Omdia
by informa techtarget

# Minimizing Data Risk in the AI Era

## Solving Compliance and Security Challenges

Author: Todd Thiemann

November 2025

This White Paper from Omdia was commissioned by Zscaler and is distributed under license from TechTarget, Inc.

# Contents

# Executive Summary

Enterprises are embracing AI to streamline operations and gain a competitive advantage. However, AI expands the threat surface in new ways and raises new data security risks. AI initiatives are inhibited by inadequate data security that can put sensitive data at risk. The multifaceted nature of securing AI means that point solutions can leave gaps in visibility and protection. Zscaler provides a holistic approach that helps facilitate enterprise AI deployments that protect data and maintain regulatory compliance.

# Sensitive Data Proliferates, and GenAI Makes It More Accessible

Enterprises are embracing the adoption of generative AI (GenAI) and large language models (LLMs). Across today's enterprise, different teams are exploring how to leverage it to increase productivity, improve processes and workflows, speed decision-making, and enhance the customer experience using GenAI.[1]

While GenAI is improving productivity and unleashing innovation, it also poses security challenges if sensitive data is inappropriately used to inform LLMs. If valuable or sensitive data

[1] Source: Enterprise Strategy Group Complete Survey Results, *AI Agents: The Game-changing Generative AI Use Case*, August 2025.

is inadvertently included in model input and subsequent output, organizations face the risk of critical data landing in the wrong hands (data leakage) or violating governance or compliance obligations by inappropriately using or sharing sensitive data.

Sensitive data, such as personally identifiable information, financial records, and intellectual property, is vital to business operations. Such data was previously locked away through data access governance solutions or inaccessible by obscurity, as insiders or adversaries could not quickly locate sensitive data. AI copilots change that dynamic by putting sensitive enterprise data at the user's fingertips when they input the right AI prompt, frequently irrespective of whether or not that user should have access.

Securing sensitive data sustains operations, maintains competitive advantage, and establishes trust with customers and suppliers. Enterprises need to secure their data to manage risk and avoid financial loss, regulatory penalties, and operational disruption. That data is scattered across a broad estate that includes on premises, SaaS apps, and the public cloud. Of particular interest is the speed with which AI data has grown as a destination for sensitive data. While 56% of respondents pointed to sensitive data residing in the relatively mature area of cloud storage, 42% of respondents reported having their sensitive data residing in the emerging area of GenAI applications (see Figure 1).[2]

Figure 1: Where Organizations Store Their Sensitive Data

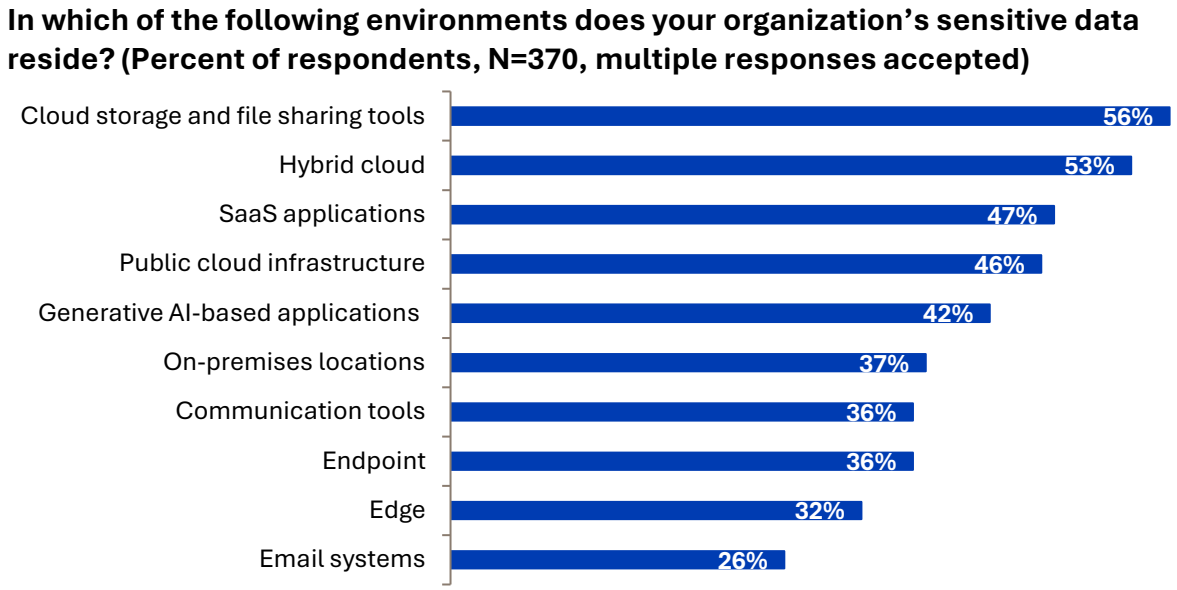


Source: Omdia

[2] Source: Enterprise Strategy Group Research Report, *Reinventing Data Loss Prevention: Adapting Data Security to the Generative AI Era*, May 2025.

Securing AI infrastructure requires a multifaceted approach to address the complex nature of AI deployments. AI deployments come in multiple forms, including public AI (ChatGPT, Gemini, Claude), SaaS copilots (Microsoft Copilot, Github Copilot), and private AI that is developed internally by enterprises.

The multifaceted nature of GenAI deployments with various elements, including training and retrieval-augmented generation (RAG) data, LLMs, copilots, and AI services, requires that enterprises think comprehensively about how they secure AI infrastructure.

It is important to have a strategic approach to data security that optimizes efficiency and minimizes gaps between solutions, helps organizations avoid a heavy integration burden, enables security teams to avoid swiveling between consoles when investigating and resolving incidents, and streamlines compliance efforts.

Gaps between solutions can require more internal resources to integrate them or manually analyze data across solutions and can increase the risk of sensitive data slipping out between data security point solutions and the amount of time needed to demonstrate regulatory compliance.

**Definitions: Public AI and Private AI**

- **Public AI** refers to general-purpose artificial intelligence models such as ChatGPT and Gemini that are trained on publicly available or large-scale internet data.
- **Private AI** is built and deployed within an enterprise's secure environment and trained on proprietary or internal data. It is designed to protect an organization's sensitive information while enabling customized AI capabilities.

## Sensitive Data Informs and Permeates AI Infrastructure

Data informs GenAI applications, and that data can include sensitive data, intentionally or inadvertently. Once sensitive data is included in GenAI infrastructure, it becomes very difficult to control or delete.

Data used to inform GenAI infrastructure can come from a variety of sources, including RAG architectures, native connectors used by Microsoft Copilot, and traditional ETL and API integrations to extract, transform, and deliver structured and unstructured data into AI-ready environments.

As enterprise teams rush to adopt AI for its benefits, teams have lacked visibility into much of their data estate, with 60% of enterprises lacking visibility into up to half of their data estate

(see Figure 2).[3] Before the advent of GenAI, this gap did not pose a significant issue because undiscovered sensitive data might not have been readily accessible to employees who did not know how to access it (i.e., security through obscurity).

Figure 2: Many Organizations Lack Visibility Into Much of Their Data Estate

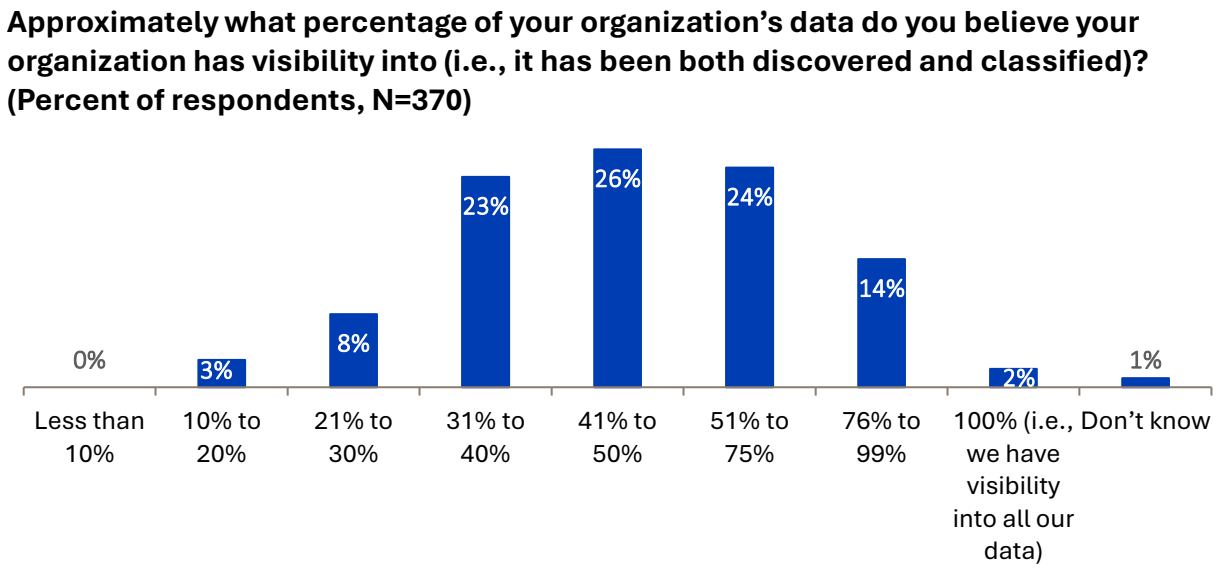


Source: Omdia

GenAI has changed that dynamic. Previously undiscovered or uncategorized data is now readily accessible when used as training data for GenAI infrastructure and can appear in response to prompts from enterprise copilots or other GenAI apps. While those apps can increase enterprise efficiency and help the bottom line, the newly accessible sensitive data heightens data security risk, and emerging AI compliance regulations raise new compliance issues.

In particular, Copilot presents an oversharing challenge. Microsoft Copilot can provide substantial boosts to productivity, while also overconsuming and sharing sensitive data with less-privileged users. Microsoft provides Copilot Security options, but security challenges remain such as improperly permissioned data being consumed and shared too widely on OneDrive. If users or systems fail to label sensitive data properly, it's open to access by Copilot prompts, which creates security risks. In addition, gaps in Microsoft 365 and Copilot configurations can expose data to risks.

[3] Ibid.

The risk lies in sensitive data that can make its way into the wrong hands, be it external adversaries or inappropriate insiders. This is not a hypothetical risk: research from Enterprise Strategy Group (now Omdia) shows that 43% of enterprises indicated that they have experienced a sensitive data loss in the last 12 months through a generative AI application.[4]

43% of enterprises indicated that they have experienced a sensitive data loss in the last 12 months through a generative AI application.

# Securing and Controlling Public GenAI

Visibility to the enterprise AI landscape is a prerequisite to controlling and securing sensitive data. While much of the AI landscape may be known to the security team, there is typically a healthy slice of "shadow AI" being used across an organization. AI app adoption has proliferated across the workforce, and teams are frequently unaware of which AI apps are sanctioned or unsanctioned.

Employees or contractors may have taken the initiative to deploy a cool AI tool, but that tool may be duplicative to a sanctioned tool, or it might represent a security risk. Permissions around the tooling may also be unknown; teams can be unfamiliar with which departments or users have access to which data and which AI app. Should an engineering team have access to public ChatGPT since an engineer prompt might contain sensitive source code that can become training data for a public AI model?

## Discovering and Classifying Data Informing GenAI Infrastructure

Sensitive data resides across the enterprise data estate, and enterprise security teams need to understand where their sensitive data lives in order to secure it and maintain compliance with external regulatory regimes and internal data governance requirements.

While information technology (IT) and information security teams have visibility into most data stores, they frequently lack visibility into unknown, hidden, or overlooked copies of sensitive information that exist outside the purview of an organization's IT security measures and data governance policies. Such data is often created, stored, or shared without being formally managed or governed by the relevant IT and security teams.

Shadow data can reside across the enterprise in unstructured files, structured databases, cloud storage, and personal devices, often without the IT department's knowledge or control. For example, it can include data copies in test environments, unmanaged backups, abandoned databases, data extracted by insiders, and data leakage via third-party apps.

[4] Ibid.

All of this data can end up being ingested into AI infrastructure of one sort or another. IT and information security teams need to understand that data and categorize it appropriately so the expected data informs GenAI infrastructure. Not doing so risks sensitive information making its way into GenAI infrastructure and subsequently leaking out of the enterprise.

### Ensuring Copilot Security Readiness

To mitigate data risks relating to Microsoft Copilot, organizations should audit and update all OneDrive data permissions and missing Purview Sensitivity Labels in order to prevent oversharing sensitive data. Additionally, security teams require visibility and context to all user prompts to Copilot and enforce inline blocking of sensitive data in Copilot. This will allow teams to monitor and control Copilot usage interactions. To prevent risky misconfigurations, continuously scanning Copilot and Microsoft 365 for dangerous posture that may expose data will further reduce risk.

### Enforcing Security Guardrails for GenAI

The saying that "you can't manage what you can't see" applies particularly well to GenAI applications. To prevent data loss from GenAI applications, it's important to understand AI usage so you can enforce policies that monitor and restrict how employees interact with public GenAI tools. A key element is implementing control mechanisms that can block or control access to applications like ChatGPT when sensitive information is at risk, while still allowing safe productivity gains. Even simple controls such as sanctioning or unsanctioning apps can go a long way in preventing data loss. But more importantly, having granular control into user access, prompt usage, and sensitive data ingested into apps can significantly improve an organization's AI security posture.

## Securing and Controlling Private AI

While GenAI opens a variety of potential data loss vectors, the major risk for significant data loss in the realm of GenAI comes from sensitive data being processed in private AI apps and infrastructure.  GenAI represents a tremendous enterprise efficiency booster, but it also can make sensitive data inside the enterprise more readily available to insiders and potential adversaries.

### Understanding Data Informing AI Tools

A key pillar for AI security lies in data security posture management (DSPM), which can improve the governance and security of data for GenAI usage by identifying and categorizing enterprise data so that sensitive data does not inadvertently make its way into an LLM model or AI tools.

Once the AI data is used in training, it is generally not possible to fully scrub sensitive data from a trained LLM. DSPM is a critical tool to ensure that the right data goes into the AI

infrastructure. LLMs do not store raw data records in something resembling a database but rather learn statistical patterns and relationships in the training data. Once that data has influenced the model weights, there is no standard, effective, provable way to “erase” it without retraining a model—something that is very costly and complex.

DSPM tools can evaluate on-premises and cloud data so that data security and data governance teams understand the data informing their AI infrastructure.

## Inventorying AI Assets – AI Security Posture Management (AI-SPM)

A comprehensive AI asset inventory provides organizations with visibility into all AI applications, models, agents, and services in use. This includes third-party and shadow AI (i.e., unofficial or unauthorized technology such as open source models). Visibility enables informed risk assessment to ensure that sensitive data is not inadvertently exposed to unregulated, unsecured, or deprecated systems.

Inventory needs to include information about the AI model infrastructure along with context, including the publisher, country of origin, licensing terms, and key risk factors.

Services inventory needs to be comprehensive and include major cloud provider AI services, such as Microsoft Azure AI Foundry, Amazon Bedrock, and Google Vertex AI.

Asset inventories are integral elements of AI frameworks and help to guide safe adoption of AI technologies. Knowing what you have reduces blind spots where sensitive information or vulnerabilities can linger without detection.

Understanding your organization’s AI security posture is also valuable for meeting regulatory frameworks such as the EU AI Act, NIST AI RMF, and ISO 42001, which increasingly require detailed AI asset inventories. As auditors come calling, security teams will need to have such information handy to answer questions and satisfy regulatory scrutiny.

## The AI Security Posture Management Process

Beyond understanding their AI assets, enterprises need to not only discover but also assess and monitor the AI security posture on an ongoing basis. AI security posture management (AI-SPM) addresses this by assessing models, data pipelines, and deployment environments to protect against vulnerabilities, ensure compliance with regulations, and minimize risks specific to AI technologies. AI-SPM works hand-in-glove with DSPM to optimize security posture. While DSPM discovers and categorizes sensitive data that informs the AI infrastructure, AI-SPM focuses on ensuring the security of AI infrastructure. AI-SPM secures the underlying AI infrastructure and looks for issues such as misconfigurations and vulnerabilities. Organizations need to feed the right data into AI infrastructure and then ensure the appropriate security posture of the AI infrastructure. Both are needed for a robust security program.

## AI Red Teaming

In the same way that red teaming helps in other areas of security, AI red teaming can proactively expose vulnerabilities in AI systems before attackers or misuse can exploit them. Unlike traditional penetration testing, which focuses on networks and applications, AI red teaming simulates adversarial scenarios unique to machine learning models—such as prompt injection, model inversion, and adversarial inputs—to uncover weaknesses in how AI interprets, generates, or secures data. By conducting structured adversarial testing, organizations strengthen resilience, ensure regulatory alignment, and build trust in AI-driven operations. AI red teaming acts as a safeguard that validates whether enterprise AI systems can withstand real-world threats, making it a significant component of responsible AI governance.

## AI Runtime Security

AI runtime security defends generative AI systems against adversarial inputs such as malicious or poorly structured prompts that can manipulate model behavior, bypass safeguards, and expose sensitive data.

Once internal AI apps are deployed, AI runtime security can secure apps in use and defend against prompt injection, jailbreaks, and malicious tool use while also enabling prompt and output monitoring. Security teams also need the ability to review prompts to detect model misuse and mitigate data exposure risks.

Security and compliance teams need to consider existing regulations like the EU AI act, US state-level laws, and emerging frameworks and standards such as ISO 42001 and NIST AI RMF that provide risk-based approaches to regulatory alignment.

## The Need for Holistic Solutions

While enterprises need to set controls to minimize risk for a wide range of AI-related threats, having a holistic solution approach covering data security avoids the potential for gaps between solutions. A holistic approach provides the shortest path to improved data security and compliance.

# AI and the Changing Compliance Landscape

Enterprises need to mitigate data security risk, but they also need to maintain data privacy and ensure regulatory compliance.

Generative AI and LLMs are subject to existing and new compliance regulations. For example, public companies in the US are subject to SEC AI governance guidance published in 2025, and companies operating in Europe are subject to the EU AI Act, which regulates high-risk AI systems by requiring mandatory conformity assessments, data transparency, and human oversight.

Compliance obligations are consequential. Not fulfilling regulatory obligations around privacy and data governance can result in fines, legal risk, increased regulatory scrutiny, and reputational and operational consequences. Noncompliance with disclosure requirements, particularly when it comes to reporting material AI risks, can trigger formal investigations and shareholder lawsuits.

Today's annual compliance processes can consume scarce internal resources. They tend to be very manual and tedious exercises for most enterprise teams. Many organizations are looking for automation of mapping to compliance frameworks to streamline or eliminate manual compliance efforts.

Whether experimenting with LLMs or deploying GenAI applications such as Microsoft Copilot or ChatGPT, organizations are relying on enterprise data. Research from Enterprise Strategy Group (now Omdia) shows that, when it comes to deploying AI agents, enterprise IT and security teams ranked data privacy issues and compliance and regulatory risks as two of the top three areas of concern (see Figure 3).[5]

[5] Source: Enterprise Strategy Group Research Report, *Identity Security at a Crossroads: Balancing Stability, Agility, and Security*, October 2025.

Figure 3: Security and Business Risks of Using AI Agents

**Which of the following security and business risks is your organization concerned about in relation to using AI agents? (Percent of respondents, N=370, multiple responses accepted)**

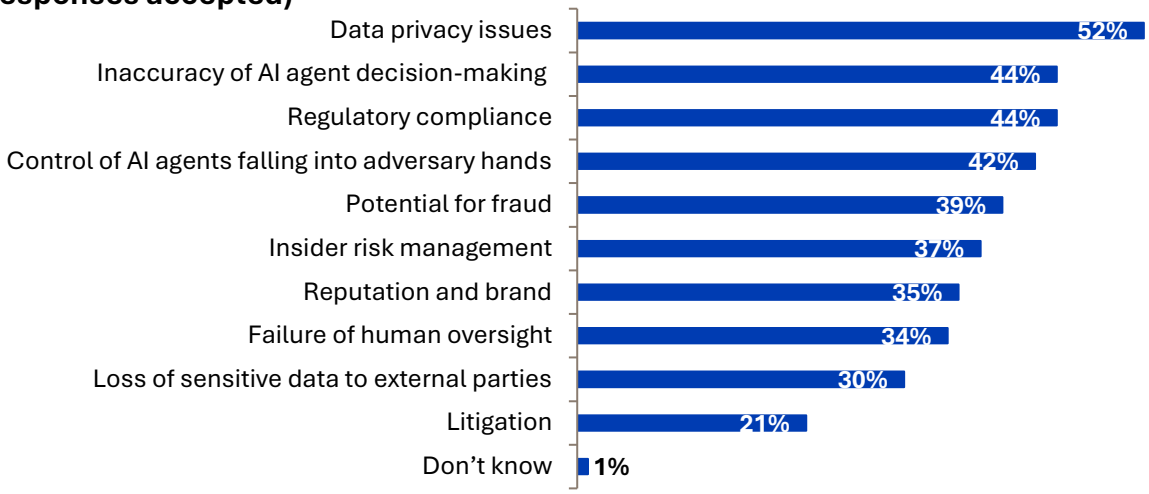


Source: Omdia

GenAI data is also subject to a growing number of regulatory regimes that focus on privacy, security, transparency, data minimization, and responsible data governance.

The regulatory landscape for AI is in considerable flux with increasing fragmentation among US states, harmonization efforts in the European Union, and increasing vertical sector oversight around the globe, making it difficult for security teams to keep up to address needed regulations as they rapidly evolve.

Different jurisdictions can have different rules, which complicates compliance efforts. While AI compliance challenges can delay innovation within an organization, non-compliance can risk fines and reputational damage to the organization.

The AI security guidance and regulations shown in Table 1 have a consistent focus on controls to ensure the integrity, confidentiality, and accountability in data used by and generated from AI.

Table 1: Prominent AI Security Guidance and Compliance Regulations

| AI Security-related Guideline or Regulation | Description of AI Security-related Guideline or Regulation |
|---|---|
| **EU AI Act (EU)** | Provides a risk-based regulatory framework designed to oversee the development, deployment, and use of AI technologies across the EU. It prohibits harmful uses of AI, classifies "high-risk" systems with documentation and oversight obligations, and mandates transparency, human oversight, and monitoring for general purpose and sectoral AI deployments.<br>The EU AI Act applies to any organizations with AI systems in the EU or whose AI systems produce any products/services used in the EU. |
| **Digital Personal Data Protection Act 2023 (India)** | Established obligations for entities handling personal data (Data Fiduciaries) around the lawful, fair, and transparent collection, processing, storage, and transfer of personal data. |
| **NIST AI Risk Management Framework (RMF – US)** | Provides a voluntary risk and governance structure to enhance trustworthiness, security, and accountability of AI systems across their lifecycle. |
| **Council of Europe Framework Convention on AI (Global)** | A global AI treaty requiring signatories to align AI use with human rights protection, democratic oversight, non-discrimination, and transparency to ensure that individuals have effective remedies with AI systems that cause harm. This is intended to complement the EU AI Act. |
| **California Consumer Privacy Act (CCPA) and US State Privacy Laws (US)** | Provides individual rights over personal data collection, sale, and automated decision-making. Updates to CCPA approved in 2025 extend to AI impact assessments, cybersecurity audits, and automated decision transparency. |
| **METI AI Governance Guidelines for Business (Japan)** | A framework for ethical and responsible AI adoption in Japan put forward by Japan's Ministry of Economy, Trade and Industry (METI), and the Ministry of Internal Affairs and Communications (MIC). The guidelines are non-binding but set expectations for organizations across all sectors deploying artificial intelligence. |
| **AI Ethics Principles (Australia)** | Voluntary, government-endorsed guidelines consisting of eight principles intended to promote safe, secure, fair, and transparent use of AI across all sectors in Australia. |

Source: Omdia

# Zscaler Data Security Platform

Zscaler helps enterprises adopt security best practices and security controls that optimize data security while also mitigating compliance risk. The Zscaler platform enables enterprises to holistically secure GenAI infrastructure and applications to ensure data security and regulatory compliance. Organizations can create and enforce policies around GenAI tools that users can utilize as they adopt and use AI while protecting sensitive data.

Zscaler provides control over all aspects of GenAI use so enterprises can securely deploy GenAI applications. The Zscaler platform is comprehensive in securing both public GenAI, including copilots like Microsoft Copilot, as well as private AI in the near term.

Zscaler secures and supports a myriad of AI use cases, including:

- **Secure Public GenAI:** Zscaler eliminates shadow AI and provides deep visibility into an organization's entire public GenAI utilization (i.e., which AI apps are in use, who has access to apps, and what data is used in apps, etc.). Zscaler also provides more controls around which apps are sanctioned or unsanctioned, along with user access to apps and data access to apps. The solution controls against having users leak data into public models.
- **Secure Microsoft Copilot/SaaS Copilots:** Zscaler enforces safe guardrails for Copilot data access while maintaining security posture. The solution ensures optimal operation for OneDrive, Purview, and Microsoft 365, while helping control Copilot oversharing and misconfigurations. Zscaler helps clean up permission hygiene and provides bulk data labeling for Microsoft Purview. In addition, Zscaler provides more granular controls against oversharing sensitive data in repositories like OneDrive and Sharepoint.
- **Secure Enterprise-managed Private AI:** Secure AI starts with securing data and AI models. Zscaler AI-SPM provides insights into an organization's AI-powered environments while proactively mitigating data and AI risks. AI-SPM ensures that inappropriate data does not leak into the model. Zscaler AI Red Teaming continuously tests AI systems and guides remediation to discovered vulnerabilities while Zscaler AI Guard AI Runtime monitor AI flows to block malicious attacks and prevent data leakage

# Conclusion

Enterprises are embracing AI to improve productivity and efficiency and open new revenue streams. Security teams need to support AI adoption with the right controls in place to minimize security risk posed by new GenAI and agentic AI applications.

Organizations can utilize the Zscaler platform for a holistic approach to GenAI security to optimize efficiency in mitigating the risk of a data breach and ensure compliance across a dynamic regulatory environment.

# Appendix

## Methodology

This white paper is based on research from Enterprise Strategy Group (now Omdia) on data security, identity security, and generative AI. The multiple research projects underpinning this paper surveyed data security, identity security, and IT leaders with enterprise organizations in the US and Canada.

**Todd Thiemann**, Principal Analyst, Data Security and Identity Security
askananalyst@omdia.com

**Omdia consulting**

Omdia is a market-leading data, research, and consulting business focused on helping digital service providers, technology companies, and enterprise decision makers thrive in the connected digital economy. Through our global base of analysts, we offer expert analysis and strategic insight across the IT, telecoms, and media industries.

We create business advantage for our customers by providing actionable insight to support business planning, product development, and go-to-market initiatives.

Our unique combination of authoritative data, market analysis, and vertical industry expertise is designed to empower decision-making, helping our clients profit from new technologies and capitalize on evolving business models.

Omdia is part of Informa TechTarget, a B2B information services business serving the technology, media, and telecoms sector. The Informa group is listed on the London Stock Exchange.

We hope that this analysis will help you make informed and imaginative business decisions. If you have further requirements, Omdia's consulting team may be able to help your company identify future trends and opportunities.

**Get in touch**

www.omdia.com
askananalyst@omdia.com