June 2026
White Paper
Sponsored by Thales

Unlocking AI with Secure, Protected Data



Jennifer Glenn
Research Director,
Security and Trust, IDC

Table of contents



Click any title to navigate directly to that page.

Executive summary 3

Situation overview 4

- Lack of visibility creates multiple challenges when securing AI implementations 7
- A unified and proactive security approach is essential for effective AI 11
  - Establish provenance and gain visibility 11
  - Implement fine-grained, identity-aware controls 11
  - Maintaining lineage with data security at runtime 12
  - Monitoring output and adjust policy and/or enforcement controls 12

Considering a unified data security framework 14

Business impact of a unified data security framework 16

Considering Thales AI security fabric 18

Challenges 20

Conclusion 22

About the IDC analyst 23

Message from the sponsor 24

Executive summary:

Data control as a strategic enabler of enterprise AI

Organizations are transitioning from experimental AI initiatives to enterprisewide deployment. In IDC's January 2026 *Future Enterprise Resiliency and Spending Survey,* nearly 44% of respondents plan to spend $2 million or more on AI efforts in 2026. This number jumps to 64% in 2027. This expenditure includes infrastructure, models, applications, and services. As this transition occurs, data security is a primary determinant of AI success and also the main impediment to rapid adoption.

AI systems require continuous access to diverse data sources across hybrid and multicloud environments. However, traditional security models, often fragmented and reactive, are not designed to support the speed, scale, and autonomy of modern AI systems.

IDC analysis shows that organizations are increasingly prioritizing data control capabilities, including data lineage, governance, and real-time classification, as foundational elements of AI strategy.

To address these requirements, organizations are adopting more integrated approaches that align data security, governance, and data intelligence. These approaches focus on ensuring that the right data is accessible to the right users and systems, under the right conditions, without introducing unnecessary risk.

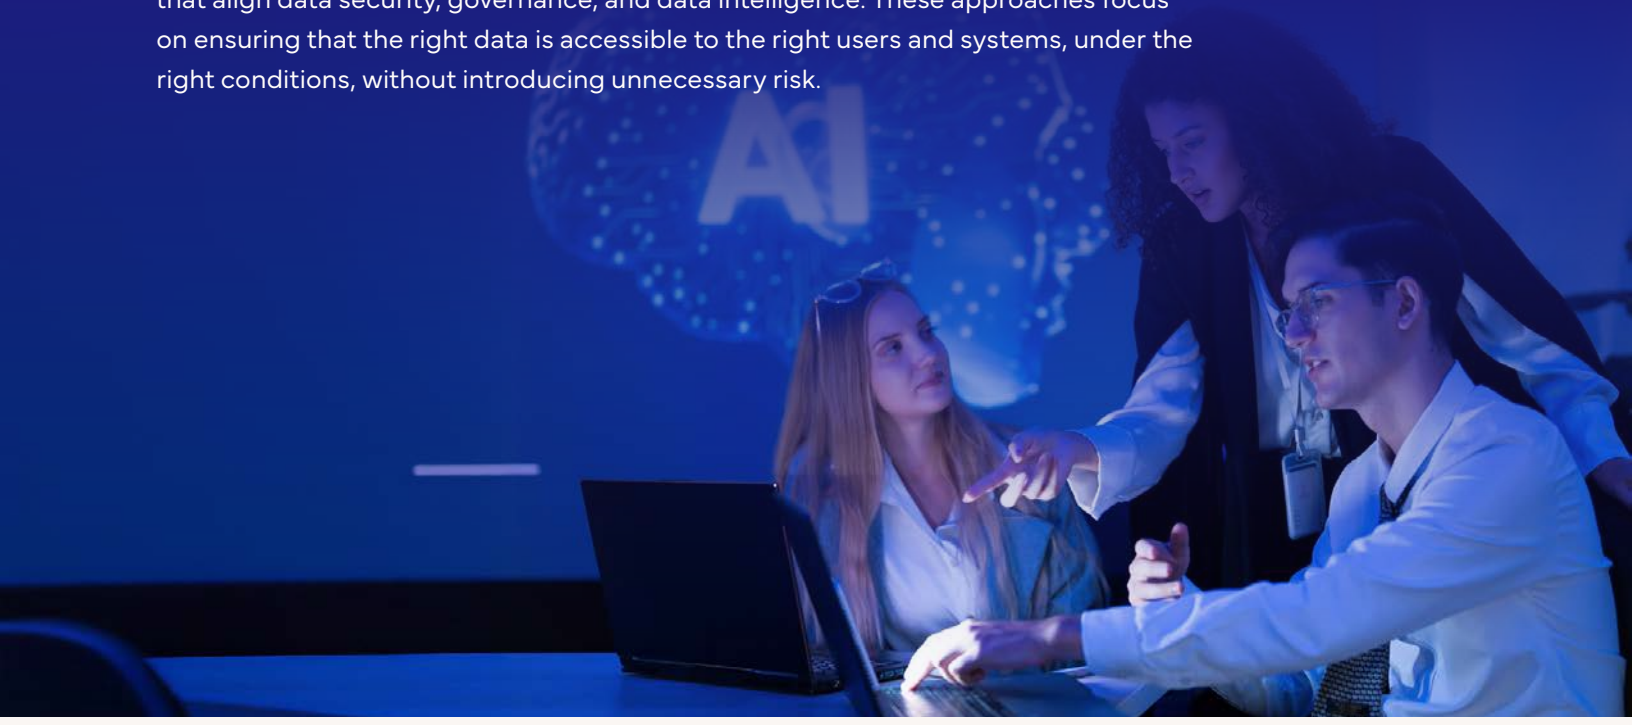



Situation overview

Organizations are rapidly investing in AI to drive innovation and business value, but many struggle to realize its full potential due to persistent challenges around data security, governance, and compliance.

High-quality, secure, and well-governed data is foundational for successful AI outcomes, yet most organizations face significant hurdles in ensuring sensitive information remains protected throughout the AI life cycle.

Key challenges include the risk of exposing personally identifiable information, intellectual property, and other confidential data through AI systems, especially with vulnerabilities such as prompt leakage and model poisoning. Many organizations lack visibility into sensitive data before it is ingested into vector databases or used for model training, creating long-term exposure risks. Fragmented security controls,

operational complexity, and inconsistent policy enforcement further exacerbate these issues, while the proliferation of AI agents and non-human identities complicates traditional access control models and expands the attack surface.

To address these challenges, a unified and proactive approach to data security is essential. This begins with establishing data provenance and visibility: knowing the source and status of data before it enters AI workflows. Organizations must implement robust data hygiene, discovery, and classification practices and apply tokenization or encryption prior to ingestion. Secure training and fine-tuning processes are critical to prevent model memorization of sensitive data.

At a glance:



44%
of enterprises plan to spend $2 million or more on AI efforts in 2026.



34%
of CDOs say ensuring security and privacy of data is a top objective.



Leading these security concerns are limited confidence in the security of data infrastructure and meeting regulatory requirements.

Fine-grained, identity-aware data controls are necessary to balance the speed of AI with the need for controlled access. Defining data ownership, distinguishing between sanctioned and unsanctioned AI agents, and enforcing policies that prevent unauthorized access or exposure of sensitive data are vital steps. Maintaining data lineage and visibility at runtime ensures that organizations can trace AI outputs back to their sources, supporting compliance and rapid issue resolution. According to IDC's *chief development officers (CDO) Survey,* 33.8% of organizations list ensuring the security and privacy of data as one of the top 3 objectives in the next 12–18 months. Similarly, 22.4% of organizations identified improving regulatory data compliance reporting, and 15.3% of organizations cited improving access to and governance of unstructured data as near-term objectives for organizations **(Figure 1, next page).**

Figure 1

Data security, governance, and compliance are critical objectives for organizations in the near term.

What are the top 3 strategic objectives for data leadership in the next 12–18 months for the use of data and analytics in the organization?

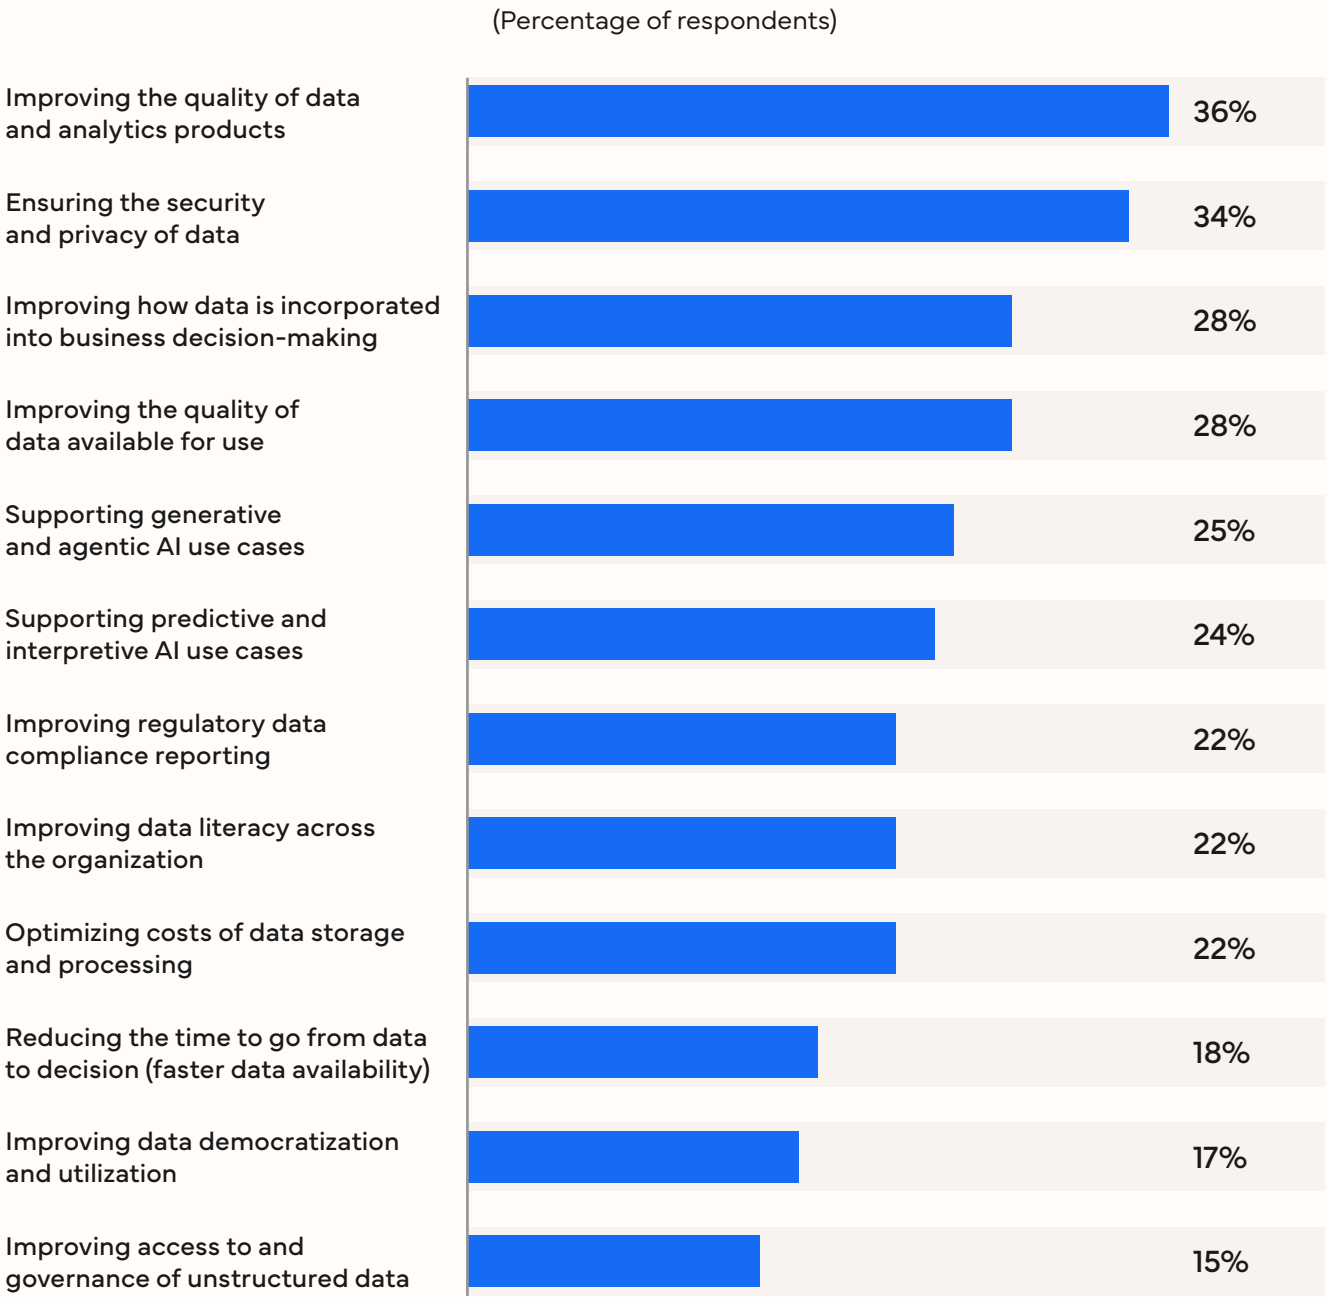

Multiple responses were allowed; n = 417; Source: IDC's *Office of the CDO Survey,* November 2025

A unified data security framework integrating discovery, classification, encryption, identity-based access controls, and runtime enforcement offers a platform-based solution that reduces operational complexity and aligns AI innovation with enterprise risk management. Such an approach accelerates AI innovation, improves output quality, reduces data breach risks, and strengthens compliance and sovereignty postures.

In conclusion, AI is now a critical enterprise initiative, and data security must evolve to enable, not hinder, its adoption. By protecting data before it reaches large language models and maintaining control throughout the AI life cycle, organizations can unlock the full value of AI while preserving trust, compliance, and resilience. A unified data security framework empowers organizations to confidently deliver the right data to the right users at the right time, ensuring AI success without compromising security.

Lack of visibility creates multiple challenges when securing AI implementations

Data is the pivot point between successful and unsuccessful AI implementations. In most organizations, data volume is increasing exponentially and is being stored and shared across endpoints, emails, applications, and infrastructure. This leads to a lack of visibility about where data exists, the sensitivity of the data, who has access, how it is being used, and if it is appropriately protected.

If data is not clearly organized, mapped to a workflow, or monitored, it can introduce a number of risks to the organization, including:

- **Sensitive data exposure:**
  By far, one of the most immediate risks associated with AI adoption is the potential exposure of sensitive data, including personally identifiable information, intellectual property, and confidential business information. Sensitive data exposure and loss of sensitive data are two significant inhibitors of AI implementation **(Figure 2, page 10).**

  AI systems can inadvertently expose sensitive data at various points throughout the AI process. These risks are exacerbated by the probabilistic nature of AI models, which may reproduce or infer sensitive information in unintended ways.

Some of these unintended ways are:

- **Data sourcing and training:**
  Too much data can obfuscate sensitive data. Further, if not identified and classified appropriately, sensitive or confidential data may be ingested into the training model. Once embedded in vector databases or models, this data becomes more difficult to track and control, increasing long-term risk.

- **Prompt misuse:**
  Data that is not properly classified or governed may be susceptible to inappropriate user prompts. For example, trusted employees may inadvertently gain access to sensitive or confidential data through poorly written prompts. On the other side, malicious users or attackers may be able to inject code into a prompt to query a database that should not be accessible.

- **Sensitive data creation:**
  The ability to dynamically assemble and contextualize data using AI models can increase the risk of unintended disclosure, particularly when data is not adequately classified or governed prior to use. For example, data used in AI models, whether in Model Context Protocol (MCP) or retrieval-augmented generation (RAG), may be benign when considered in isolation. However, RAG architectures can dynamically retrieve and combine data from multiple sources, potentially surfacing sensitive information in new contexts. Similarly, MCP-based interactions may aggregate multiple nonsensitive data elements, such as a name and address, resulting in the creation of a sensitive data record.

- **Integration with external tools or APIs:**
  Improperly configured APIs may inadvertently share or expose sensitive or confidential data to the wrong user or application. Attackers or malicious users may be able to manipulate a prompt that will bypass well-intended protections, allowing access to inappropriate information.

- **Fragmented security controls:**
  The challenges with data security are not unique to AI or AI agents, though they are making the problem worse. Many organizations have invested heavily in multiple data security tools to address different issues that have arisen over the past decade, including compliance, new attacks, and insider risks. Often, these tools are purchased and managed separately, resulting in inconsistent policy enforcement, siloed visibility, operational complexity, and integration challenges. The nature of AI and AI agents is one of purpose: The tool wants to provide the right answer and will go to great lengths to get it. Any workarounds are likely to exploit the individual policies and enforcement capabilities of multiple tools.

- **Governance, compliance, and data sprawl:**
  As mentioned previously, enterprises are managing rapidly growing volumes of structured and unstructured data. Excessive, outdated, and/or nonvaluable data increases risk and reduces AI value. Without clear purpose, provenance, or data lineage, organizations do not have the visibility to trace AI outputs back to model training or source data. This adds complexity and time to the audit process and also puts the organization at risk for potential privacy, sovereignty, or data security compliance violations.

  Further, AI introduces nonhuman identities (agents, APIs, tools) that require governance. Data ownership and metadata management have become central to accountability. Rapid adoption of autonomous AI agents (copilots, chatbots) in collaboration platforms creates ungoverned attack surfaces. Existing access control models become ineffective in agent-driven workflows. Agents can also unintentionally expose weaknesses in data governance and monitoring systems.

Unstructured data: The hidden risk

In IDC's December 2025 *State of Content and Knowledge Survey,* 88.5% of respondents cited using retrieval-augmented generation for any content or unstructured data AI use cases. The challenge for most organizations is ensuring unstructured data is well protected.

Unstructured data, such as emails, images, and text documents, lacks a predefined model or format, making it fundamentally harder to govern than structured data. The scale of the challenge is significant: Unstructured data accounts for more than 90% of enterprise data. Compounding this, unstructured data is distributed across file shares, collaboration platforms, cloud storage, endpoints, and archives. This sprawl stretches discovery and classification efforts to their limits, hindering the visibility of sensitive data.



Figure 2

Data is a primary theme across security challenges limiting AI implementation

Which of the following data security concerns, if any, limit or slow your organization from implementing AI and agentic AI?

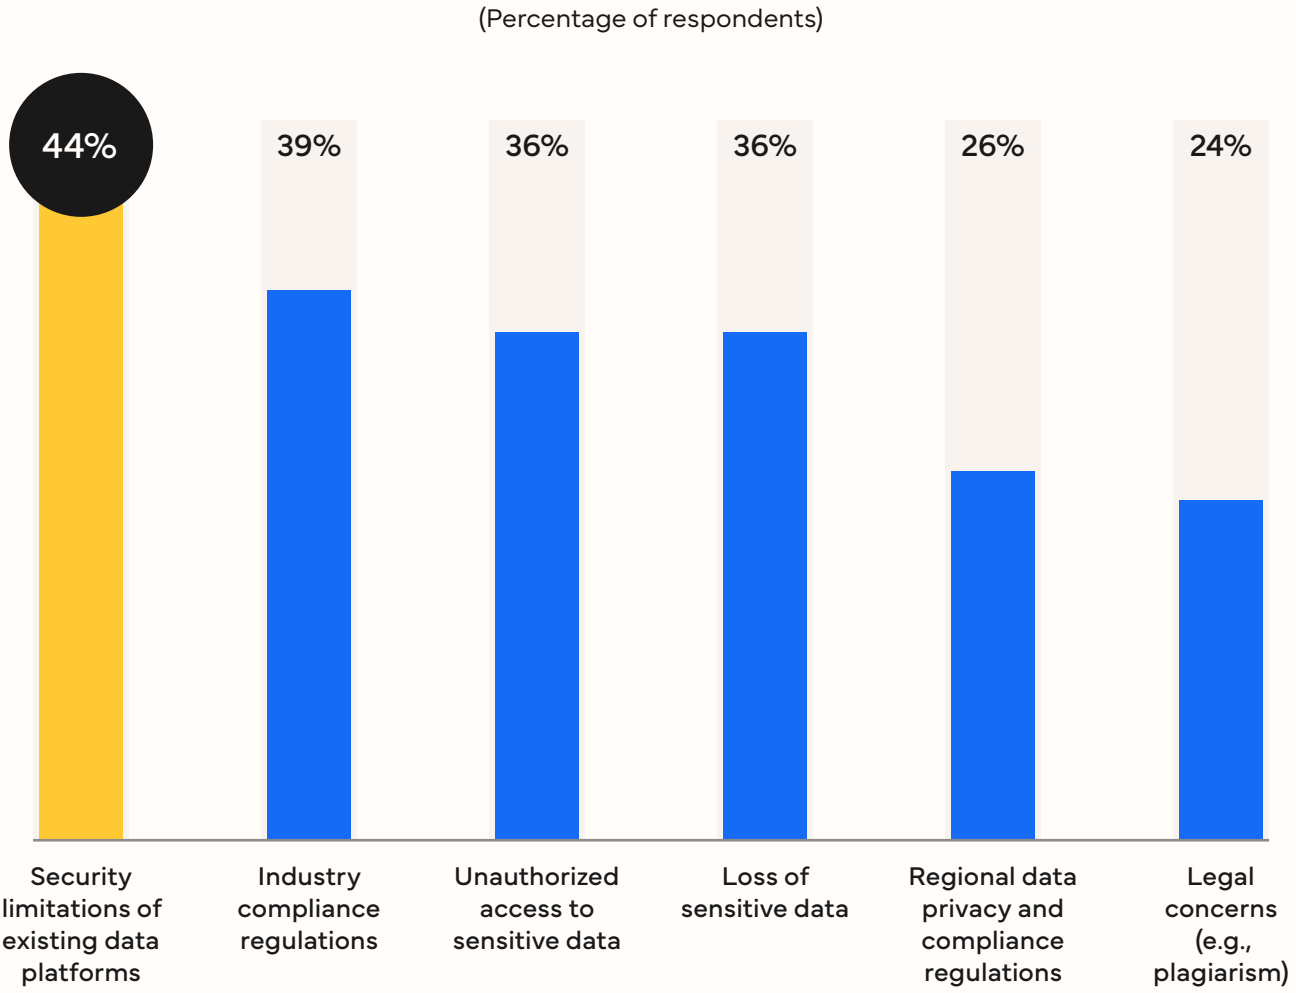

n = 107; Source: IDC's *IT Data Management Quick Poll*, Q1, February 2026

Traditional security approaches are often reactive, perimeter based, and system centric. They are not designed for distributed, hybrid, and multicloud environments; dynamic data flows across AI pipelines; and/or real-time interactions between agents, users, and systems. As a result, organizations relying solely on legacy approaches struggle to balance innovation with risk management.

A unified and proactive security approach is essential for effective AI

To address these challenges, organizations must move beyond reactive security models and adopt a unified, proactive approach. High-quality, secure data is foundational to AI success, and achieving this requires embedding security and governance controls earlier in the data life cycle.

This approach emphasizes protecting data before it enters AI systems rather than attempting to mitigate risk after exposure. By integrating security into the early stages of data preparation and ingestion, organizations can reduce risk while maintaining the speed required for AI innovation.

The sections that follow outline how to build and maintain optimal data and AI practices for successful outcomes.

Establish provenance and gain visibility

Trust in AI outputs depends on trust in the underlying data. Organizations must be able to understand where data originates, how it is transformed, and how it is used throughout the AI life cycle.

Establishing this level of visibility requires capabilities such as data discovery, classification, cataloging, and lineage tracking. These capabilities provide the context needed to assess data quality, enforce governance policies, and ensure that data is used appropriately.

In addition, organizations should apply protection mechanisms such as tokenization and encryption before data is ingested into AI systems. This reduces the risk of exposure during downstream processing and helps ensure that sensitive data remains protected at all stages.

Implement fine-grained, identity-aware controls

AI systems operate at high speed and scale. Security cannot be the bottleneck. Access controls must evolve to become more granular and context aware. Organizations must define data ownership and accountability while implementing policies that govern both human and nonhuman identities.

This also includes distinguishing between sanctioned and unsanctioned AI agents and applying controls that determine what data can be accessed and how it can be used. In situations where full restriction is not feasible, techniques such as dynamic masking or tokenization can be used to limit exposure while maintaining usability.

These identity-aware controls enable organizations to balance accessibility with security, ensuring that data is available to authorized users and systems without introducing unnecessary risk.

Maintaining lineage with data security at runtime

Data security must extend beyond ingestion. Organizations must be able to detect anomalies, enforce policies, and respond to potential risks in real time.

Visibility is a recurring need throughout the AI process. As data moves through pipelines, is transformed, and is used to generate outputs, organizations must maintain visibility and control to ensure sensitive data is appropriately protected at all times.

This includes curating high-quality and "clean" (e.g., no confidential information or sensitive data is obfuscated) data for RAG models, as well as adequately securing data within vector databases with encryption and infrastructure controls to prevent MCP risks. In addition, organizations should apply endpoint protection or API-level controls to prevent sensitive data loss.

Observability plays a critical role in this process. Organizations need to continuously monitor and enforce the security of user behavior and prompts to prevent misuse. This helps ensure that downstream processes only interact with protected data.

By maintaining lineage, organizations can trace AI outputs back to their source data, supporting both compliance and troubleshooting efforts.

Monitoring output and adjusting policy and/or enforcement controls

The final steps in a secure AI process include logging and monitoring output as well as detailed reporting. The data preparation and observation activities outlined previously make it much easier (and faster) to answer audits and comply with industry and/or privacy regulations.

Output monitoring is critical. If data security controls fail, and sensitive data is exposed in AI responses, organizations will be able to rapidly trace back the data path to make adjustments. In addition, monitoring AI output can help organizations identify areas where outdated or incorrect information may be creating hallucinations or skewed results.

While not explicitly a security problem, incorrect output could cause data exposure issues. An example here might be determining if AI tools are surfacing documentation from insurance or banking programs that no longer exist.

Data hygiene and lineage can help organizations track backward to where problems exist and make the necessary fixes. Continuous improvement also enables organizations to refine data governance practices and improve AI performance over time. It also strengthens compliance efforts by providing clear audit trails and documentation of how data is used.

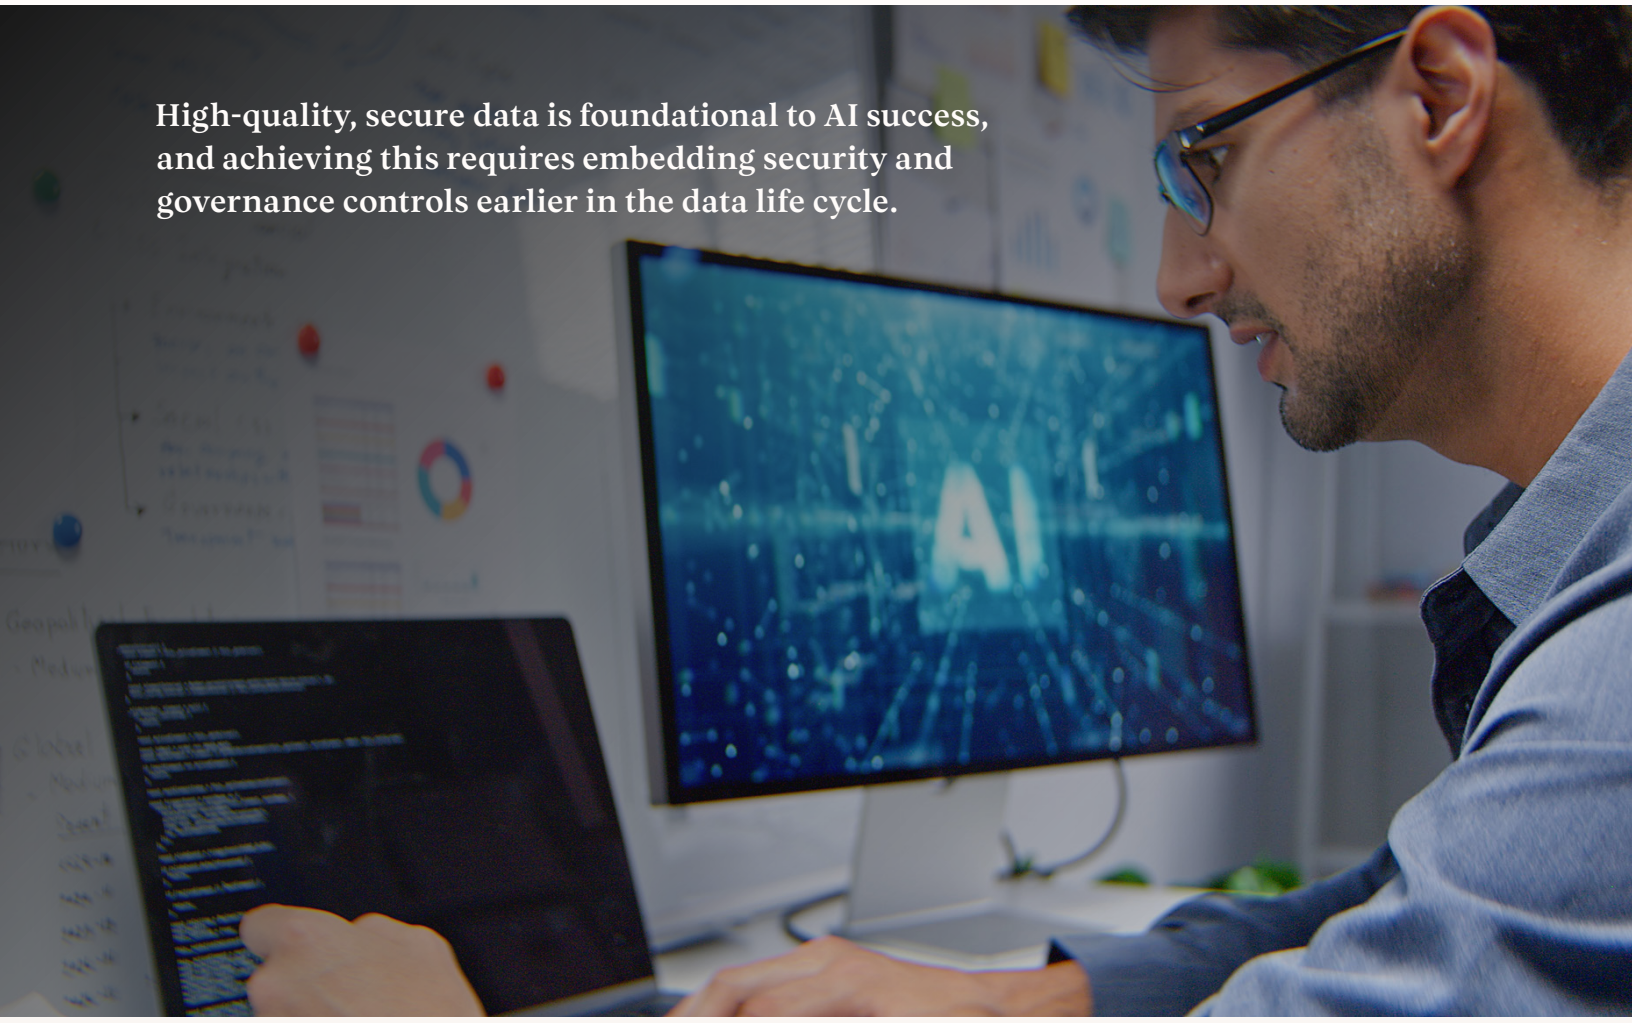

High-quality, secure data is foundational to AI success, and achieving this requires embedding security and governance controls earlier in the data life cycle.



Considering a unified data security framework

A unified data security framework integrates governance, security, and data intelligence into a cohesive architecture.

Rather than relying on disparate tools, organizations can adopt a platform-based approach that provides consistent visibility and control across all environments.

This framework typically includes end-to-end data security from source to output, including:

- **Visibility and control:**
  Centralized data discovery and classification capabilities that identify different data types and locations across the enterprise to provide broad visibility, detailed context, continuous management, and granular control

- **Real-time monitoring:**
  Continuous inspection of data activity, model interactions, and pipeline behavior as it occurs, enabling security teams to identify anomalous or malicious activity and intervene before exposure escalates
- **Automated policy enforcement and remediation:**
  Consistent enforcement of policies and automated data security remediation to ensure data is secured at all stages, in all environments, by default
- **Identity-based access controls:**
  Appropriate actions that protect data from user or device exposure using different obfuscation techniques such as encryption, masking, or tokenization
- **Support for multicloud and hybrid environments:**
  Effective security across all endpoints, devices, operating environments, and applications and wherever AI is deployed and consumed





Business impact of a unified data security framework

Organizations that adopt a unified approach to data security are better positioned to scale AI initiatives while managing risk and operational complexity.

By aligning data visibility, governance, and protection across the life cycle, this approach delivers measurable business and operational benefits:

- **Accelerated AI innovation:**
  Organizations can move from risk avoidance to risk-managed enablement. By embedding security controls before data is ingested into AI models, organizations reduce the need for downstream remediation and rework. This allows development and data science teams to operate with greater confidence, knowing that data has been properly discovered, classified, and protected. As a result, AI initiatives can progress more quickly from experimentation to production without introducing unacceptable risk.

- **Improved AI output quality:**
  AI systems are only as effective as the data they rely on. Establishing strong data hygiene, provenance, and governance practices improves the accuracy, relevance, and consistency of AI outputs. When data is properly classified, deduplicated, and aligned to business context, organizations reduce the likelihood of biased, outdated, or incorrect results. In addition, maintaining lineage throughout the AI life cycle enables organizations to validate outputs and trace issues back to the source data, supporting continuous improvement in both data quality and model performance.

- **Reduced data breach and leakage risk:**
  Protecting sensitive data before it enters AI systems reduces the risk of sensitive data being exposed later in the AI process. A unified framework enables organizations to identify and secure sensitive data across structured and unstructured sources, applying controls such as encryption, tokenization, and dynamic masking. Fine-grained, identity-aware access policies limit unnecessary data exposure by ensuring that both human users and AI agents are only able to access data appropriate to their role and usage needs. This is particularly important in AI environments where data may be dynamically retrieved, combined, or inferred, increasing the risk of unintended disclosure.

- **Strengthened compliance and sovereignty posture:**
  Regulatory requirements, such as GDPR, are constantly changing, and organizations must demonstrate greater control over how data is accessed, processed, and shared. Regulations such as the EU AI Act require transparency of AI training data as well as model evaluations to demonstrate trustworthy AI. A unified data security framework provides the visibility required to support these and other privacy and compliance initiatives. Maintaining clear data lineage and centralized policy enforcement enables organizations to track data back through its usage and sources, allowing for quicker responses to audits. In addition, broad support for hybrid and multicloud environments helps organizations meet emerging data sovereignty requirements.



Considering Thales AI Security Fabric

Thales introduced the AI Security Fabric in December 2025 to deliver runtime protection for agentic AI, LLM-powered applications, the enterprise data these systems consume, and the human and nonhuman identities that interact with them.

The platform targets AI-specific threat vectors across enterprise deployments (prompt injection, data leakage, model manipulation, and insecure RAG pipelines) with the objective of allowing organizations to scale AI without a corresponding expansion of the attack surface.

Mapped against the five framework capabilities outlined previously, the offering aligns as follows:

- **Real-time monitoring:**
  The AI Application Security component of AI Security Fabric inspects prompts and model responses in line, flagging and blocking malicious activity as it occurs. The RAG Security module extends that telemetry across the retrieval layer, tracking every interaction with structured databases and unstructured repositories throughout the RAG life cycle.

- **Centralized and automated policy management:**
  The policy authoring and enforcement capabilities are consolidated in a single control plane. This is designed to allow security teams to codify custom rules and map them to regulatory regimes such as GDPR, HIPAA, and PCI DSS. The automated remediation enforcement function reduces the manual overhead that typically accompanies multi-regulation environments and shrinks the gap between policy definition and runtime behavior.

- **Identity-based controls:**
  The platform brokers data set access for agentic AI and GenAI workloads, applying identity context to each interaction so that both human users and nonhuman agents operate within defined boundaries. Thales has committed to extending this model in 2026 with end-to-end runtime access control and a Model Context Protocol security gateway that will govern how autonomous agents connect to downstream data sources.

- **Multicloud and flexible deployment:**
  The AI Application Security component can be deployed across cloud-native, on-premises, and hybrid footprints, serving enterprises with AI pipelines that span multiple clouds and regulated on-premises environments. Uniform controls across these environments are designed to help organizations avoid policy drift that may emerge from fragmented security tools.

- **Visibility and control:**
  The RAG Security component discovers and classifies sensitive enterprise data, both structured and unstructured, ahead of ingestion into retrieval-augmented applications, then applies encryption and managed keys so that protection persists through the retrieval path. The module is also designed to secure the channel between the LLM and its external data sources, helping reduce the exposure gap.



Challenges

The Thales AI Security Fabric reflects the company's strengths in data encryption, key management, and application and identity security, capabilities that create a strong foundation for AI security. The offering announced in December 2025 represents the company's initial offering in a rapidly evolving market segment.

The Thales AI Security Fabric draws on capabilities that extend beyond the company's traditional encryption and key management legacy. The 2023 acquisition of Imperva added an established data security platform, including data discovery and classification, access monitoring, risk analytics, and intrusion detection, capabilities that map directly to the needs of securing AI data flows.

The acquisition of Imperva also added a robust application security portfolio to Thales' existing product line and is the basis for the AI Application Security capability within the AI Security Fabric. Also available now is the OneWelcome identity and access management portfolio that offers the ability to manage complex nonhuman and/or agentic identities.

For organizations evaluating the platform, it is important to understand how well Thales' core encryption capabilities have been integrated with Imperva's data visibility and application security solutions.

A portion of the platform's road map remains in development. Key capabilities, including sensitive data leakage across AI data flows, the MCP security gateway, and end-to-end runtime access control, are planned for delivery in 2026 but were not yet generally available at the time of writing this document. For organizations with active agentic AI deployments or obligations under strict regulatory frameworks, this gap may affect the platform's fit for near-term requirements and could influence deployment timing.

Finally, as outlined previously, data governance and the management of nonhuman identities are some of the most complex challenges in AI security, areas where marketwide solutions continue to mature alongside organizational readiness. Thales' progress in this space will be supported by its long-standing enterprise presence and its sustained investments in R&D but will ultimately depend on the pace at which full platform capabilities are delivered.

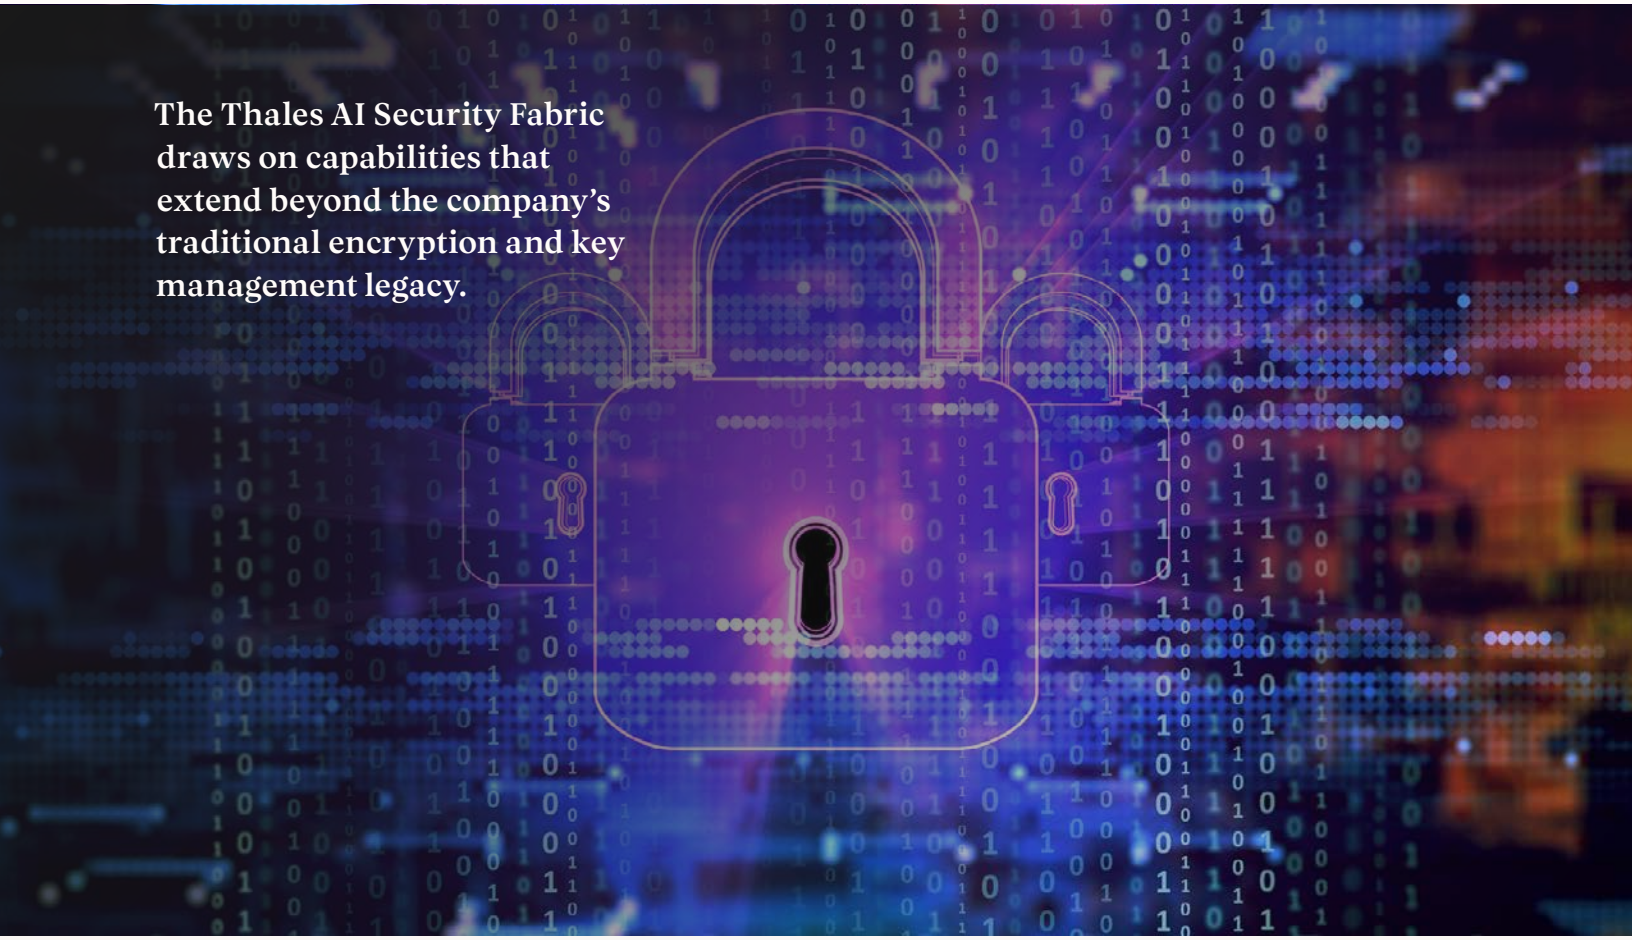

The Thales AI Security Fabric draws on capabilities that extend beyond the company's traditional encryption and key management legacy.

Conclusion

- AI has become a core enterprise initiative, but its success is increasingly determined by how effectively organizations manage and secure their data. As AI adoption expands across use cases and environments, data serves as both the foundation for innovation and a growing source of risk.
- IDC believes that organizations must move beyond fragmented and reactive security approaches to adopt a unified, proactive model. Establishing data visibility, enforcing governance, and applying protection controls before data enters AI workflows are critical to reducing exposure and enabling trusted outcomes.

IDC believes that organizations must move beyond fragmented and reactive security approaches to adopt a unified, proactive model. Establishing data visibility, enforcing governance, and applying protection controls before data enters AI workflows are critical to reducing exposure and enabling trusted outcomes.



- Organizations that implement a unified data security framework, combining data intelligence, identity-aware controls, and continuous enforcement, will be better positioned to scale AI with confidence. This approach enables improved data quality, reduced risk of sensitive data exposure, and stronger compliance and sovereignty outcomes.
- Ultimately, aligning data security with AI strategy allows organizations to accelerate innovation while maintaining control. Those that can consistently deliver trusted data to the right users and systems will be best positioned to realize the full business value of AI.

About the IDC analyst



Jennifer Glenn

Research Director,
Security and Trust, IDC

Jennifer Glenn is research director for the IDC Security and Trust Group and is responsible for the Information and Data Security. Glenn's core coverage includes a broad range of technologies, including messaging security, sensitive data management, encryption, tokenization, rights management, key management, and certificates. As part of this research, Glenn will demonstrate the critical role of data security in top enterprise initiatives, such as generating customer trust and digital transformation.

More about Jennifer Glenn →

Message from the sponsor

THALES
Building a future we can all trust

Thales is a global leader in cybersecurity, helping the most trusted organizations protect critical applications, data, and identities anywhere, at scale.

Through Thales' integrated platforms, customers achieve better visibility of risks, defend against cyber threats, close compliance gaps, and deliver trusted digital experiences for billions of consumers every day.

Thales AI Security Fabric offers enterprises specialized tools to secure AI applications while minimizing operational complexity. Supported by decades of security expertise, Thales enables businesses to accelerate their AI adoption, safeguard sensitive data, and unlock the value if AI confidently.

IDC Custom Solutions

IDC Custom Solutions produced this publication. The opinion, analysis, and research results presented herein are drawn from more detailed research and analysis that IDC independently conducted and published, unless specific vendor sponsorship is noted. IDC Custom Solutions makes IDC content available in a wide range of formats for distribution by various companies.

This IDC material is licensed for external use and in no way does the use or publication of IDC research indicate IDC's endorsement of the sponsor's or licensee's products or strategies.

IDC.com | IDC on LinkedIn | IDC Blog

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications, and consumer technology markets. With more than 1,300 analysts worldwide, IDC offers global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries. IDC's analysis and insight helps IT professionals, business executives, and the investment community to make fact-based technology decisions and to achieve their key business objectives.

©2026 IDC. Reproduction is forbidden unless authorized. All rights reserved. CCPA