GUIDE

proofpoint.

Securing and Governing AI in the Modern Enterprise

A guide to protecting sensitive data in the age of AI

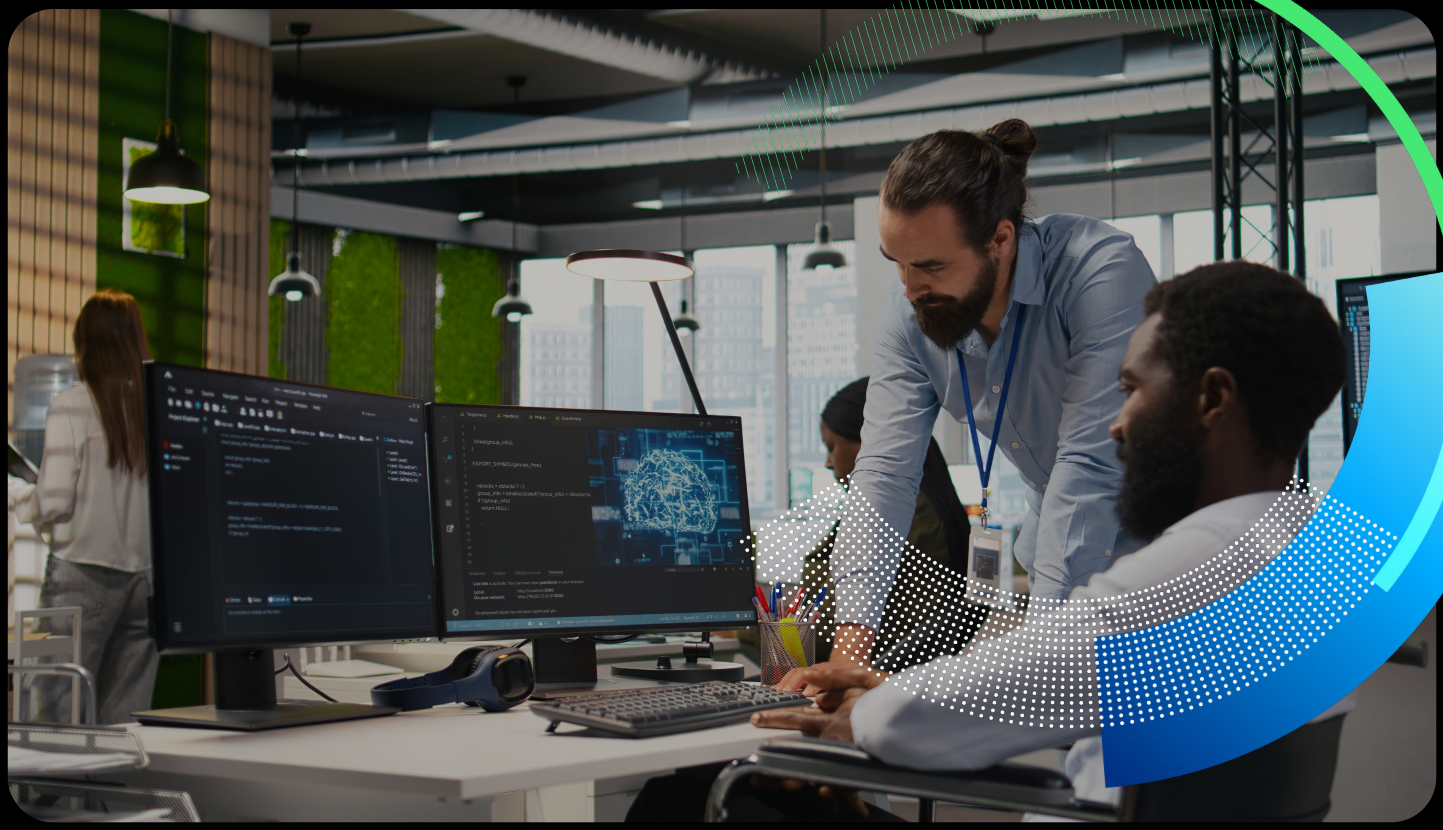

Introduction

A brave new world

January 2025. A high-profile hacker adds a post with some extraordinary claims to a notorious cybercriminal forum. The hacker boasts about breaching the defenses of OmniGPT, an AI aggregator trusted by over 30,000 users. Their alleged prize is a staggering haul of 34 million AI chat messages and 6,000 confidential file links. Those allegedly reveal API keys, crypto wallets, credentials, 30,000 email addresses and thousands of phone numbers—all put up for sale on the dark web.[1]

AI is rapidly reshaping society and industry, rewriting the rules of creativity, work and innovation. But with its transformative benefits come complex new data security challenges.

AI tools pose unique risks because of their potential to ingest and later expose sensitive information. The alleged OmniGPT attack wasn't just a breach; it was a rallying cry for hackers and a wake-up call for security practitioners. The alarming numbers in this attack show that without robust defenses, even the most advanced AI systems can become treasure troves for cybercriminals.

This guide:

- Examines the emerging data security risks that modern enterprises face from unsecure use of AI
- Explains how weak AI security governance opens the door to those risks becoming damaging breaches
- Lays out a framework for strong AI security governance
- Explores a range of tools, techniques and practices to secure AI applications and data in this new AI age

1. SecureWorld. "OmniGPT Data Breach Exposes 30,000 Users and Millions of Chat Messages." February 2025.

Chapter 1

Data risks from unsecure and ungoverned AI use

AI is transforming how organizations create, manage and use data. Powerful additions to enterprise IT environments include generative AI (GenAI) tools, such as ChatGPT; enterprise AI tools, such as Microsoft Copilot; and AI agents. While these types of tools promise huge gains in productivity and operational efficiency, they've also introduced a range of new data security risks. Compared with traditional software applications, AI tools use sensitive data in more opaque ways. They create new threat vectors that legacy security models simply weren't designed to handle.

This chapter explores some of the data security risks—and damaging outcomes—that have already emerged from unsecure and ungoverned AI use, even as the technology continues to evolve.

Transformative gains, serious risks

Modern enterprises are rushing to deploy AI to transform their business workflows. But rapid adoption has also paved the way for accidental leakage of sensitive data. This might be valuable or confidential information such as intellectual property (IP), personally identifiable information (PII), payment card information (PCI) or protected health information (PHI). Its loss can be particularly costly in highly regulated industries such as healthcare and financial services.

Accidental data loss through AI can happen in several ways:

- **User prompts**: Employees might input proprietary data or upload confidential files to public GenAI tools, risking later exposure.
- **Model training or tuning**: Enterprise data is often used to train or fine-tune custom large language models (LLMs) and agentic (that is, agent-based) AI applications. It's also used in retrieval-augmented generation (RAG) systems. But there's a risk that models later expose such information, especially if queried with clever prompts.
- **Overprivileged AI applications**: Overprivileged custom, enterprise or agentic AI applications might have broad access to corporate data. This creates the risk of accidental disclosure. For example, an AI agent with read access to internal documentation and a misconfigured prompt guardrail might reveal confidential product information. Without tight controls, AI can become a superuser that's difficult to audit or constrain.

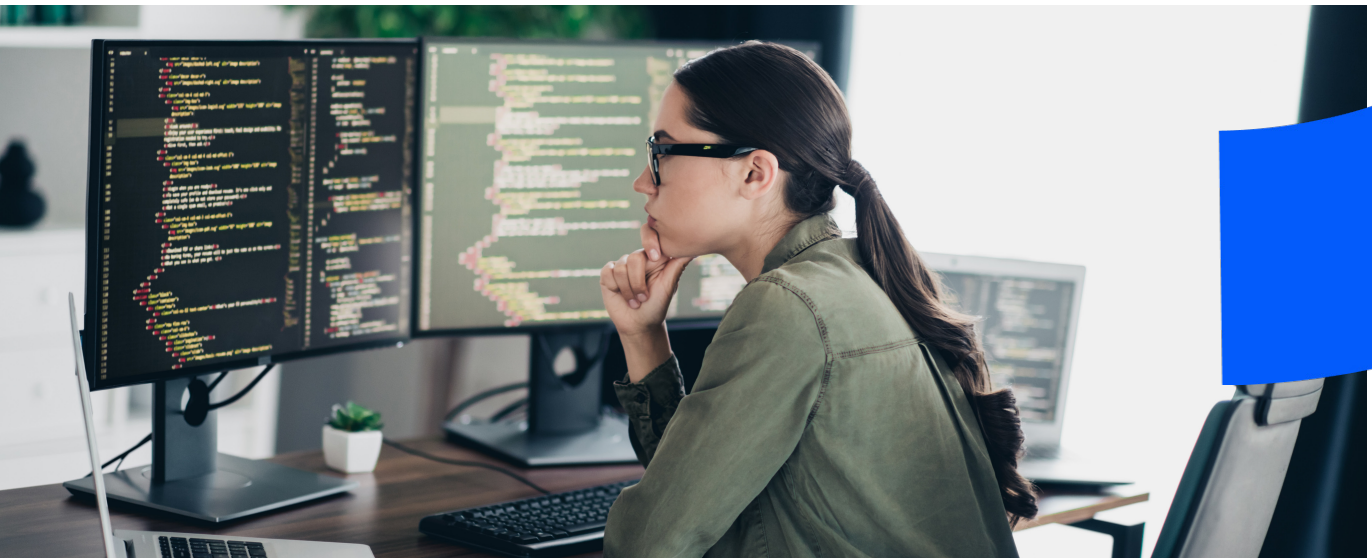

The following figure illustrates these risks.

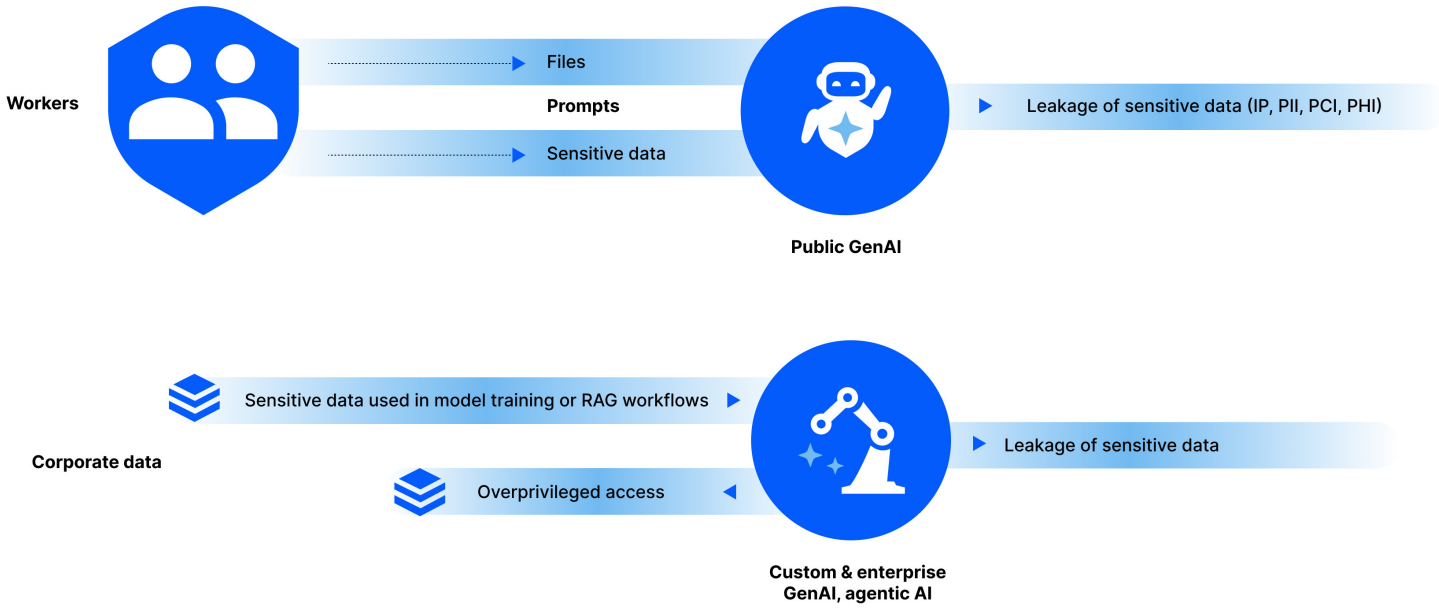

Figure 1: Accidental data loss through AI can happen in many ways.

New exploitation techniques, same old aims

Accidental data loss is not the only category of risk. Threat actors have already developed exploitation techniques tailored to the architectures and behaviors of AI tools. The techniques might be new, but the aims are familiar: exposing and stealing sensitive data for financial gain.

AI applications that are trained on, or have continued access to, broad sets of corporate data are especially vulnerable. Exploitation techniques include adversarial prompting, prompt injection, data poisoning and other forms of memory manipulation. These use clever prompts or malicious training data to override the AI model's intended behavior or access its memory to extract sensitive data.
Threat actors also exploit the architectures of custom and agentic AI applications. These often use APIs and third-party plugins. If those additional components allow too many actions, use weak authentication or have data access that is too wide, they become significant weak points.

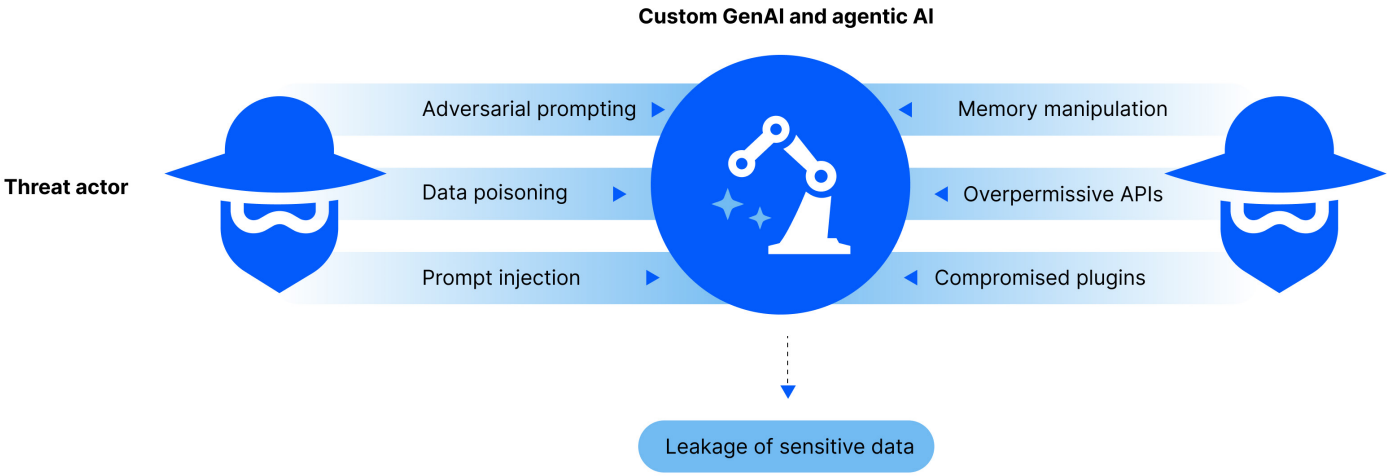

Figure 2: Threat actors have already developed sophisticated exploitation techniques to target custom and agentic AI applications.

Insider risk in the age of AI

Traditional security models have long considered insiders—whether careless, compromised or malicious—a major risk. However, AI increases these risks in the following ways:

- **Careless** insiders might input sensitive data into unauthorized tools or misuse AI capabilities in ways that violate policy.
- **Compromised** insiders (that is, employees whose accounts are taken over) can have their AI access weaponized to enable data discovery, exfiltration or manipulation.
- **Malicious** insiders might deliberately exploit AI access privileges, feed AI tools misleading data or use them against internal systems.

Insider risks from agentic AI

AI agents can make decisions and act autonomously in dynamic environments. Their rapid emergence is giving rise to a new agentic workspace, one in which humans and agents both interact with sensitive data. But, as digital co-workers, agents inherit many of the same risks that affect humans. Like humans, they can be overprivileged. They can also have too much autonomy in deciding what actions to take. This condition is known as *excessive agency*.

Without appropriate guardrails and correctly configured permissions, agents might process manipulated instructions, access malicious URLs, leak sensitive data, use compromised plugins, share credentials, or manipulate systems and data in unintended ways. Depending on the situation, agentic systems can behave like careless or compromised insiders.

In one concerning example, an AI agent made unauthorized changes to live infrastructure and wiped out a software company's production database. The error reportedly deleted data for more than 1,200 executives and 1,190 companies.[2]

The high costs of AI-driven data loss

Just like breaches caused by humans, AI-driven data leakage can have serious consequences. These include financial damage, reputational harm and erosion of customer trust.

And those financial costs can be very high. In 2025, the average cost of a data breach was $4.44 million. Malicious insider incidents averaged $4.92 million.[3]

$4.44M

was the average cost of a data breach in 2025, rising to $4.92M for malicious insider incidents.

Source: IBM

Regulatory consequences continue to grow as well. Breaches can lead to fines, audits and legal exposure under frameworks such as the General Data Protection Regulation (GDPR), Health Insurance Portability and Accountability Act (HIPAA) and California Consumer Privacy Act (CCPA).

For example, under GDPR, fines can be as high as €20 million or 4% of global revenue.[4] HIPAA violations carry penalties up to $1.5 million per violation category per year.[5] And other, AI-focused regulations are emerging. Those include the EU AI Act and US Executive Order on AI.

Simply, organizations must proactively prevent AI-driven data loss before it becomes a costly business crisis.

2. Fortune. "An AI-powered coding tool wiped out a software company's database, then apologized for a 'catastrophic failure on my part'." July 2025.
3. IBM. *Cost of a Data Breach Report 2025*. July 2025.
4. IT Governance. "GDPR Fines & Penalties." Reviewed October 2025.
5. American Medical Association. "HIPAA violations & enforcement." Reviewed October 2025.

Chapter 2

How weak governance opens the door to AI-driven data loss

Public, custom, and enterprise GenAI, as well as agentic AI applications, are becoming critical tools in modern organizations. In a 2024 McKinsey survey, 78% of respondents said their organizations use AI in at least one business function, a significant increase from 55% the previous year.[6]

78%

of respondents in a 2024 survey said their organization uses AI in at least one business function.

Source: McKinsey

But the rate of adoption and broad data access of AI tools make them a growing challenge for security teams. Without clear policies, visibility and controls, organizations risk exposing sensitive data through their own AI use. This chapter outlines some AI security governance deficiencies that can open the door to serious data loss incidents.

Lack of visibility and control of public GenAI use

In McKinsey's 2024 survey, 71% of respondents said their organizations regularly use GenAI, a jump of six percentage points from earlier the same year.[7] High-profile public GenAI tools include ChatGPT, Gemini, Perplexity and Claude. But without governance to monitor and control their use, employees can easily expose valuable or confidential information.

For example, in a widely reported case, Samsung employees leaked a range of company secrets, including proprietary source code, through ChatGPT.[8]

When GenAI use happens outside IT's view, organizations can't track what data workers are submitting. Adding to this problem, some public tools retain or learn from user inputs, making disclosure of confidential information irreversible. Without oversight, user awareness and technical guardrails, every interaction becomes a potential data exfiltration moment.

Use of shadow AI tools

A related problem is the use of "shadow" AI tools that haven't been vetted or authorized by IT or security teams. Workers often use these in good faith. But lack of oversight creates significant risks. Shadow AI tools might have unclear or unsafe data usage policies or weak security. They might be integrated with other third-party tools in ways that further expose enterprise data. And because they're outside the official IT stack, organizations can't monitor, audit or restrict their use.

Unauthorized or overprivileged access

Even authorized tools are risky if they have broad or inappropriate access to internal systems and data. For example, AI agents often require wide access to enterprise data to make decisions. But without strong identity and access management (IAM) policies, they might be overprivileged. This means they can retrieve, reproduce or output data beyond their intended scopes.

6. McKinsey. "The state of AI: How organizations are rewiring to capture value." March 2025.
7. Ibid.
8. Mashable. "Whoops, Samsung Workers Accidentally Leaked Trade Secrets via ChatGPT." April 2023.

Similar risks exist for applications using Open Authorization (OAuth) privileges to interact with enterprise data. As OAuth applications evolve—particularly to introduce AI capabilities—organizations must continue to evaluate their privileges and how they process and handle data.

Clearly, in the absence of strict access controls, organizations are vulnerable to internal leaks and accidental exposure.

Weak data security posture and inadequate classification of sensitive information

To secure its data, an organization must know where all of that data is and how much of it is sensitive. Unfortunately, many still operate with inadequate data discovery and classification systems. According to Gartner, 60% of organizations will fail to realize the anticipated value of their AI use cases by 2027 due to incohesive data governance frameworks.[9] And because AI tools need clear classification to distinguish between public, internal and sensitive data, those governance gaps might lead to damaging leaks.

60%

of organizations will fail to realize the anticipated value of their AI use cases by 2027 due to incohesive data governance.

Source: Gartner

Even when organizations do classify their data, weak data security posture can still lead to exposures. For example, a company might grant Microsoft Copilot permission to index SharePoint and OneDrive data. However, a simple misconfiguration might expose an executive's personal drive to the entire company.

Lack of visibility of training data

As organizations deploy custom, agentic or RAG applications, a new blind spot has emerged: the datasets used to train or augment those applications. Model training or RAG workflows often pull from internal repositories and document stores. But without visibility and controls, data pipelines can easily ingest sensitive or regulated data. For example, in February 2025, security researchers found more than 12,000 API keys and passwords in a public dataset used for LLM training.[10]

Limited adoption of synthetic data

Synthetic datasets are artificial datasets that mimic the properties of real data without containing actual sensitive information. They're a powerful way to reduce privacy risks during model training. A growing array of commercial and open source tools can generate synthetic data of various types. Despite this, its use is not yet widespread. According to Gartner, "although there has been significant investment in synthetic data, user skepticism, reliance on real data, and a lack of standards, trust and awareness still impede its acceptance." [11]

9. Gartner. "Adopt a Data Governance Approach That Enables Business Outcomes." July 2025.
10. The Hacker News. "12,000+ API Keys and Passwords Found in Public Datasets Used for LLM Training." February 2025.
11. Gartner. *Emerging Tech Impact Radar*. 2024.

Chapter 3
A framework for AI security governance

As AI tools become firmly embedded in business processes, companies need a proactive, structured approach to AI security governance. This means developing the roles, policies and oversight mechanisms that regulate AI deployment and use. This chapter provides a framework for effective AI security governance in modern organizations.

Implementing dedicated AI security governance roles

Traditional data governance and security functions—led by roles such as the Chief Information Security Officer (CISO) and cybersecurity architects—have focused on priorities such as securing infrastructure, enforcing access controls and ensuring regulatory compliance. But AI is a different security paradigm. It changes the focus from passive protection to active oversight and dynamic controls.

And because AI has brought fresh risks and challenges, it's natural that new data security and governance roles have followed. Indeed, AI-mature organizations are already implementing roles such as those shown in the table.

Table 1: Example AI security governance roles

ROLE	EXAMPLE ROLE DESCRIPTION
Chief AI Officer (CAIO)	Oversees deployment and governance of AI systems across an organization. Balances innovation with oversight, ensuring that AI makes secure use of enterprise data.
AI Governance Lead and AI Compliance Officers	Ensure that AI systems comply with legal requirements, international standards such as ISO/IEC 42001, and emerging regulations such as the EU AI Act and US Executive Order on AI.
Prompt and Model Compliance Analysts	Review prompts, analyze model outputs for sensitive content and ensure compliance with acceptable AI use policies.
AI Privacy and Security Architects	Design and implement security architectures tailored to AI systems. Focus on securing training data, inference pipelines and deployed models.

Defining an acceptable AI use policy

An acceptable use policy is critical to an organization's overall AI governance program. The policy defines how employees and partners are allowed to use AI tools. It must cover all types of AI in use in the organization, including custom applications, public GenAI tools, enterprise AI tools and agentic AI.

The policy defines user responsibilities and scenarios for appropriate AI use. Those might include drafting reports or emails, creating summaries or researching non-confidential topics. The policy must also describe prohibited uses of AI. Examples are inputting sensitive or valuable information to AI tools, using unapproved applications or bypassing access controls.

Improving visibility and monitoring of AI use

Acceptable use policies are effective only if they're enforceable. Therefore, another critical governance layer is visibility of how AI tools are being used and what data is flowing through them. Key capabilities include:

- **Reviewing user interactions with GenAI.** Use a data loss prevention (DLP) tool to capture the use of public and custom GenAI tools. Gather insights such as most active users, most used sites and the sensitive data types being submitted. Detect the use of shadow AI tools.
- **Elevating monitoring for risky users.** Use behavioral analytics from an insider threat management (ITM) tool to identify users with AI misuse patterns. Elevate monitoring and DLP policies for users in high-risk roles or departments.
- **Building audit trails of user interactions with GenAI.** Use a prompt logging tool to record prompts and responses involving internal or third-party GenAI tools. Include user identities and time stamps for compliance investigations.
- **Identifying and classifying all sensitive information.** Use a data security posture management (DSPM) tool to build an inventory of structured and unstructured data across on-premises and cloud environments. Classify data into tiers (for example, *public*, *internal*, *confidential*, *restricted*) and tag with metadata for governance.
- **Detecting sensitive data used by custom model training and RAG workflows.** Assuming you have classified your data, use a DSPM tool to scan training sets and detect sensitive data being fed into custom models and RAG workflows.
- **Tracking data lineage.** Use a unified data security solution or dedicated data lineage tool to track the origin of data used in model training or inference. Visualize data lineage from source to output, including transformations and policy decisions.

Applying technology-driven policy enforcement

Adding to improved visibility and monitoring, organizations must also consider technology-driven policy enforcement. This means technical controls that enforce usage standards in real time. Important capabilities include:

- **Visualizing and governing data access for AI tools.** Use a DSPM tool to visualize data accessed by AI tools in your environment. Get alerts about unauthorized use of sensitive data. Enforce approval workflows for connecting enterprise data stores to custom model training or RAG pipelines. Use an IAM tool to enforce role-based access control (RBAC) and attribute-based access control (ABAC) for AI applications. Limit access of AI tools to only non-sensitive, authorized data.
- **Blocking sensitive data in GenAI prompts.** Use a DLP tool or secure web gateway to detect and block sensitive data included in prompts made to GenAI tools. Monitor browser extensions and third-party plugins that might bypass upload restrictions.
- **Preventing sharing or processing of sensitive information by enterprise AI tools.** Use a DSPM tool to apply information protection labels that protect data accessed by enterprise copilots.

Driving behavior change

Policies and enforcement technologies are key AI security governance components. But sustainable governance also requires complete, company-wide buy-in. Drive long-term changes in security attitudes and behaviors by:

- **Educating workers on responsible AI use.** Train workers on enterprise-approved AI tools and acceptable use. Share real-world examples of misuse, emphasizing the business and legal implications.
- **Automating training for high-risk users.** Use AI usage reports to identify users that interact heavily with GenAI tools. Trigger real-time coaching or quizzes when risky behaviors are detected. Incorporate security awareness training into remediation plans for DLP or ITM policy violations.

Chapter 4

Securing AI applications and data

To complement strong governance, organizations must also apply rigorous AI security, forming a multilayered strategy. While governance is about oversight, regulation and policy enforcement, AI security involves the technical measures that protect data from unauthorized access, leakage or misuse.

What's more, AI security is also multilayered. It must consider both AI applications and the data those applications use. This chapter delineates those two security goals, describing best practices and technologies aligned to each.

Securing AI applications

Securing AI applications involves protecting models, inference processes and deployment environments from manipulation and compromise. This means protecting the inputs and outputs of AI systems from being exploited or causing harm.

When it comes to securing an AI application, the burden on an organization depends on the basis of that application. The organization might be directly using a third-party AI tool, has integrated a third-party AI tool with its own systems, or has built its own, custom AI application. The following sections describe common security techniques for each of these adoption patterns.

Security techniques for directly using third-party AI tools

Because the most important security concerns are handled by the provider, directly using a third-party AI tool puts the lowest security burden on an organization. However, the techniques shown in the table are still valuable for ensuring secure use.

Table 2: Security techniques for directly using third-party AI tools

TECHNIQUE	DESCRIPTION
Access control	Defines who can interact with the AI, how often and in what ways. Can apply authentication, authorization and rate-limiting rules to protect the AI against malicious use.
Input filtering	Validates, sanitizes and restricts user input if necessary before it reaches the AI. Helps to block malicious or non-compliant prompts.
Output filtering	Filters model outputs to block disclosure of harmful or compliance-violating information. Helps ensure that outputs meet ethical, legal and policy standards.
Security gateways	Act as intermediaries between users and AI systems. Can use context awareness to dynamically enforce policies such as rate limits, content moderation and access control. Can also help to detect misuse patterns.

Security techniques for integrating third-party AI tools

Integrating a third-party AI tool means exposing it to the organization's data and users. It might also mean connecting it with other systems, APIs and services. As a result, the security onus on the organization is greater. In addition to the preceding techniques for directly using third-party tools, the following additional practices are important to consider.

Table 3: Additional security techniques for integrating third-party AI tools

TECHNIQUE	DESCRIPTION
Penetration testing	A type of simulated attack that's applied to the whole AI deployment environment, including APIs, models and endpoints. Helps detect system-wide technical vulnerabilities before real-world attackers do.
Prompt hardening	Involves implementing well-crafted system prompts, which are the predefined instructions that determine the AI's behavior, boundaries and responses. These help models to recognize and resist attempts at manipulation. They're also critical for ensuring appropriate behavior of agentic AI systems.
Red teaming	Simulates attacks on an AI system to uncover vulnerabilities. Tests how a model might be manipulated to disclose confidential data or generate dangerous content. Compared with penetration testing, which tests for specific technical vulnerabilities, red teaming looks for weaknesses across people, processes, controls and technology.

Security techniques for building custom AI applications

Custom AI applications can include custom LLMs, RAG systems and agentic applications. Because an organization handles development and training right through to deployment, these impose the heaviest security burden. As well as securing applications with the preceding techniques, organizations must also build robustness into model training. Adversarial training—and emerging variants of it—is a key technique. This is described in the table.

Table 4: Additional security techniques for building custom AI applications

TECHNIQUE	DESCRIPTION
Adversarial training	Trains a model on a mix of standard data and data designed to deceive it. Improves robustness and security by teaching the model to distinguish between safe inputs and maliciously modified ones.

Securing data in AI use

Custom and enterprise GenAI, as well as agentic AI, commonly train on, or retain access to, broad sets of enterprise data. To prevent breaches, misuse or poisoning of training and inference processes, organizations must ensure their data is secure at rest, in transit and in use.

The following security techniques are all important to consider. They apply regardless of whether an organization is directly using third-party AI tools, is integrating them or is developing its own, custom applications.

Table 5: Techniques for securing data in AI use

TECHNIQUE	DESCRIPTION
Clean rooms	A privacy-preserving technique, sometimes used together with federated learning, that enables multiple parties to contribute data to train a shared model. Each party controls its own raw data in a secure sandbox environment and does not share it with others.
Data access mapping and visualization	Provides visibility of who or what accessed data, when, and from where. Helps security teams detect anomalies and enforce least privilege principles.
Data discovery and classification	Scans on-premises and cloud environments to identify and label sensitive data (for example, IP, PII and PHI) within structured, semi-structured and unstructured formats. Key to ensuring safe data handling in AI applications.
Data loss prevention	Redacts or blocks AI inputs and classifies and labels outputs that contain sensitive data. Protects against disclosure of PII, customer data, trade secrets and compliance-sensitive terms.
Data masking	Obscures data values in non-production environments, while maintaining data usability. Prevents exposure of sensitive data during AI development and testing.
Differential privacy	Introduces statistical noise into datasets or queries to prevent later re-identification. Key to protecting personal data in model training.
Enterprise AI and GenAI data security posture management	Ensures that data stores used by enterprise copilots, custom LLM training and RAG workflows are properly configured and securely accessed for safe AI deployment.
Federated learning	Trains models using decentralized devices or data stores, without transmitting raw data. Helps reduce risks from centralized data.

Table 5: Techniques for securing data in AI use

TECHNIQUE	DESCRIPTION
Synthetic data	Generates artificial datasets that mimic real data without containing actual sensitive information. Used in model training to reduce the risk of privacy violations, while still retaining model accuracy.
Tokenization and encryption	Replaces sensitive data elements (for example, Social Security numbers or credit card numbers) with tokens, or encrypts them, while still preserving the data format. Useful for anonymizing datasets for AI use.

The following figure summarizes the techniques that apply to both securing AI applications and data.

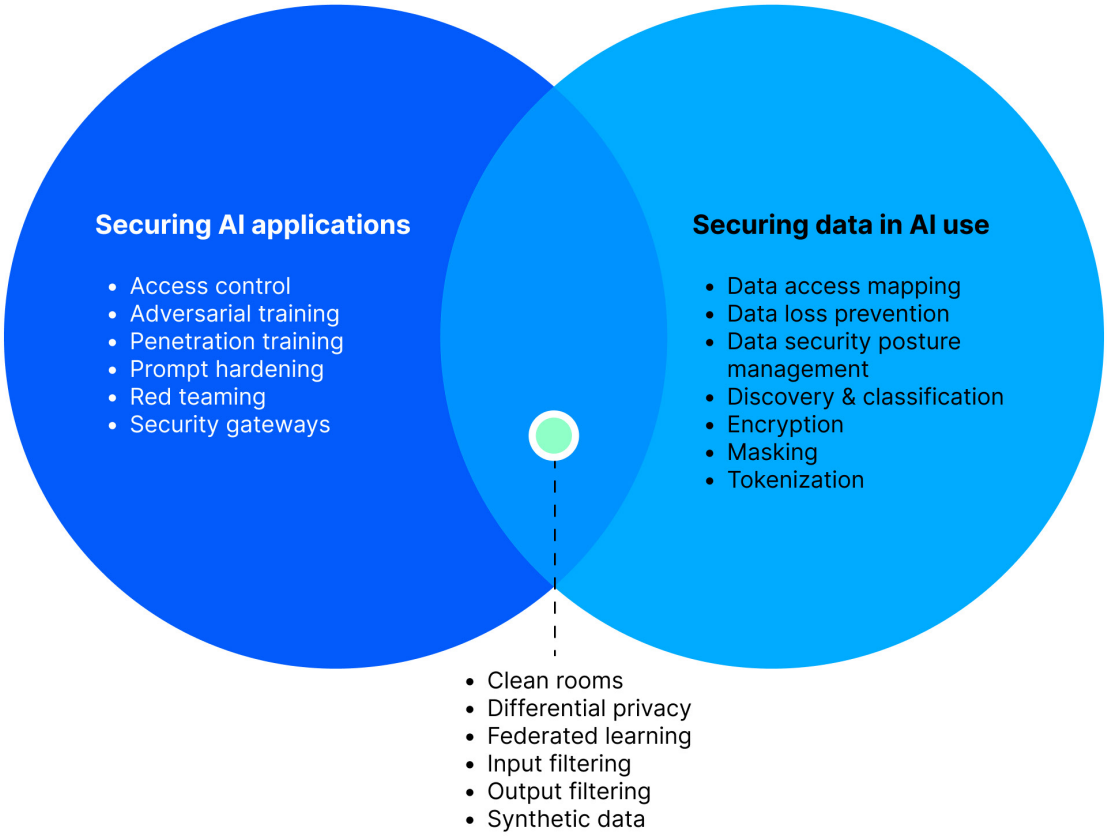

Figure 3: Although they can overlap, securing AI applications and securing data in AI use are separate challenges that require some distinct techniques and tools.

Conclusion

Proofpoint as your foundation for AI security and security governance

As this guide has explored, enterprise data security programs must urgently evolve for the AI age. By applying complementary layers of AI security and security governance, your organization can keep sensitive data protected without hindering AI-driven innovation and productivity.

Combined with Proofpoint ZenGuide™, our security awareness education product, Proofpoint Data Security Complete, Proofpoint Secure Agent Gateway, and Proofpoint AI Data Governance provide a strong foundation for this multilayered approach.

Proofpoint Data Security Complete is a unified solution that consolidates DLP, DSPM, ITM and data lineage in a single architecture, providing streamlined deployment, operation and administration. Proofpoint Secure Agent Gateway is a purpose-built solution that secures sensitive information flowing into and out of every agentic workflow. Working together, these provide the **industry's first solution** for securing data across both people and agents.

Components of our comprehensive solution set are described as follows:

Proofpoint Data Security Posture Management discovers and classifies sensitive data in your cloud and on-premises environments and enhances the security posture of the data used by AI. Detects AI services in your environment and alerts you about unauthorized use of sensitive data. Labels data to protect information accessed by Microsoft Copilot. Secures custom LLMs and applications on cloud AI services by detecting when sensitive data is used in training and RAG workflows.

Proofpoint Enterprise DLP shows who is using AI tools and whether sensitive data is leaking into public tools or custom LLMs. Tracks usage and enforces endpoint DLP policies for more than 600 AI tools. Can block web uploads, copy and paste activity and prompts made to AI websites. Restricts the inclusion of sensitive data in AI prompts.

Proofpoint Insider Threat Management provides visibility of risky behaviors by careless, malicious and compromised users. This can include unusual or risky interactions with sensitive data. Monitors for insider threats with dynamic policies that capture metadata and screen captures

Proofpoint Secure Agent Gateway controls how AI agents access data, monitors agent activity, enforces policies for data use and blocks or redacts sensitive data before it's shared with humans or other agents. Built using Model Context Protocol (MCP), it unifies agent controls with comprehensive data security policies across your environment.

Proofpoint AI Data Governance incorporates data classification, data governance, DLP and security posture management capabilities in a dedicated solution for GenAI governance. Enables the safe use of enterprise copilots and AI applications by identifying sanctioned and unsanctioned use, applying controls to prevent exfiltration and privacy violations, and governing access with automated workflows for security teams and content owners.

Proofpoint ZenGuide transforms employee behavior with continuous learning that adapts to risk. Fosters security consciousness by continuously detecting risks, intervening with guidance, encouraging more secure actions and evaluating effectiveness.

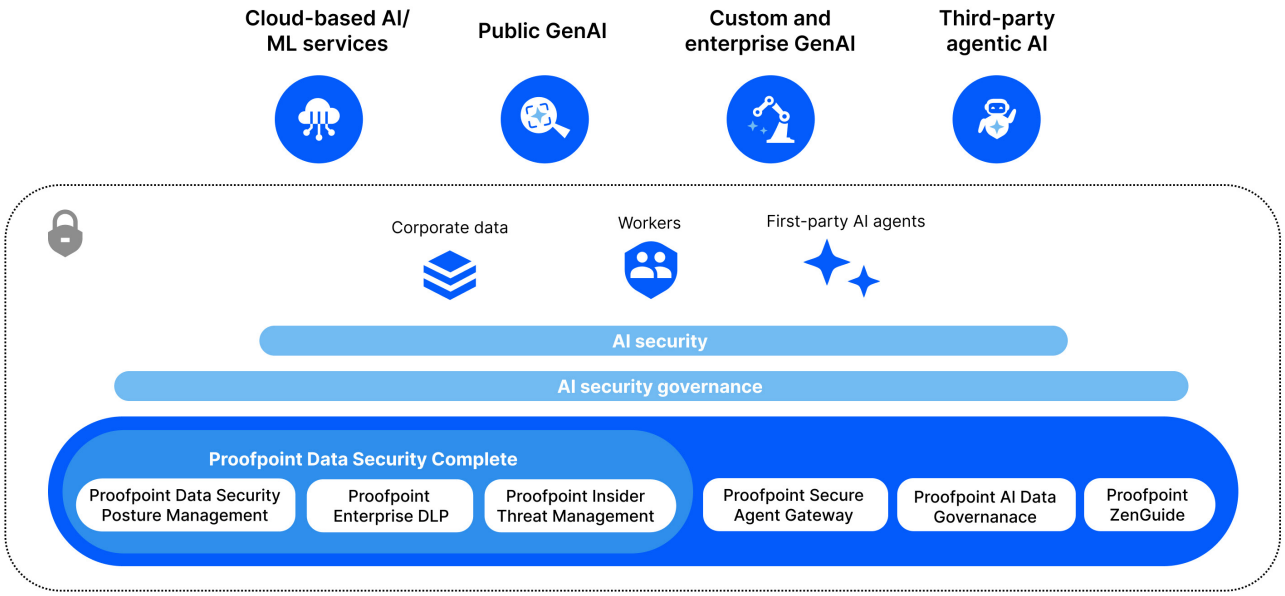

Figure 4: Proofpoint is the foundation of your multilayered approach to AI security and security governance. From the secure environment that our solutions help create, your workers, AI agents and data can interact safely with both third-party and public AI tools.

Your next steps

- To understand how Proofpoint solutions ensure safe GenAI adoption by driving visibility, control and education, read our solution brief.
- To learn more about our unified data security solution, see Proofpoint Data Security Complete or contact your Proofpoint representative to schedule a demo.

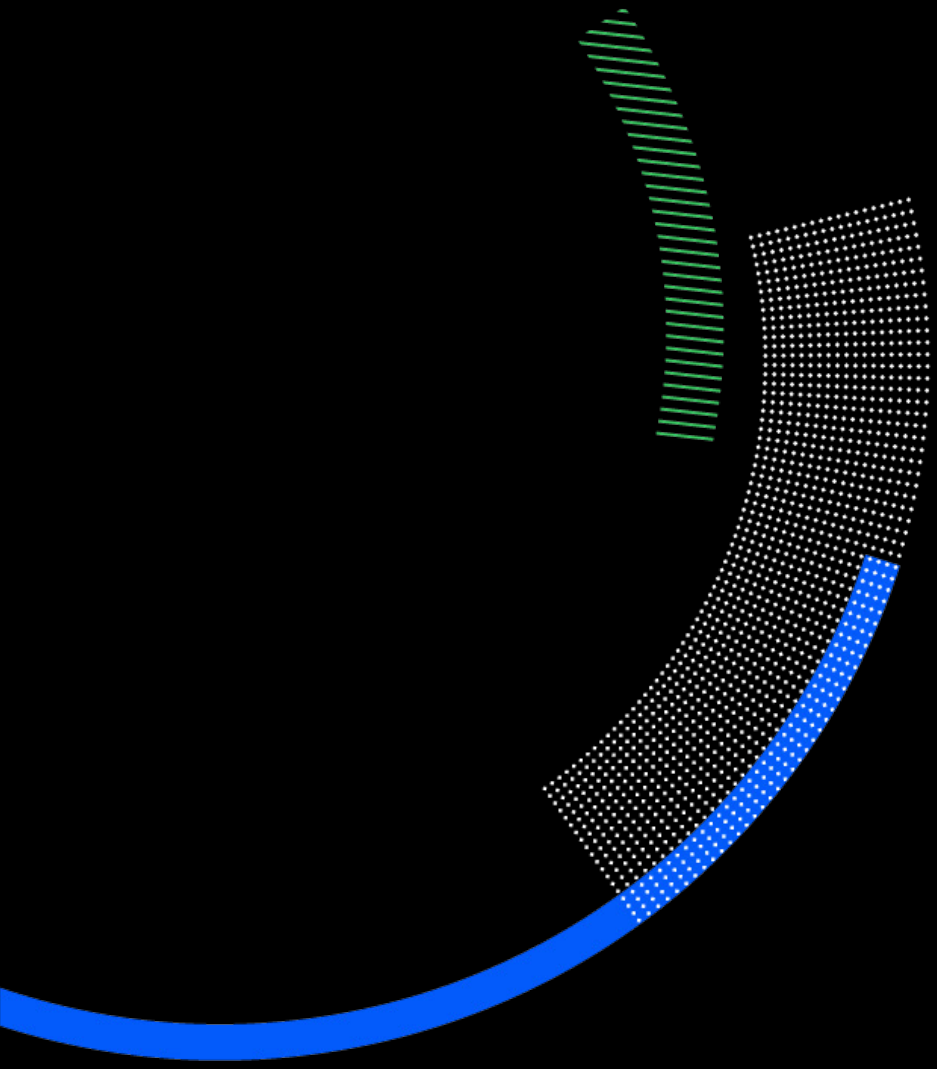

proofpoint®

Proofpoint, Inc. is a global leader in human- and agent-centric cybersecurity, securing how people, data and AI agents connect across email, cloud and collaboration tools. Proofpoint is a trusted partner to over 80 of the Fortune 100, over 10,000 large enterprises, and millions of smaller organisations in stopping threats, preventing data loss, and building resilience across people and AI workflows. Proofpoint's collaboration and data security platform helps organisations of all sizes protect and empower their people while embracing AI securely and confidently. Learn more at www.proofpoint.com.

Connect with Proofpoint: LinkedIn