SPOTLIGHT

Sponsored by: Netskope

AI success relies on secure, compliant, and well-prepared data. A unified security platform offers broad visibility and granular control, enabling organizations to confidently deploy AI without compromising business opportunities.

Improving AI Readiness with Unified, Data-Centric Security

December 2025

Written by: Jennifer Glenn, Research Director, IDC Security and Trust

Introduction

AI is reshaping the enterprise with compelling benefits: automating routine work, accelerating project completion, and processing massive amounts of data at unprecedented speed.

The rush to implement AI is pushing the boundaries of modern security. AI projects are being greenlit quickly, in some cases prioritizing speed over safety, increasing the risk of exposing sensitive data, and/or violating compliance or privacy regulations.

Business leaders need security teams to enable AI initiatives or at least keep pace with them. However, current security solutions are not designed for the AI era. Data is sprawled all over the organization in different applications, on endpoints, and in multiple clouds. Even before AI adoption ramped up, many security teams were working on identifying, mapping, and protecting sensitive data to comply with industry and privacy regulations. The increased volume of data created by digital business — combined with the complexity of AI — adds new challenges for security teams to keep organizational data protected. Yet security teams know they can't be the impediment to these projects.

AT A GLANCE

WHAT'S IMPORTANT

- AI-ready data requires a balance of protection and availability, leading to misalignment of priorities between security teams and the business.
- Securing data is at the core of securing AI, but as organizations use data within AI systems, it introduces entirely new ways that data is accessed, shared, and manipulated — requiring a broader, multilayered security approach.
- Effective data security requires enforcing controls on data entering public applications, validating the security posture of AI SaaS environments, tracking and managing data used for training, and securing the movement and interactions of AI agents.
- Visibility and context are critical elements for creating AI-ready data.

This misalignment is evident in IDC's April 2025 *Data Security and Privacy Survey,* which finds that alignment between security goals and company business objectives is shrinking when it comes to AI. In April 2024, the same survey found that 57% of respondents felt their security goals "somewhat aligned" with the business objectives for AI, and 36% felt that security and the business "completely aligned." In 2025, those numbers dropped: 45% said "somewhat aligned" and 29% said "completely aligned" (see Figure 1).

FIGURE 1: ***Security Teams Not Seeing Eye to Eye with the Business When It Comes to AI***

Data from 2024 and 2025 shows a decrease in alignment on AI between security teams and business goals.

Q *How aligned are the company's generative AI business objectives and goals with those of security?*

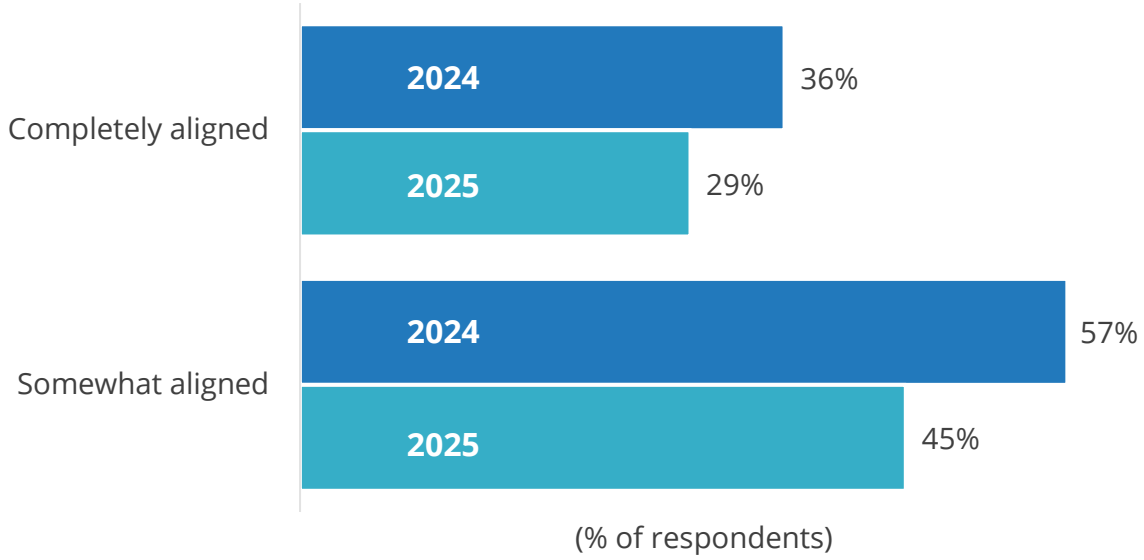

n = 619 in the 2024 survey; n = 618 in the 2025 survey

Note: This figure includes survey results from 2024 and 2025. The question was the same across both years, and the results show a marked decrease in alignment from one year to the next.

Source: IDC's Data Security and Privacy Survey, March 2024 and April 2025

This misalignment can have significant consequences for AI projects, as well as security teams, including:

- **Project derailment:** With business teams moving so quickly, security teams may not be consulted until deep into the project. If this happens and security uncovers risks to data or assets, the project may be delayed — or worse, abandoned. This waste of time and resources is the antithesis of what AI promises.
- **Exposure of sensitive data:** One of the most critical success factors for any AI initiative is data — source data, processed data, and data generated as an output. The transition to digital business has already increased the amount of data being generated within an organization, making proper data hygiene a continuous mission. AI is increasing this volume even more. If AI project teams do not consult with security until the project is complete, these teams run the risk of exposing intellectual property, personally identifiable information (PII), and/or other regulated and sensitive information to unauthorized users or third-party organizations.
- **Violation of privacy regulations:** Regulatory bodies are working hard to create new standards for AI use, and existing data protection legislation has new relevance as AI demands to be fed. Existing and emerging requirements need to be considered in advance to ensure AI projects are not derailed by compliance requirements.

Data security is built on four principles: What data does the organization have, and is it considered sensitive or confidential? Where does that data live? Who or what devices have access to that data, and what can they do with it? Finally, is that data adequately protected from threats? When it comes to AI readiness, organizations must be able to answer these questions to ensure data is being used appropriately.

Ensuring data is ready for AI applications and initiatives requires a new way of thinking that focuses on data readiness (i.e., uncovering sensitive data, establishing data provenance, using a detailed lineage to see where data has moved, and having effective data access and usage controls). The fine-tuning of all these details and controls is equally important to avoid blocking data that will impact business or AI operations.

Creating a Foundation of Secure Data for AI Readiness

Organizations are clearly aware of the value of business data. Data security, privacy, and compliance are not new concepts, and organizations already work to find, organize, and protect sensitive data. In that regard, AI implementations are similar: Data has to be protected in accordance with privacy and compliance regulations, and confidential business data needs to be available only to the appropriate users.

Where AI security differs from these traditional measures is in the overall volume of data to manage. This includes AI source data, data within the model, the output, and any other data or metadata created as part of the AI process. Further, organizations need to prove the lineage of data used for AI tools to ensure the integrity of that information.

There are a number of steps to prepare business data for use in AI applications, processes, and other implementations. These steps include:

- **Discovery:** To effectively implement the right protections, organizations should first prioritize finding where sensitive data lives. Data exists all over the organization — within applications, in shared files, on endpoints, and elsewhere. Uncovering where data exists is the essential starting point to ensure its security. In addition, knowing the location of important business data offers more flexibility in how organizations can use it.

- **Classification:** Once data residence is established, organizations need to understand what data types they have. This starts with the classification of the data. Is it sensitive or restricted (e.g., in the case of product designs)? Is it regulated data, such as PII or HIPAA? Is it internal data that is neither regulated nor restricted but still needs monitoring and controlling, such as product road maps? In any of these cases, classifying data appropriately is essential for implementing effective enforcement, ensuring only the right data is available to the right users.

  Data classification has traditionally been one of the biggest challenges for security teams. Identifying and classifying regulated data is mostly clear-cut. However, classification becomes more of a struggle when it comes to unregulated information, including internal confidential business or company data. Part of the issue is simply the high volume of data, plus all of the duplicates, revisions, and shared data that live across the business. Another complication is that what one department may consider to be internal data, another department may consider to be sensitive. Classification requires context to be useful.

  AI not only exacerbates the aforementioned legacy security issues but it also opens the door to even more risk. For example, if the data classification is incorrect, organizations risk inadvertently exposing high-value, confidential company data. In addition, improperly classified data may result in the use of low-value internal data for AI model training, affecting the quality of output and the success of the project.

- **Context:** Understanding data context is an important element for building a foundation of secure data. Context offers insight that distinguishes similarly formatted data. One example is correctly identifying a 16-digit credit card number instead of a customer ID that may also have 16 digits. This context comes from knowing who has access to the data and/or what machines or services can use it. Armed with the right classification and contextual information, implementing the right guardrails for which departments, roles, or applications can access or use the data becomes much clearer.
- **Control:** Finally, it's important to know how the organization is protecting its sensitive corporate information and whether that approach is enough to meet security and/or compliance requirements. This requires full visibility of data and all of its attributes. Point solutions offer insight and visibility at the data level, but when it comes to AI, this is not adequate. Full AI data readiness requires a much broader view of data within the context of how it moves through the network, its use in web-based applications, and how the organization is storing and using it on each endpoint.

Benefits of a Unified, Data-Centric AI Security Platform

Currently, organizations are piecing together security technologies to achieve this level of data visibility and control. While this may offer some insight, it takes time to uncover the relevant information for each product and piece it together into something useful. The time required to do this can mean delays in protection and the accrual of extra costs. A unified solution that can pull context and information from each of these control points — endpoints, web, applications, emails, and networks — offers comprehensive security at the speed of AI.

Consolidating multiple security technologies to secure AI data offers a number of business benefits, including:

- **Reduced costs:** Procuring multiple individual security tools can be expensive in terms of time and money. Licenses may expire at different times, and working through sales and support issues for each product can be a tedious task. Further, each solution may require different expertise, and additional staffing can incur extra costs. Finally, pulling information from each of these solutions also requires significant investment to build effective connections to achieve the desired level of visibility for an AI-driven business.
- **Simplicity:** A unified platform for AI security also becomes much easier to manage and maintain. Security data coming from multiple control points (endpoint, data, identity, web, and cloud) arrives in similar formats, making it easier to share and analyze. Further, any control/enforcement policies that a business implements are more consistent across each modality.
- **Stronger business posture:** Security incidents and privacy/compliance violations can damage an organization's brand. Not only do these situations weaken the organization at the moment of the attack and its immediate aftermath but they can also set back key innovations and disrupt support operations, which can impact a consumer's view of an organization.

Considering Netskope

Netskope One is a security and networking platform designed to offer secure access and real-time, context-based data security. Combining multiple data and network security point solutions offers organizations a modern approach to

reducing security risk while consolidating cost and support functions. Each technology offers new insights and context that businesses can use to fortify the other solutions.

Netskope's platform is designed for unified data security to protect sensitive information sitting on endpoints, within cloud applications and collaboration tools, and during its journey throughout the organization. The platform uses DLP and DSPM capabilities to understand where data exists, how well it's protected, and who is exfiltrating company data. The incident evidence coming from these capabilities is combined with risk information from other tools in the portfolio to offer broader visibility and context around the data. This information can be used for building, adjusting, or enforcing policies across the entire business.

Netskope offers these capabilities across multicloud environments to help organizations address numerous use cases for security and AI, including, but not limited to:

- **Understanding of AI application risk:** The Netskope One platform is designed to offer comprehensive visibility and understanding of application risk from all angles. It aims to find what applications are in use and what type of data supports them. The platform looks at what the application is doing with the data, such as sharing it with a third party. This includes gaining insight into the terms and conditions of the third-party provider to determine whether it is using the data responsibly.
- **Securing access to AI tools:** Netskope is designed to ensure that the right people have access to the right AI tools and applications, including third-party contractors and external collaborators. By enforcing granular, role-based access controls, Netskope aims to ensure users can only reach the apps and data necessary for their specific tasks. This reduces the risk of unauthorized access or data exposure while maintaining productivity. Combined with continuous monitoring and contextual security policies, organizations can safely extend AI capabilities to external partners without compromising sensitive information.
- **Securing data in custom-built AI applications:** Netskope DLP on Demand was created to help organizations build their own AI applications while ensuring that attacks, such as prompt injection or even jailbreaking, do not compromise the AI engine. The organization is also responsible for ensuring that output from this custom application does not send back incorrect or inappropriate responses, such as biased, racist, or malicious information.
- **Protecting the AI data pipeline:** Organizations know that to get optimal AI output, it's important to understand the data and context to ensure that only high-quality and non-sensitive information is fed into the model. Netskope's platform includes capabilities for identifying, labeling, and understanding the risk of exposure to the data before it goes into the model.
- **Protecting the organization from emerging AI threats:** The AI threat landscape is evolving just as quickly as most businesses, and for many of the same reasons: automation and efficiency. As new threats emerge, Netskope can monitor these attacks and their vectors to strengthen data protection throughout the AI life cycle.
- **Maintaining a secure and compliant AI posture:** As organizations mature in their AI journey, many want an overall view of their entire AI ecosystem — including SaaS applications that rely on AI for private apps and agentic interactions — to view all the risks and vulnerabilities. Netskope aims to offer insight and understanding of these interactions to prove compliance with industry and/or privacy requirements.

Challenges

While the market is moving toward consolidation, it may be several years before organizations fully rely on a single platform for data and/or AI security. First, many organizations have invested heavily in point solutions and may not be anxious to scrap these tools. They also have long-term licenses that they cannot change or cancel without further control. Finally, some organizations may be reluctant to trust a single vendor for all of their security needs.

Another threat to a unified AI security platform from a traditional security vendor comes from hyperscalers. Large cloud providers offer multiple security services natively with their deployments. While organizations still rely on technology vendors for many solutions, the simplicity (and cost and labor savings) of native security tools can be incredibly attractive, particularly to the midmarket.

Netskope has a solid customer base that offers an opportunity for expansion; however, the company will need to clearly articulate how a module in the platform has the same functionality as the point solution(s) it is replacing. This also means Netskope will need to show the definitive value of the platform to justify replacing individual tools.

Conclusion

AI readiness requires comprehensive visibility of data across every modality in the organization to provide the granular control that secures without disrupting business.

AI is rapidly plowing through enterprise budgets and plans. Deploying AI requires alignment across all disciplines within the business, and security cannot be the exception. Security practitioners cannot be the speed bump in getting AI projects to production. Security requires a better way to ensure only the right data is available for these projects, meaning carefully curated information that has been checked for regulated, sensitive, or confidential content, as well as redundant, obsolete, or trivial data.

It will not always be easy. Data is dynamic; organizational data is always in motion, being shared, stored, and used across all parts of the business. AI amplifies both the volume and use of data, as well as spreading it to different applications. Further, AI continually exposes gaps in data security. While point solutions can help address such gaps, the resulting web of information from these tools is often fragmented and useful only in specific situations.

A platform to secure AI-ready data as it interacts with AI systems and processes offers a quicker and more organized approach to keep pace with business projects while still conforming to data security and compliance requirements.

AI-ready data requires broad visibility into how organizational data is used and managed. With visibility as the foundation, organizations can easily uncover sensitive data wherever it lives and ensure appropriate classification. Further, combining information and telemetry from the organization's endpoints, clouds, and networks offers significant context around how data is being used — and insight as to whether it is appropriately protected. Armed with this information, organizations can also unify management and control of data, offering consistent policy enforcement and consolidated reporting.

A unified approach to data security offers enterprises simplicity and effectiveness. Data security practitioners have the tools and support to keep pace with AI changes and ensure corporate data is ready for the AI business era.

About the Analyst

Jennifer Glenn, *Research Director, IDC Security and Trust*

Jennifer Glenn is research director for the IDC Security and Trust Group and is responsible for the Information and Data Security practice. Ms. Glenn's core coverage includes a broad range of technologies including messaging security, sensitive data management, encryption, tokenization, rights management, key management, and certificates.

MESSAGE FROM THE SPONSOR

AI readiness requires security that can protect sensitive information, ensure compliance, and preserve data integrity without slowing down business adoption. A data-centered approach to security helps organizations establish the guardrails needed to reduce risk while enabling responsible AI use across applications and workflows.

To learn more about how to strengthen data security as part of your AI readiness strategy, read our AI Security Playbook.

IDC Custom Solutions

The content in this paper was adapted from existing IDC research published on www.idc.com.

IDC Research, Inc.
140 Kendrick Street
Building B
Needham, MA 02494, USA
T 508.872.8200
F 508.935.4015
blogs.idc.com
www.idc.com