations – all point toward a common theme: the need for comprehensive, adaptive security monitoring. Organizations in BFSI, manufacturing, healthcare, government and beyond are realizing that security is no longer just about keeping the bad guys out at the perimeter; it's about continuously observing what's happening inside, making sense of vast data in real time, and reacting swiftly to the faintest signs of trouble.

Looking ahead to 2026, several key developments are anticipated:

Wider Adoption of AI-Driven Analytics

AI and machine learning will play an even larger role in cybersecurity operations. We expect increased use of behavioral analytics not as standalone tools but embedded in broader platforms. Advanced algorithms, including those powered by emerging GenAI techniques, will help predict attacker moves and sift through anomalies faster than humanly possible. For example, AI may identify a previously unseen pattern of insider exfiltration by comparing against learned behavior across many divisions within the organization. The cat-and-mouse game will continue – attackers are also arming themselves with AI – but defensive AI offers a chance to level the playing field, especially for under-resourced teams.

Convergence of Security Tools

The era of siloed security products is waning. Enterprises will increasingly demand **smart integrations** that streamline and simplify logging, detection (SIEM/UEBA/NDR), response (SOAR), and even recovery workflows. This consolidation is driven not only by efficiency needs but also ROI – studies show that organizations consolidating and integrating security tools see significantly higher returns on security investment and faster incident handling [31][32]. One can imagine a future SOC platform where an alert automatically triggers enrichment from threat intel, checks compliance reporting templates, and suggests remediation steps – all within one interface. The big players in the industry are already moving this direction, and by 2026 many mid-to-large enterprises will likely have migrated to a more integrated security architecture. This should gradually reduce the issues of alert fatigue and blind spots, as everything feeds into one brain, so to speak.

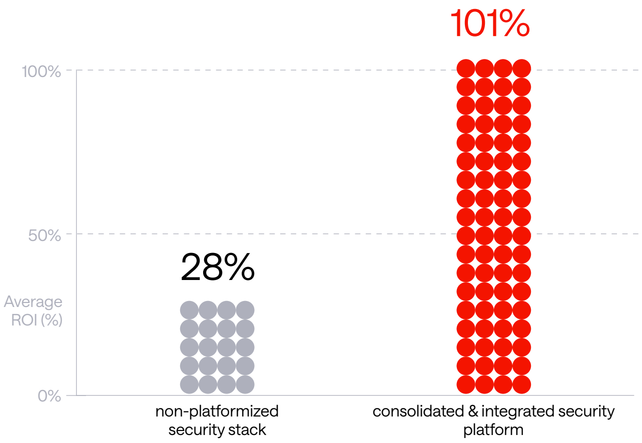

Organizations with consolidated security platforms are seeing an average ROI 4x better than non-adopters.

Zero Trust and Identity-Focused Security

With insiders and credential abuse being top threats, organizations will double down on Zero Trust principles. Identity will be the new perimeter, meaning continuous monitoring of user and service account activity is crucial. UEBA is a cornerstone of Zero Trust because it validates that each identity is behaving as expected. In practice, we'll see more deployment of identity analytics, privileged access monitoring, and conditional access controls that adapt based on real-time risk scoring of user behavior. This ties back to our discussion on missed insider threats – by 2026, many firms aim to catch abnormal user actions before they escalate (for instance, automatically challenging a user with MFA or locking an account when their behavior score spikes). This proactive stance minimizes damage from inside threats.

Compliance as Catalyst for Improvement

The new regulatory requirements (DORA, NIS2, etc.) coming into effect will reach their enforcement phase by 2025–2026. We anticipate some high-profile regulatory actions (fines or sanctions) against organizations that suffer incidents and are found non-compliant, which will serve as cautionary tales. Conversely, organizations will seek out certifications or attestations of resilience as a competitive differentiator – for example, a bank that can say it meets DORA and has state-of-the-art monitoring might inspire more trust from customers and partners. In the broader market, spending on cybersecurity is projected to continue rising (global security spend is forecast to reach $213 billion in 2025, up from $193B in 2024 [51]). The focus of that spend will likely shift toward solutions that not only prevent breaches, but also streamline compliance. Automation of reporting, audit-ready logging, and clear metrics will be selling points. Businesses will favor vendors that offer **data residency options** and can prove their tools don't run afoul of privacy laws.

Spending is estimated to increase 12.5% in 2026 to total $240 billion.

Information Security End-User Spending Worldwide 2024-2026 (Millions of U.S. Dollars)

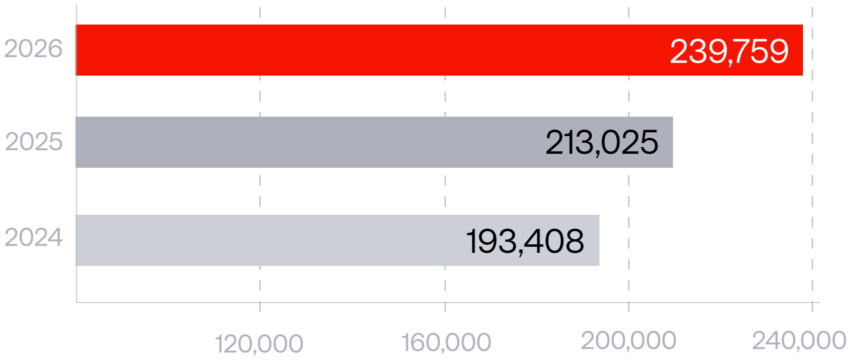

Human Element and Process

Finally, despite all the technology, the human element remains pivotal. There is growing recognition that **cultivating skilled cyber talent** and avoiding analyst burnout is part of resilience. In 2026 and beyond, expect companies to invest in their security teams' well-being – whether through better training (so they can utilize advanced tools effectively), workload management, or even mental health support given the stress levels reported. A culture of security will permeate organizations more deeply, with regular drills, company-wide awareness (so even non-IT staff understand things like shadow IT risks or how to spot insider red flags), and executive-level oversight of cyber risk (boards in DACH and elsewhere are already more engaged due to regulatory pressure [52][53]).

In conclusion

The next years are poised to be the ones of significant transformation in cybersecurity practices. Companies that proactively adapt – by shining light on formerly dark corners of their IT environment, embracing intelligent automation, and aligning with the new rulebooks – will not only thwart threats more effectively but also operate more efficiently. Those that lag may find both attackers and regulators knocking on their door.

For a CISO or Head of IT Risk reading this, the message is clear: Now is the time to close the blind spots, catch the malicious insiders, declutter your SOC, and ensure your defenses (and incident reports) are ready for whatever comes next.

The stakes in terms of financial losses, reputational damage, and compliance penalties have never been higher, but so too are the opportunities to strengthen cyber resilience through smart investment and strategic focus. By leveraging the trends and insights outlined in this report, organizations can navigate the evolving threat landscape of 2026 with confidence and control.

Sources

01. **Auvik, 50 Essential Shadow IT Statistics for 2024 – Shadow IT prevalence and risks** (unapproved SaaS apps, spending) [2][3][14].

02. **Vectra AI, 2023 State of Threat Detection Report – SOC alert volumes and blind** spots (average daily alerts, % ignored; lack of visibility concerns) [8][54].

03. **DeepStrike (Khalil, 2025), Insider Threat Statistics 2025 – Insider threat frequency and cost** (83% experience rate, $17.4M avg cost, 81-day containment, causes) [7][5].

04. **Cybersecurity Insiders, 2024 Insider Threat Report (via Gurucul) – Increase in insider incidents** (48% saw rise in last year)[6] and complexity as a driver [24].

05. **Trend Micro Survey (via IT Europa, 2021) – SOC overload stats** (51% teams overwhelmed by alerts, 55% not confident in response, 27% time on false positives) [9].

06. **IBM Institute for Business Value, Security Platformization Report (2024) – Tool sprawl and complexity** (83 security solutions on avg; 52% say complexity impedes ops; 74% say workload excessive; benefits of consolidation) [10][33].

07. **Observo Blog (Turriff, 2023), Cut SIEM Costs & Reduce Noisy Data – Log value and blind spots** (80% of logs low value; dropping high-fidelity sources due to cost creates blind spots) [17][55][16].

08. **Check Point Research, Q1 2025 Cyber Attack Report – Threat surge data** (47% YoY increase in attacks per org, 1925 weekly avg) [1].

09. **InnReg, DORA Regulation Explained (2023) – DORA requirements** (Jan 2025 in force; incident reporting within hours; need robust monitoring) [37][11].

10. **Snare Solutions, Review of NIS2 Directive (2023) – NIS2 requirements** (24-hour incident notification vs 72h; mandate for better detection tools like SIEM) [56][12][43].

11. **FutureMarketReport, UEBA Software Market Outlook 2032 – UEBA trends** (cloud vs on-prem adoption, privacy-preserving analytics for compliance, vertical-specific needs)[57][36].

12. **ManageEngine, Evolution of SIEM in 2025 (2025) – Modern SIEM capabilities** (AI-driven analytics, integration of UEBA for advanced threat detection) [58][59].

13. **Check Point, 2025 Cyber Security Report – Global spending and priorities** (Projected $213B security spend in 2025; AI and unified platforms emerging – referenced in summary) [51].

14. **Wing Security, Countering AI Security Threats in SaaS Environments by Lia Ciner,** Oct 2023 [60].

exeon