eBook

THALES

CYBERSECURITY

Strategic Guide: Accelerate Innovation with Security Foundations

cpl.thalesgroup.com

Contents

This e-book is designed to is to help CISOs to proactively shape business growth through security by continuously assessing emerging trends and making informed decisions whether to adopt, monitor, or deprioritize them— while the organizations strengthens its security foundations and risk posture —to sustain trust in an evolving threat landscape.

Introduction

The Gartner 2026 Cybersecurity Trends report reinforces what Thales sees across global enterprises: cybersecurity has become a strategic enabler of digital trust, not just a defensive function.

As forces like GenAI, post-quantum risk, regulatory pressure, and identity sprawl reshape the threat landscape, CISOs must move from reactive controls to proactive, systemic risk management.

Thales aligns with Gartner's findings—organizations must secure data, identities, and cryptographic foundations at scale, embedding trust into how technology evolves rather than relying on human vigilance alone. This shift is essential for enabling innovation while maintaining resilience, accountability, and long-term confidence in the digital enterprise.

From this e-book, you will gain clear, actionable guidance on how to translate Gartner's 2026 cybersecurity trends into practical decisions—helping you prioritize risk, modernize controls, and build scalable trust across data, identities, AI, and emerging technologies.

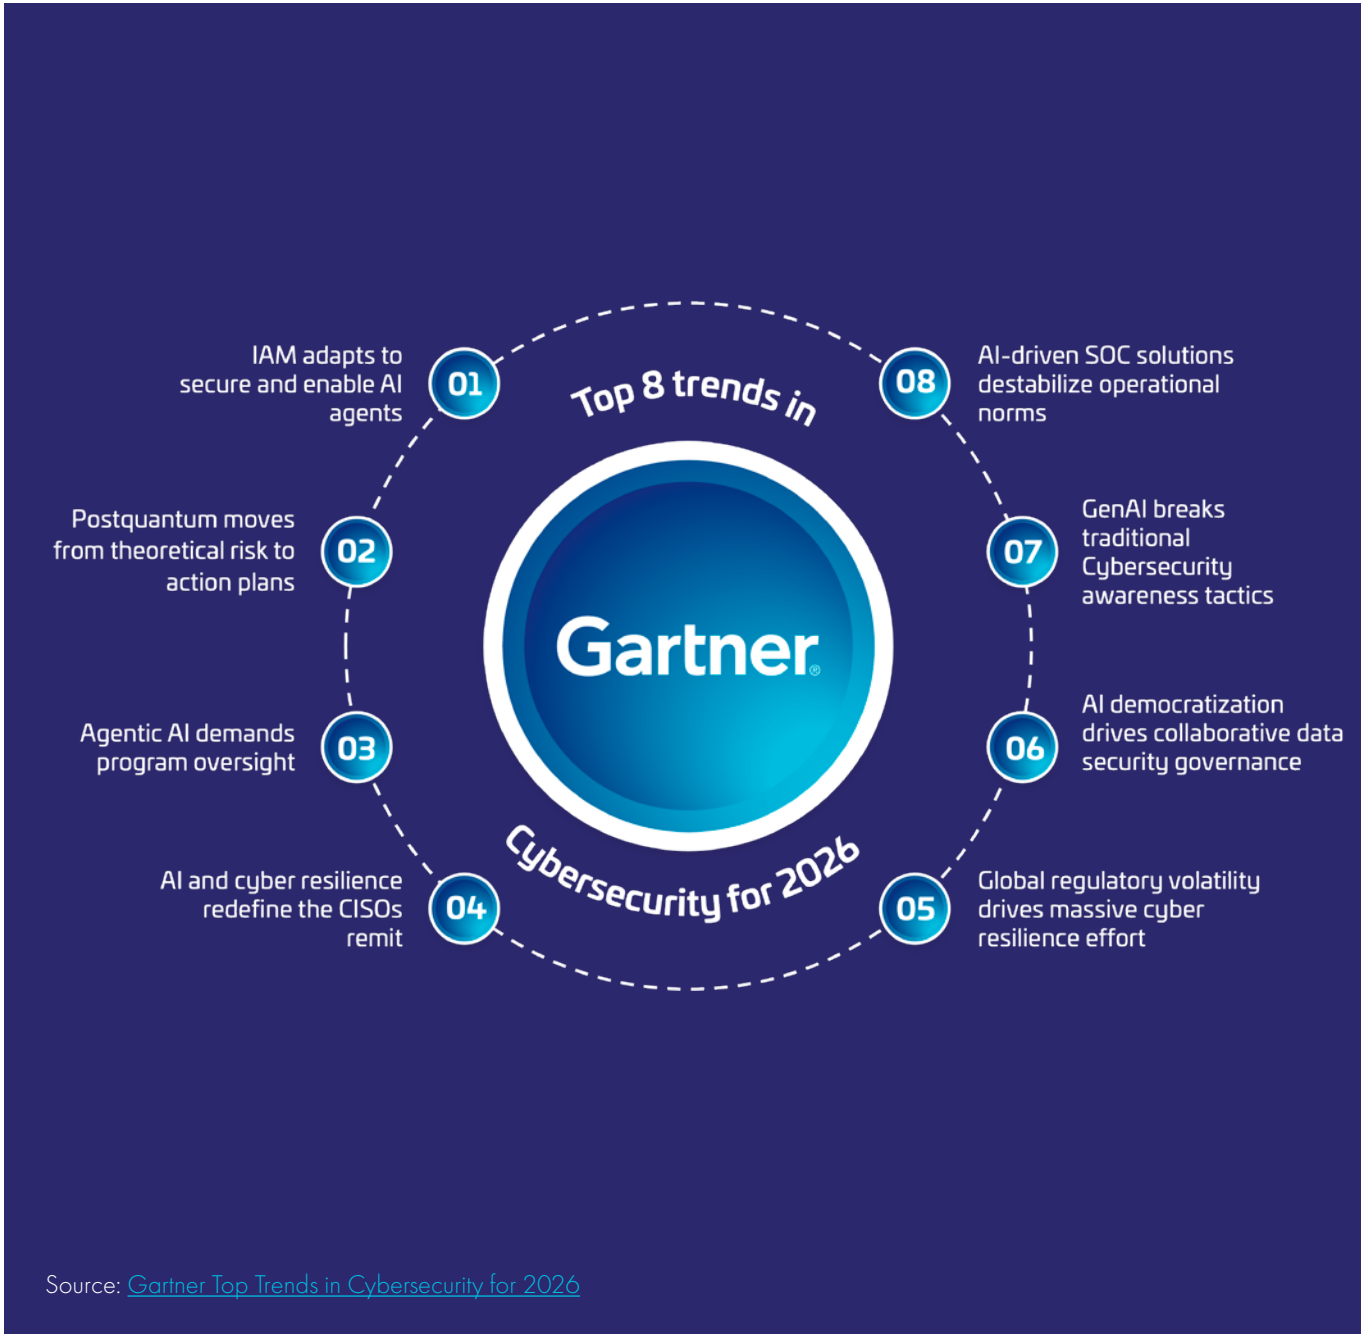

Source: Gartner Top Trends in Cybersecurity for 2026

IAM Adapts to Secure and Enable AI Agents

Protect against synthetic and non-human identities.

GenAI-enabled attacks are highly contextual, adaptive, and indistinguishable from legitimate interactions by removing clear indicators of malicious activity, introducing autonomous AI actors that operate beyond direct human control, and exposing sensitive data regardless of user intent.

At the same time, embedded AI assistants and automated workflows now retrieve, process, and act on sensitive enterprise data without explicit human decisions. As a result, security can no longer rely on trained security staff and application owners to detect malicious behavior.

Adapt IAM to secure and enable AI agents
Protect agentic AI and LLM-powered applications, enterprise data, and identities from emerging AI-specific threats, including prompt injection, data leakage, model manipulation, and insecure RAG pipelines.

Move protection closer to the data.
When attacks are adaptive, contextual, and indistinguishable from legitimate activity—and when AI agents access and act on sensitive data without explicit human intent—security can no longer depend on users recognizing red flags. Instead, protection must move closer to the data and the AI agents themselves.

Deploy a targeted, risk-based security model.
Govern GenAI through data-centric protection, identity-driven controls, and continuously enforced policies aligned to data sensitivity, agent autonomy, and business impact.

Top OWASP Risks for LLMs

Prompt Injection
User prompts that alter the LLM's intended behavior or output

Sensitive Information Disclosure
Sensitive data can be exposed through the LLM output

Data and Model Poisoning
Manipulated data introduces vulnerabilities, backdoors, or biases

Unbounded Consumption
Malicious AI resource use and denial of service attacks that drive up costs

Improper Output Handling
Insufficient validation, sanitization, and handling of LMM outputs before moved to other systems

System Prompt Leakage
LMM prompts or instructions can steer the LMM behavior to reveal sensitive information

Vector & Embedding Weakness
Generation, storage, and retrieval creates vulnerabilities for all the type of AI attacks

The Challenge

GenAI introduces highly convincing interactions, non-human actors, and autonomous workflows that bypass traditional awareness controls. AI agents can access, process, and act on sensitive data at machine speed, creating risk that cannot be mitigated through user training alone and often outpaces static security controls.

Strategy

Adopt a risk-based security layer for AI ecosystems, enabling enterprises to scale AI adoption across cloud and on-premises environments by:

- Classifying data and aligning protection to sensitivity and regulatory impact
- Assigning strong, verifiable identities to AI agents and enforcing least-privilege access
- Applying policy-driven controls with agent autonomy and intended function
- Continuously monitoring access, behavior, and outcomes across human and AI interactions

Outcome

Maximize AI's business value while mitigating risks.

Post Quantum Moves From Theoretical Risk to Action Plans

Preparations can't wait.

Post-quantum security is no longer a distant, theoretical risk—it is a present-day planning imperative. Advances in quantum computing, combined with long data lifecycles and "harvest now, decrypt later" threats, mean organizations must begin transitioning cryptographic foundations now to protect future confidentiality and trust.

The shift in regulatory posture confirms that now is the time. When regulators and standards bodies such as NIST, NSA, and the EU require concrete post-quantum cryptography (PQC) migration plans, quantum readiness becomes a board-level responsibility tied to risk management, compliance, and long-term trust.

Embed post-quantum readiness into security foundations. Organizations must move from awareness to action by building crypto-agility, gaining visibility into cryptographic dependencies, and embedding post-quantum readiness into their security foundations.

Act now to build crypto-agility to reduce future complications. Get started inventorying cryptographic assets, prioritizing high-value data and long-lived secrets, as well as adopting crypto-agility, and integrating standardized post-quantum algorithms as they emerge.

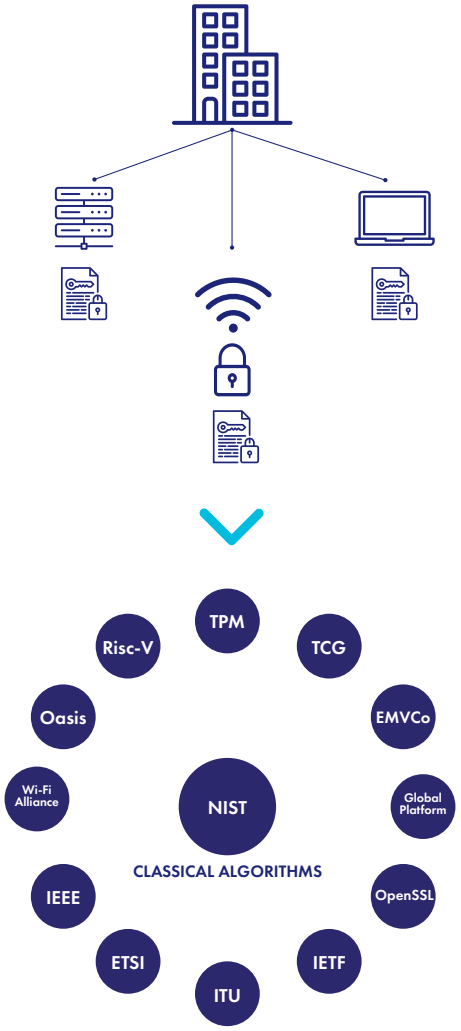

The Challenge

Post-quantum risk as a timing and complexity challenge—not a distant theory. Organizations face long-lived data, deeply embedded cryptography, and limited visibility into where vulnerable algorithms are used, all while adversaries actively harvest data today for future decryption.

Strategy

Begin PQC migration now by identifying quantum-vulnerable cryptography, prioritizing high-risk and long-lived data, and preparing for hybrid environments using both classical and NIST-standardized PQC algorithms.

Discover and classify cryptographic assets, prioritize protection for sensitive and long-term data, and integrating standardized post-quantum algorithms into key management and data protection.

Apply quantum key generation and leverage a quantum random number generator to generate quantum-enhanced keys.

Implement quantum resistant algorithms. Agencies need to ensure that their encryption solutions are using the standardized PQC algorithms—or are crypto-agile and have an upgrade path to utilize PQC algorithms in the near future.

Outcome

Reduce future migration risk, meet emerging regulatory expectations, and ensure cryptographic resilience.

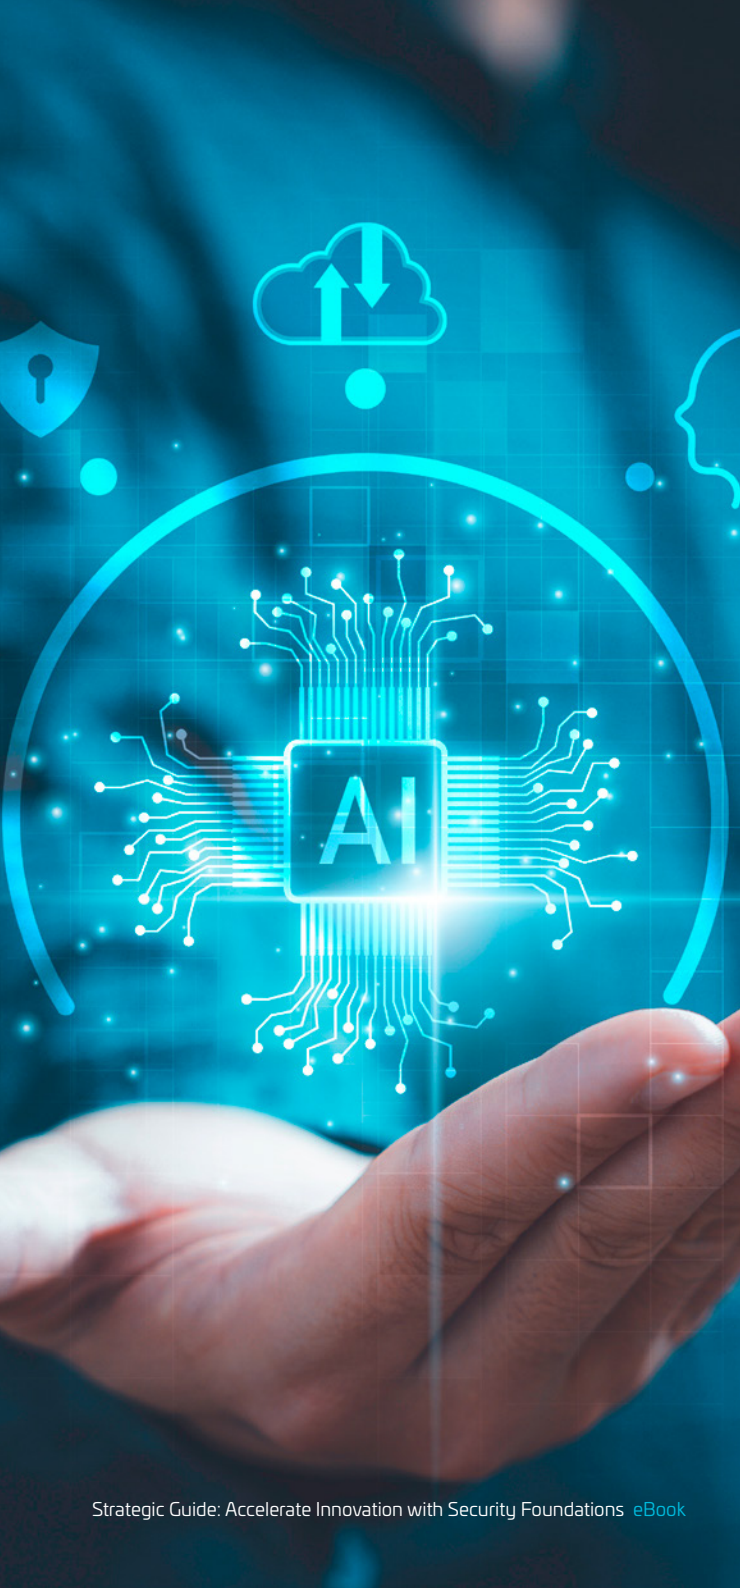

Agentic AI Demands Program Oversight

Develop and operate within enforceable AI guardrails.

The rapid, unsanctioned creation of agentic AI through no-code and low-code platforms is creating a new class of rogue automation. In this environment, agentic AI can turn access mistakes into high-speed business actions—where an unauthorized agent activity becomes an immediate operational cost. Budgets are drained, sensitive data is exposed, and incorrect workflows or decisions are executed before a human ever sees the damage.

Treat agents as first-class actors.
AI agents operate as delegated, non-human identities. They rely on machine-to-machine credentials, APIs, tools, and workflow permissions to act on behalf of the business.

Human > Agent > Scoped Token > Tool/API > Data > Action
As a result, governance must treat agents as first-class actors within the enterprise security model—ensuring that every agent action is clearly attributable, explicitly authorized, strictly bounded, continuously monitored, and quickly interruptible.

Secure agentic AI without creating a new identity silo.
When autonomous agents proliferate without oversight, they expand the attack surface and operate with excessive privilege beyond the visibility of security teams. The answer is not to create a new identity silo for AI, but to integrate agent governance into existing identity stacks, secrets management systems, API gateways, and CI/CD identity flows—so agent access is consistently visible, governed, and enforced.

The Challenge

AI agents can combine prompts, tools, APIs, data access, and static credentials into action chains that legacy IAM cannot fully govern. Without runtime authorization, identity sprawl becomes operational risk.

Strategy

Enable secure agentic AI with identity-aware access control, strong key and secrets protection, and policy enforcement that discovers shadow agents, binds actions to business context, and replaces static keys with short-lived tokens.

Outcome

Reduce agent identity sprawl and contain the blast radius of autonomous AI actions.

AI and Cyber Resilience Redefine the CISO's Remit

Prove AI-enabled operations can survive incident, audit, and regulatory scrutiny.

The CISO's remit is undergoing a fundamental shift—from protecting systems and data to ensuring the resilience, trustworthiness, and continuity of AI-driven business operations.

Manage risk across the AI decision chain—not just the infrastructure. As AI becomes embedded in decision-making, risk now flows through decision chains, not just infrastructure, with models, copilots, agents, APIs, data stores, identity systems, and automation workflows all participating directly in business execution. The CISO must define the accountability model and ensure evidence is automatically collected, normalized, retained, and mapped across identity, data, AI runtime, engineering, and business systems.

Extend the CISO's role to secure and govern AI at scale. To securely scale AI, the CISO's remit must expand beyond traditional infrastructure protection to include AI governance, runtime protection, data integrity, identity, and cryptographic trust, ensuring that AI systems are secure by design, continuously verifiable, and resilient against both cyber threats and operational failures.

Provide clear evidence of oversight of AI across the lifecycle. Cyber resilience becomes a proof problem. The enterprise must be able to reconstruct AI-driven events, contain delegated machine actions, prove what data was exposed, and restore confidence with regulators, auditors, the board, and business owners.

The Challenge

As AI systems become autonomous and embedded in critical business processes, security must evolve from reactive and siloed, to proactive and business-aligned, providing continuous oversight of AI behavior and outcomes to support audit, investigation, regulatory confidence, and operational resilience without requiring the CISO to manually own every workflow.

Strategy

Establish AI-era regulatory resilience program by unifying security, identity, data protection, runtime controls, and accountability—enabling continuous evidence collection, decision-chain visibility, blast-radius containment, and regulatory-ready proof across the full AI lifecycle.

Outcome

Improve audit survivability by proving a verifiable trail to restore and maintain confidence with regulators, auditors, the board, and business leaders.

Global Regulatory Volatility Drives Massive Cyber Resilience Efforts

Maintain operational agility in an evolving regulatory landscape.

Mandates for cyber resilience, privacy, and sovereignty are in large part based on the same standards for cyber security best practices such as NIST Cyber Security framework 2 and ISO 27001. Strong centralized and automated regulation enforcement is key to achieve compliance, by enforcing common rules (and managing rule exceptions), seamlessly across hybrid environments.

Reduce complexity and increase security effectiveness Organizations have an average of 7 data protection solutions[1], often siloed and lacking integration. Consolidating fragmented tools into a single unified ecosystem or platform can simplify operations, strengthen protection, and improve visibility and control.

Understand, prioritize, and act on risk. Only 34% of organizations have complete knowledge of where their data is stored[1]. You can't protect what you can't find. Organizations need to understand which assets are at risk, prioritize remediation, and enforce protection across hybrid IT at a global scale.

Automate enforcement and leverage AI. It is impossible for hard-pressed security personnel to effectively enforce regulation across environments at a global scale. Policy enforcement has to be automated, and AI used to assist through the process, from data discovery to real-time monitoring of threats.

Detect and mitigate threats. Finally organizations need to detect and prevent cyber threats, ensuring seamless operations and peace of mind. It is critical to safeguard critical applications from DDoS or Bad Bots attacks while continuing to allow legitimate traffic.

[1] Thales Data Threat Report, 2026

The Challenge

Efforts by regulators to improve cyber resilience have produced a growing number of mandates such as DORA, NIS 2, CRA, and many others under development globally. On top of that, geopolitical tension has increased the focus on digital sovereignty and exacerbated the challenges for organizations trying to maintain compliance in a volatile environment.

Strategy

Replace isolated tools with a single solution to centralize data security, combining functions such as data discovery and classification, risk assessment, policy definition and enforcement, key management, encryption, and data activity monitoring. Detect and prevent cyber threats in seconds, maintaining operational effectiveness of key business applications.

Outcome

Improved resilience, business continuity, operational efficiency, and mitigate organizational liability amid regulatory change.

AI Democratization Drives Collaborative Data Security Governance

Unlock AI-driven productivity at scale.

AI democratization moves sensitive data into more hands, more tools, and more automated workflows. Business users, developers, analysts, low-code teams, and AI builders can now retrieve, summarize, transform, and act on enterprise data outside traditional application paths.

Strengthen visibility across the data estate. Without visibility creates systemic risk. Organizations must maintain continuous visibility into where sensitive data resides, how it is accessed, and how it is used by both humans and AI—across cloud, SaaS, on-premises environments, data lakes, and AI platforms. Security teams must be able to govern data use and hold business teams accountable for how AI consumes data.

Answer the question: What data can this user, agent, model, or workflow reach, infer, combine, expose, or act on? This question reveals a core weakness in many RAG and agentic AI deployments. A user or agent may be authorized to access multiple datasets that appear low risk in isolation, but become high risk when combined. These toxic data combinations create inference risk, privacy violations, and regulatory exposure that traditional access controls fail to detect or prevent.

Enforce AI-system governance to assure accountability. Approval and accountability depend on enforceable—not just documented—governance. Organizations must be able to discover the data, classify the risk, enforce policy, and block or approve the action.

The Challenge

Enterprises have long failed at comprehensive data classification—and AI raises the stakes by exposing sensitive data through prompts, vectors, copilots, agents, and shadow workflows.

Strategy

Implement DSPM-led data security governance by: Continuously discover and classify sensitive data wherever it resides, map AI access to data sensitivity, ownership, consent, and regulatory impact; enforce encryption and policy-based controls that give security and data owners the ability to block, quarantine, or approve access, as well as monitor prompt, retrieval, agent, and output access.

Outcome

Accelerate responsible AI adoption by shifting to collaborative, business-led governance that embeds cybersecurity early, reduces unnecessary friction, and increases accountability.

GenAI Breaks Traditional Cybersecurity Awareness Tactics

Shift from awareness-led to control-driven security.

GenAI fundamentally breaks cybersecurity awareness models by removing the signals humans rely on to identify malicious activity. The scale, speed, and sophistication enabled by GenAI create lapses in judgment and can no longer reliably distinguish safe from unsafe interactions.

Redefine how threats go unnoticed. GenAI removes obvious threat signals, operates autonomously, and routinely acts on sensitive data without direct user intervention. By eliminating the very signals those programs rely on, GenAI introduces non-human actors that operate beyond direct user control, and exposes sensitive data regardless of user intent— enabling highly fluent, contextual, and adaptive attacks that no longer look suspicious.

Mitigate the new risk surface by governing AI execution. AI assistants and automated workflows now act on behalf of users—retrieving data, making decisions, and triggering actions, shifting risk from human error to uncontrolled access and execution. By enforcing governance at execution, organizations can contain AI-driven risk and preserve accountability.

Operationalize trust across all systems. Trust must be enforced through data-centric controls, identity governance, and systemic security mechanisms that limit risk regardless of human or AI behavior. This ensures GenAI systems operate with least privilege and within defined governance boundaries, independent of user intent.

Source: 2026 Thales Data Threat Report

The Challenge

GenAI produces highly convincing, contextual interactions and introduces non-human actors that bypass human judgment entirely. Awareness alone cannot scale to manage AI-driven risk. Trained users cannot prevent oversharing, data leakage, or unintended data retention once GenAI systems are granted access.

Strategy

Shift from awareness-led security to control-driven security by protecting data at its source, enforcing identity-based access for both humans and AI, and applying continuous, policy-based controls.

Outcome

Safely enable GenAI at scale while reducing reliance on fallible human judgment, limiting data exposure risk, preserving accountability, and maintaining trust as automation and autonomy increase.

AI-Driven SOC Solutions Destabilize Operational Norms

Govern AI in the SOC as critical infrastructure.

The disruption seen in AI-driven security operations stems from a fundamental misunderstanding: AI is being treated as a tool, when it must be governed as critical infrastructure.

Anchor AI decisions in trust. Every AI decision—from alert prioritization to automated response—must be cryptographically secured, identity-bound, and policy-controlled. Strong identities for models and agents, integrity protection for training data and outputs, and verifiable evidence of where an action, decision, or output came from and how it was produced. Without this, AI becomes an unaccountable authority rather than a trusted partner.

Accelerate response but not replace accountability. AI-driven SOCs are not "human-out-of-the-loop" systems; they are human-on-the-decision-boundary systems. High-impact actions must include enforced approvals, role-based limits on autonomy, and tamper-proof audit trails linking AI recommendations to human decisions. This ensures regulatory accountability, forensic integrity, and executive confidence. AI may accelerate response—but it must never obscure responsibility.

Treat AI itself as a high-value asset. As AI becomes central to SOC operations, it becomes a prime target for attack, including data poisoning, model manipulation, and privilege escalation. AI security is inseparable from data security. AI will redefine the SOC however without governance, it erodes trust.

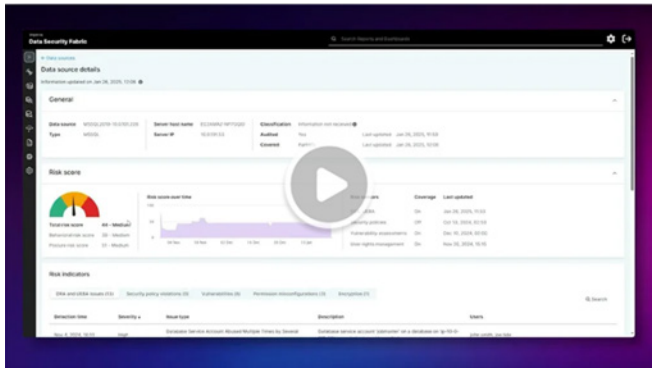

Day in the Life of a SOC Manager: Prioritizing high-risk incidents to boost SOC efficiency.

The Challenge

To distinguish the noise from the real threats—because lost amid the chaos of false positives, the genuine incidents that could impact security and the business lurk in the shadows, waiting to be uncovered.

Strategy

Move faster against attacks while preserving regulatory compliance and forensic integrity by streamlining operations, prioritizing high-risk data incidents, and deploying cutting-edge threat mitigation techniques through AI- and data security posture management, multi-dimensional access analytics and automated incident remediation.

Protecting AI requires encryption, strong key management, access controls, and continuous access monitoring across the AI lifecycle.

Outcome

Deliver speed and scale without surrendering control, accountability, or confidence.

Why the most trusted organizations in the world rely on Thales?

Security for what matters most.

92%
of the worlds 100 largest bank rely on our solutions

6,000
cybersecurity experts and developers

113 Billion
application attacks blocked every month

184
the number of countries we monitor for threats everyday

$150 Trillion
the value of interbank money transfer we protect every year

8
Security Operation Centres (SOC)

From AI to quantum computing, Thales helps organizations embed security, from design and training to deployment and use.

Next Steps

Download Thales Sponsored Reports

Visit Thales Website

Read Thales Data Security Product Collateral

THALES

CYBERSECURITY

Contact us

For contact information, please visit
cpl.thalesgroup.com/contact-us

cpl.thalesgroup.com