Skyhigh Security

REPORT

2025 Cloud Adoption & Risk Report

Enterprise AI Adoption & Security Risk

Table of Contents

3 Introduction

4 AI Usage is Growing Rapidly

5 The Shadow AI Problem

6 Increased AI Usage Triggers Compliance Risks

7 Data Leakage via AI Applications

9 Sanctioned AI Apps

11 LLM Risk—Key Element of AI Security

13 Rise of Private AI Applications

14 Conclusion

Introduction

The enterprise cloud landscape has undergone a seismic shift in just the past 12 months. What began as cautious experimentation with generative AI has evolved into full-scale integration across business units. From marketing and legal to customer support and software engineering, AI-driven tools like ChatGPT, Microsoft Copilot, Claude, and Gemini are transforming how work gets done—faster, smarter, and at unprecedented scale.

At the heart of this transformation is the democratization of AI. No longer limited to data scientists and developers, AI is now in the hands of every employee. And while this unlocks massive potential for productivity, creativity, and automation—it also introduces new and rapidly evolving security, compliance, and governance risks.

Today's security leaders are being asked to support AI innovation while mitigating threats, often with limited visibility into how AI tools are being used—or misused—across their organizations. Traditional DLP and access control models simply aren't enough to address the nuances of prompt-based data exposure, AI learning risks, and the Shadow AI explosion.

That's why Skyhigh Security launched the **2025 Cloud Adoption and Risk Report**, powered by anonymized telemetry data across **3M+ users, 40,000+ cloud services, and 2B+ daily events**. This year's report dives deep into:

- The rise of Shadow AI and the unsanctioned use of generative apps in the enterprise,
- The growing attack surface of large language models (LLMs), including prompt injection, hallucination, and data poisoning threats,
- The emergence of private AI workloads and the need for secure LLM deployment,
- And how leading organizations are adopting AI-aware policies to balance innovation with protection.

In short, this report is a blueprint for securing the modern AI-powered enterprise—backed by real-world insights, trends, and best practices from across the globe.

As you turn the pages, we invite you to reflect not only on where your organization stands today, but where it must go next.

Because in the new era of cloud and AI,
security is not a blocker—it's a business enabler.

AI Usage is Growing Rapidly

There is an undeniable increase in productivity with the use of AI apps by corporate employees. These applications have streamlined workflows and reduced manual errors, giving employees time back to focus on high-value, strategic activities.

- An MIT study[1] found that access to ChatGPT reduced the time workers took to complete writing tasks by 40%, while the quality of their output improved by 18%.
- JPMorgan Chase reported[2] that an AI coding tool increased the productivity of its software engineers by up to 20%, allowing them to focus more on high-value projects.

This has led to a significant increase in the use of Shadow AI applications within the enterprise. Skyhigh Security found a 200% increase in the traffic of AI apps in the last year compared to only a 23% increase in traffic to non-AI apps. Data uploaded to AI apps is up 80% while other categories have registered 13% growth.

200%

increase in the traffic of AI apps in the last year compared to only a 23% to non-AI apps

UP **80%**

data uploaded to AI apps, while other categories have registered 13% growth

1 https://news.mit.edu/2023/study-finds-chatgpt-boosts-worker-productivity-writing-0714
2 https://nypost.com/2025/03/14/business/jpmorgan-credits-coding-assistant-tool-for-boosting-engineers-efficiency/

The Shadow AI Problem

A lot of the increase in the use of AI applications in enterprises is in the form of 'Shadow AI' apps, which represents the use of unsanctioned AI applications and chatbots that operate outside of IT oversight.

The Shadow AI problem is an extension of the Shadow IT problem that enterprises have been dealing with for the large part of the last decade. On average, an enterprise uses 3,500 cloud services, a marked increase from the number in 2023, which stood at 2,400 and a dramatically different picture from the average number in 2014, which stood at 831.

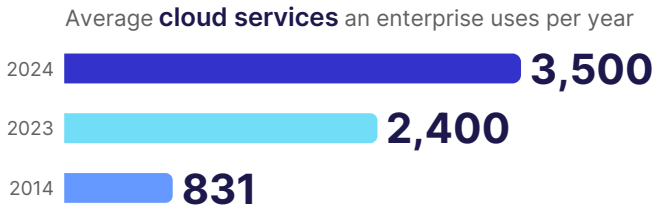

Skyhigh Security's data reveals that the average enterprise uses 320 AI cloud apps.

These include chatbots, content generators, and automation tools. The top AI apps used by enterprises include:

1. Copilot for Microsoft 365
2. OpenAI - ChatGPT
3. Beautiful.AI
4. Hugging Face
5. UiPath
6. UserWay
7. Perplexity AI
8. Sardine
9. Poe
10. Anthropic - Claude
11. Yellow.ai
12. Google Gemini (formerly Bard)
13. Rask AI
14. LivePerson AI
15. Forethought
16. Securiti

Customer Solution

Customers are using Security Service Edge (SSE) solutions to gain full visibility into all AI applications, along with usage metrics such as user counts, upload data, and request count. In addition, the SSE solutions provide risk information calculated using a set of controls. Organizations use this combination of AI usage and risk data to understand AI risk and define governance policies.

Increased AI Usage Triggers Compliance Risks

As enterprises adopt AI solutions across functions and geographies, aligning with regional and industry-specific compliance frameworks is critical to minimize regulatory exposure and avoid costly penalties. Each regulation brings distinct expectations around data handling, auditability, encryption, bias mitigation, and access control.

Top regulations that have expanded their coverage to include AI apps include General Data Protection Regulation (GDPR), Health Insurance Portability and Accountability Act (HIPAA), and EU AI Act. They mandate that AI applications meet key security and privacy requirements that include:

- End-to-end encryption
- Transparent AI model logic and audit trails
- Robust access policies and controls
- Bias testing and mitigation protocols

However, most AI applications today don't have enterprise-grade security requirements, which put enterprises at risk of violating the compliance and regulatory frameworks of the industry they operate in. Skyhigh Security's analysis finds that 95% of AI apps are at Medium or High risk for EU GDPR.

When it comes to other compliance certifications such as HIPAA, PCI, ISO, FISMA, and FedRAMP, which are important to enterprises, **only 22% of all AI applications** are in adherence to one or more of these compliance standards.

When it comes to security controls required by several compliance frameworks, AI applications fall short. Skyhigh Security's analysis finds that 84% AI apps don't support 'Data Encryption at Rest'and that 83% of AI apps don't support integration with Multi-factor authentication (MFA) tools.

84%
of AI apps **don't support** 'Data Encryption at Rest'

Allowing unrestricted use of AI applications not only puts enterprises at the risk of compliance violations, but it also puts them at risk of security threats such as data breaches, phishing attacks, malware, and unauthorized data access.

83%
of AI apps **don't support integration** with Multi-factor authentication (MFA) tools.

Customer Solution

Customers are using Secure Services Edge (SSE) solutions to enforce governance and compliance controls on AI usage within their enterprises. The simplest control includes blocking services that violate policies, for instance, all services that are High Risk and do not support Encryption-at-Rest. But SSE solutions offer more granular controls that include blocking specific activities such as Upload or Download. More recently, there are controls to block specific activities within applications, such as restricting 'Learning' or 'Conversation Sharing' within ChatGPT.

Data Leakage via AI Applications

The risk of sensitive corporate data being exposed through AI applications has become a top concern for large enterprises. As employees increasingly rely on AI tools to summarize presentations, transcribe meetings, and analyze internal communications, Security teams are deeply concerned about the potential leakage of confidential information to unauthorized external systems. A notable example is the incident involving Samsung engineers who unintentionally exposed proprietary semiconductor source code by using a public AI tool for debugging purposes.[3] Skyhigh Security's analysis finds that 11% of the files uploaded to AI have sensitive corporate content in them.

11% of the files uploaded to AI have **sensitive corporate content** in them

This figure is expected to rise as AI tools become more seamlessly integrated into daily workflows. This growing trend is placing significant pressure on enterprise security teams to implement safeguards that prevent data from being inadvertently shared or misused. However, many organizations are in the initial stages of their AI security journey.

Per Skyhigh data, less than 10% of the enterprises have implemented data protection policies and controls on data going into AI apps.

Conversations with enterprise customers indicate that several have established dedicated "AI Security Strategy" teams tasked with developing comprehensive policies over the next 12 to 18 months. These teams are exploring ways to leverage existing SSE capabilities—such as web security, DLP, and zero-trust frameworks—to monitor and control data interactions with AI applications, aiming to strike a balance between innovation and risk mitigation.

Customer Solution

Customers are using SSE solutions to apply or extend DLP existing controls to AI applications. This includes standard out of the box DLP controls to identify sensitive information such as PII or source code, but also advanced controls such as fingerprinting to identify data exfiltration from structured data bases (Exact Data Matching) or unstructured data sources (Index Document Matching). These controls can also be applied on data uploaded from managed and unmanaged devices and on data-at-rest in AI applications.

3 https://in.mashable.com/tech/50407/whoops-samsung-workers-accidentally-leaked-trade-secrets-via-chatgpt

DeepSeek - Driving Shadow AI Use

DeepSeek, a Chinese artificial intelligence startup founded in 2023, has experienced a meteoric rise in popularity. Not only did it surpass ChatGPT[4] to become the highest-rated free app in the U.S. on Apple's App Store, but the AI assistant also had a profound market impact, as major technology stocks experienced significant declines.

In January 2025, Skyhigh saw DeepSeek usage by 43% of Skyhigh customers who uploaded 176 GB of data into the AI chatbot.

Skyhigh Security also found that DeepSeek has significant security vulnerabilities that can put enterprise data at risk.

1. No support for Multi-factor authentication
2. No support for data encryption at rest
3. Allows anonymous use
4. No support for audit logging for users and administrators

From a compliance standpoint, the service is at high risk for EU GDPR and does not provide any evidence of supporting compliance certifications such as HIPAA, PCI, SOC2, and ISO.

4 https://www.cnbc.com/2025/01/27/chinas-deepseek-ai-tops-chatgpt-app-store-what-you-should-know.html

Sanctioned AI Apps

Given the risk associated with unsanctioned AI apps, enterprises are increasingly licensing selected AI applications for employee use to maintain control over data security, compliance, and operational efficiency. By carefully evaluating and licensing trusted AI tools, companies can ensure that these applications meet corporate security standards, comply with data protection laws, and integrate safely into existing IT systems.

Microsoft Copilot is Dominating Enterprise AI Use

Enterprises are increasingly adopting Microsoft 365 Copilot to enhance productivity and streamline operations. Microsoft Copilot seamlessly integrates with familiar Microsoft applications like Word, Excel, PowerPoint, Outlook, and Teams, and enables employees to automate repetitive tasks, generate content, and analyze data more efficiently.

Skyhigh Security's data confirms that Microsoft Copilot is the runaway leader in sanctioned AI apps used by enterprises with 82% of all customers using Microsoft Copilot within their enterprise, up from 18% last year.

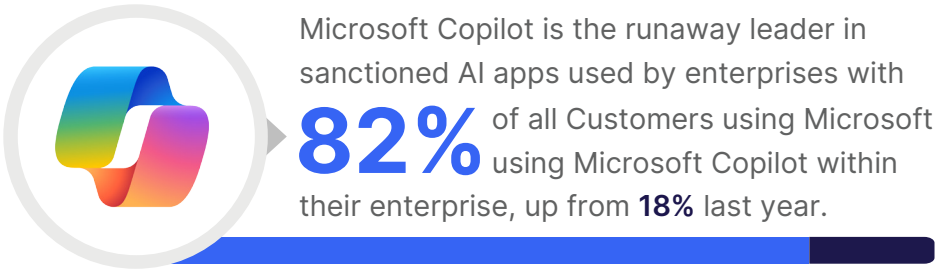

In the last year, the traffic to Microsoft Copilot increased 3,600x, with data uploads increasing 6,000x.

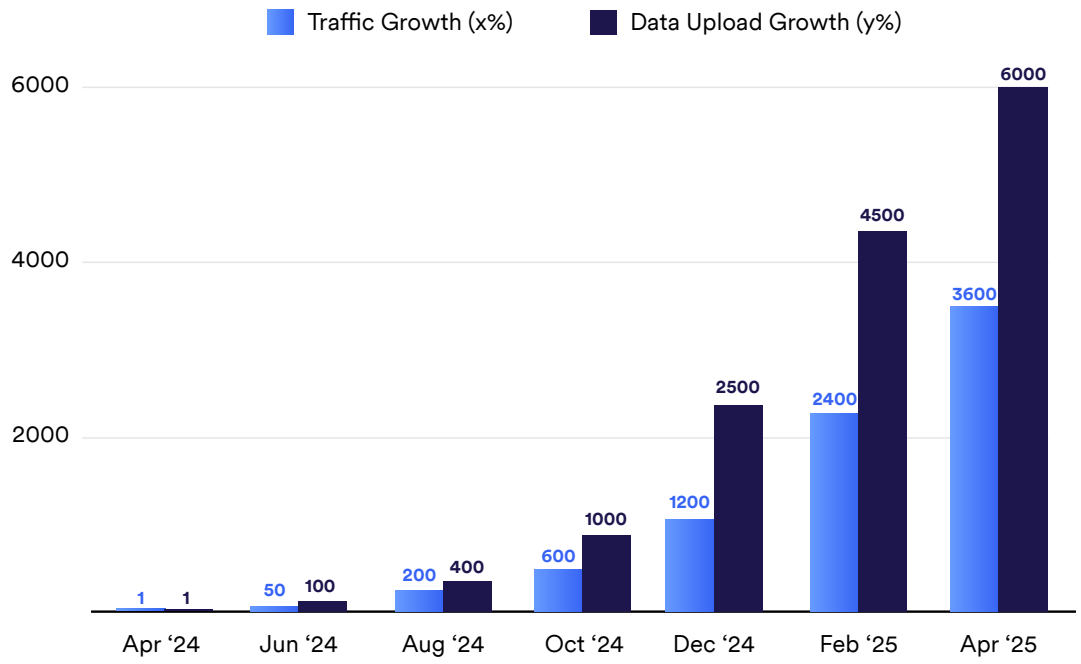

As Copilot adoption accelerates across the enterprise, organizations are increasingly prioritizing the extension of their existing security controls to protect sensitive data within Copilot environments. This includes:

1. **Data Loss Prevention (DLP) for Prompts and Files:** Ensuring that user inputs (prompts) and Copilot-generated outputs are continuously monitored and governed by DLP policies to prevent accidental or intentional exposure of sensitive or regulated data.
2. **Data-at-Rest Scanning:** Implementing scanning mechanisms to identify and classify sensitive information stored within connected applications and repositories, even after being processed by Copilot, enabling proactive risk management.
3. **Prevention of Sensitive Data Ingestion:** Leveraging data classification labels and sensitivity metadata to restrict Copilot from accessing, learning from, or suggesting content that contains confidential or regulated information—thereby preventing leakage or unauthorized use during prompt interactions.

Together, these controls empower organizations to securely embrace AI copilots while maintaining compliance, minimizing data risk, and aligning with their broader data protection strategies.

LLM Risk—Key Element of AI Security

As enterprises evaluate the security of AI applications, one element crucial to this process is the security of the underlying Large Language Model (LLM). This matters because the LLM is at the core of how the application processes, stores, and generates content.

While commonly used AI chatbots have the LLM built into the core of the application, other AI applications could be accessing hosted LLMs via APIs, enterprise licensing, or via private LLM deployments. Enterprises also commonly tailor publicly available LLM models via fine-tuning or Retrieval-Augmented Generation(RAG) to customize the LLM to their use cases.

Irrespective of how the AI application uses the LLM, the risk associated with the LLM is an integral part of the overall risk of the AI application. Common risks associated with LLMs include:

- **Prompt Injection / Jailbreak**

  Example: In February 2023, users discovered ways to bypass ChatGPT's safety controls by using techniques like "Do Anything Now" (DAN) prompts. These prompts manipulated the system into generating harmful content, including phishing emails and violent text.[5]

- **Malware Generation**

  Example: In 2024, cybersecurity researchers from HYAS revealed an LLM-assisted malware called BlackMamba that used generative AI to create polymorphic keylogging code on the fly, making detection extremely difficult.[6]

- **Toxicity**

  Example: Meta's BlenderBot 3, during its public release in 2022, produced toxic and offensive content, including racist remarks, conspiracy theories, and misinformation, due to insufficient content filters and user interaction loops.[7]

- **Bias**

  Example: A 2024 Stanford study found that popular LLMs such as GPT-4 and Claude exhibited racial and gender bias in medical advice outputs, often giving less accurate responses for historically marginalized groups.[8]

5 Wired – "People Are Jailbreaking ChatGPT With a Simple Trick"
6 HYAS – "BlackMamba: AI-Generated Polymorphic Malware"
7 The Verge – "Meta's BlenderBot AI chatbot has some opinions on Mark Zuckerberg"
8 Stanford HAI – "Measuring bias in large language models for clinical use"

Skyhigh Registry tracks the LLM risk of AI applications. Based on Skyhigh Security data, 94% of all AI services are at risk for one of the LLM risk vectors.

94% of all AI services are at risk for **one of the LLM risk vectors.**

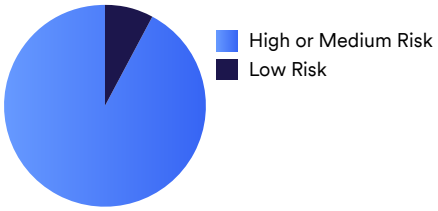

This shows that while capabilities of LLM models are impressive, they are still unpredictable and risky. This should not be surprising, as this technology can still be considered 'early stage.' As enterprises adopt LLM models as part of SaaS or private applications, they need to be cognizant of these risks and implement the necessary controls to protect their corporate data and systems.

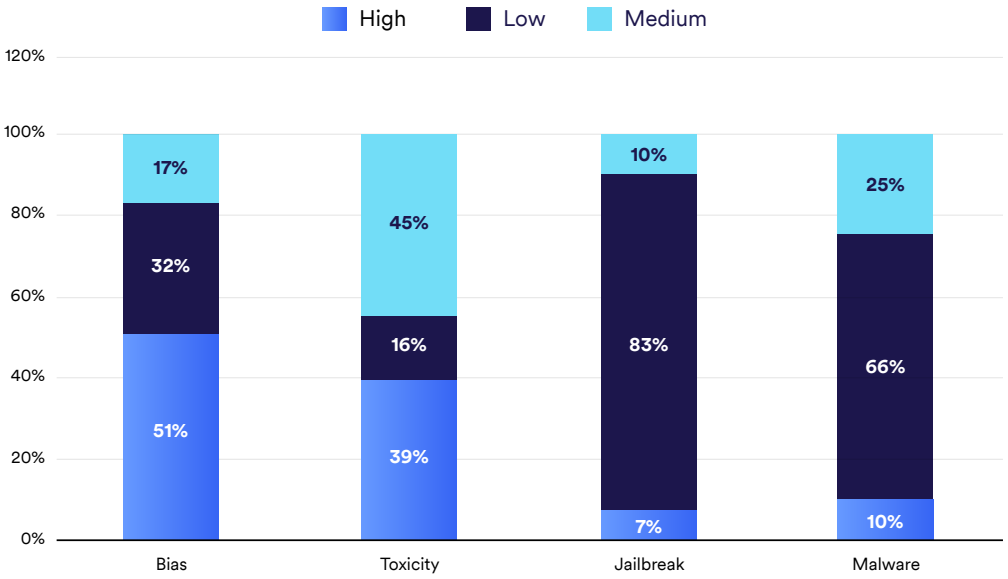

Figure. LLM model risk data sourced from Enkrypt AI.[9]

1. Risk of jailbreak remains the top risk with 90% of the applications being at high or medium risk.
2. The increasing use of AI copilots for code generation and optimization, 76% of the applications being at risk for malware is a key enterprise concern.
3. As enterprises increase their use of sanctioned AI apps and use LLM models in home grown private applications, the risk of toxicity and bias in AI applications jumps to more than 83%.

Customer Solution

Customers are using SSE solutions to track LLM risks across the use of AI apps used by their employees. They are building LLM risk attributes into their governance policies so that AI apps which are risky from an LLM standpoint get blocked or have specific DLP policies applied to prevent exfiltration of sensitive data.

9 https://www.enkryptai.com/llm-safety-leaderboard

Rise of Private AI Applications

Enterprises are increasingly developing private, in-house AI applications as part of a broader trend toward greater control, security, and customization in the adoption of artificial intelligence. These efforts are driven by benefits of data privacy, greater customization, and lower operational costs in the long term.

These private apps are often targeted towards streamlining internal workflows, such as employee chatbots that answer HR related questions or provide IT troubleshooting workflows. Software engineers are using chatbots to assist with tasks such as code completion, bug fixing, and generating infrastructure scripts. Sales teams are using these chatbots trained on company products to reduce RFP response times and efforts.

Many private AI apps cater to an internal audience, a growing number of apps are being built, which are customer-facing. These apps serve multiple use cases, such as to provide enhanced and personalized customer experience. These range from retail chatbots that help customers choose the products that are tailored to their preferences to banking or finance chatbots that offer budgeting, credit, tax planning, and investments.

High Adoption of Private Apps

Skyhigh data shows strong adoption of private AI apps by customers. **78% of Skyhigh customers have deployed private, in-house AI apps.** This increased adoption by large, often conservative enterprises, indicates the significant value that AI is driving for their employees and customers and the competitive pressures they are facing. But the increased use of private apps also indicates the greater regulatory constraints faced by these organizations.

These applications are commonly hosted on public cloud platforms like AWS, Azure, and Google Cloud, which allow them to scale these applications, while also making it easy to deploy sophisticated AI models within virtual private clouds (VPCs), behind firewalls, and under strict access controls. Customers also deploy private apps on their own data centers.

Skyhigh data found that AWS was the platform of choice for AI apps, with 66% of the customers choosing this platform. This included services such as AWS Sagemaker and AWS Bedrock. Azure and GCP had a total of 34% share when it came to AI apps.

Platform	Customer Split
AWS	66%
Azure	25%
Google	9%

Customer Solution

Skyhigh customers are using the Private Access solution to restrict user access to in-house AI apps, scan data for sensitive content. Furthermore, customers are performing data-at-rest scans on source repositories such as S3 buckets to prevent the use of sensitive company data in training AI apps.

Conclusion

Artificial Intelligence is no longer a fringe innovation—it's central to enterprise transformation. From copilots that turbocharge productivity to private AI assistants that streamline operations, AI is driving undeniable value across industries. But with this rapid adoption comes a new wave of risks: unsanctioned Shadow AI use, regulatory non-compliance, data leakage, and security vulnerabilities tied to generative models.

Skyhigh Security's 2025 data makes it clear:

- AI adoption is surging, with traffic to AI apps up 200%,
- Shadow AI use is widespread, with over 320 unsanctioned AI apps being used on average per enterprise,
- And LLM risks are pervasive, with 94% of AI services vulnerable to one or more LLM threat vectors.

While innovation is inevitable, security and governance cannot be optional.

That's where Skyhigh Security's cloud-native Security Service Edge (SSE) platform makes the difference. It provides the visibility, control, and enforcement capabilities needed to enable safe and scalable AI use—whether that's monitoring prompt-level activity, enforcing "no-learning" policies, or ensuring that only compliant AI apps are used across your hybrid workforce.

Forward-looking organizations are:

- Blocking high-risk AI tools like DeepSeek,
- Govern AI copilots like Microsoft Copilot with granular policies,
- Protecting private AI apps hosted on AWS, Azure, or GCP,
- And embedding AI-aware governance policies into their broader security framework.

The enterprises that will succeed in this new era are those who recognize that AI and security must go hand-in-hand. By taking a proactive approach today—powered by the right tools, real-time intelligence, and policy automation—organizations can confidently embrace AI without compromising on trust, compliance, or data protection.

About Skyhigh Security

When your sensitive data spans the web, cloud applications, and infrastructure, it's time to rethink your approach to security. Imagine an integrated Security Service Edge solution that controls how data is used, shared, and created, no matter the source. Skyhigh Security empowers organizations to share data in the cloud with anyone, anywhere, from any device without worry. Discover Skyhigh Security, the industry-leading, data-aware cloud security platform.

For more information visit us at skyhighsecurity.com

Skyhigh Security

3099 North First Street
San Jose, CA 95134
888.847.8766
skyhighsecurity.com

RPT-EN-000040-02