



CYBERSECURITY
PARTNER OF CHOICE



The State of Generative AI

2025



Table of Contents

Executive Summary. 3

 Methodology.3

Growth Observations: AI Is a Force Multiplier 4

 GenAI Traffic Increased by More than 890%4

 Most GenAI Is Used for Writing, Conversing, and Coding5

 Introducing the Top 10 GenAI Apps6

 Grammarly Takes the Top Spot Among Writing Assistants—and All Apps for That Matter.....7

 Microsoft Leads Enterprise Adoption of AI Agents with MS Copilot.....8

 Technology and Manufacturing Use AI Coding Tools a Lot.9

 GenAI Is Reshaping the SaaS Landscape 11

Data and Security Observations: A New Era of Risk. 12

 Shadow AI Is an Emerging Source of Risk12

 Expect at Least Six High-Risk GenAI Apps12

 GenAI-Related DLP Incidents Are Rising.....13

Key GenAI Challenges and Risks: AI’s Growing Pains 14

 Navigating the Sensitive Data Minefield14

 The Evolving AI Regulatory Landscape15

 Staying Ahead of AI Security Blind Spots15

Recommendations and Best Practices: The AI Game Plan 16

Summary: Innovate Fast, but Secure Faster 17

 Authors.....17

About Palo Alto Networks 18

Executive Summary

In our *State of Generative AI* report, the Palo Alto Networks product and research teams explore the adoption and use of generative AI (GenAI) technology across our customers. Our data yielded some noteworthy trends based on the tens of thousands of tenants we observed.

Here's what we found:

- **GenAI traffic surged more than 890% in 2024.** This explosive growth may be attributable to maturing AI models, greater enterprise automation, and higher adoption due to more evident productivity gains. Increased adoption and usage marks a definitive shift from GenAI as a novelty to an essential utility.
- **Data loss prevention (DLP) incidents for GenAI more than doubled.** In 2025, the average monthly number of GenAI-related data security incidents increased 2.5X, now comprising 14% of all data security incidents. GenAI apps amplify a growing data loss vector, as unsanctioned or careless usage can lead to intellectual property leaks, regulatory compliance concerns, and data breaches.
- **Organizations saw on average 66 GenAI apps, with 10% classified as high risk.** The widespread use of unsanctioned GenAI tools, lack of clear AI policies, and pressure to adopt AI quickly—without proper security controls—can expose organizations to significant risk, especially from high-risk GenAI apps.

This report dives into the world of AI adoption to explore its benefits, challenges, and how to balance AI innovation with security to safely transform how businesses operate for years to come. We'll analyze what we're seeing at Palo Alto Networks to provide valuable insights that will position you for success in this AI revolution.

Methodology

We employed a comprehensive approach to analyze GenAI traffic across our servers to help ensure the reliability of our findings. Our methodology compiled historical traffic logs across a sample of 7,051 global customers, spanning the entirety of 2024. Anonymized DLP incidents were also collected and analyzed from January to March 2025. The totality of our analysis was conducted with strict adherence to privacy and security guidelines, ensuring no sensitive customer data was exposed during the research process.

For this report, we observed trends for third-party-provisioned GenAI apps accessed across Prisma® Access and Next-Generation Firewall (NGFW) deployments.

Growth Observations: AI Is a Force Multiplier

Imagine walking out of a meeting where the minutes are already drafted courtesy of a GenAI tool that listened in and summarized everything. Picture an AI agent autonomously managing your sales process—from initial customer outreach to follow-ups—analyzing the propensity to buy, crafting personalized messages, and determining the optimal timing for communication. This is the reality of AI in enterprises today.

Companies are embracing AI at an incredible pace. Think of it like the early days of cloud computing but with even more potential for disruption. For instance, some organizations that leverage GenAI are seeing productivity gains yield an incremental economic impact of up to 40%.¹ With so much promise, companies are exploring off-the-shelf AI solutions—attracted by their rapid deployment, predictable costs, and access to cutting-edge features—in addition to internally developed AI systems and apps.

GenAI Traffic Increased by More than 890%

When ChatGPT launched to the public, it proved to be the fastest aggregator of users to date due to its novelty, broad accessibility, and ease of use.² Today, hundreds of AI models and apps are introduced in timespans of only months, reflecting the breakneck speed at which GenAI apps are hitting the market.




Figure 1. Monthly trend in GenAI transactions

1. *The economic potential of generative AI: The next productivity frontier*, McKinsey & Company, June 14, 2023.

2. Krystal Hu, "ChatGPT sets record for fastest-growing user base - analyst note," Reuters, February 2, 2023.

In 2024, we saw GenAI traffic increase by more than 890%. The notable surge in October 2024 may suggest organizations bringing more GenAI initiatives to production,³ a trend fueled by increased investments arriving in the final quarter. It's no surprise that GenAI apps are being adopted at a record pace and are even being introduced into enterprise workspaces, as with any public consumer application, without the approval or oversight of IT departments.

Most organizations will see 66 GenAI apps on average in their environments. This number can range from a handful of apps to hundreds, depending on the organization's size and needs. Adoption and usage also spike around the release of new AI models or updates as businesses and users evaluate the latest advancements for their productivity and innovation needs.⁴ For instance, with heightened demand following the release of DeepSeek-R1 in January 2025, we saw a significant 1,800% spike in DeepSeek traffic in the two months that followed.⁵

Most GenAI Is Used for Writing, Conversing, and Coding

GenAI apps are usually organized by use case, with each type designed to meet the specific needs of its field.




Figure 2. GenAI transactions distributed by use case

66

Average number of GenAI apps discovered per company

1,800%

Surge in DeepSeek traffic after its R1 release

3. Lev Craig, "Survey: Enterprise generative AI adoption ramped up in 2024," TechTarget, October 31, 2024.
4. Sarah Perez, "ChatGPT doubled its weekly active users in under 6 months, thanks to new releases," TechCrunch, March 6, 2025.
5. Charles Choe and Himani Singh, "DeepSeek Unveiled — Exposing the GenAI Risks Hiding in Plain Sight," Palo Alto Networks, February 28, 2025.

The vast majority (83.8%) of GenAI transactions stem from four core use cases—writing assistants, conversational agents, enterprise search, and developer platforms. These AI tools are popular among employees because they directly address everyday, high-demand tasks that many people and businesses rely on:

- **Writing assistants** help people with various stages of writing, whether it's drafting emails, generating blog posts, or crafting reports.
- **Conversational agents** offer instant responses in natural language to a wide range of queries, making them useful in customer service, learning, and productivity.
- **Enterprise search** makes it easier and faster for employees to find relevant and more personalized information within vast amounts of enterprise data.
- **Developer platforms** provide a suite of tools and services to help developers build, train, and deploy software apps, including modern AI apps.

Less common GenAI tools, like image generators, meeting assistants, video editors, and other niche apps, don't yet serve the everyday needs for most people. As a result, these tools will see more targeted adoption, aimed at specific creative or operational tasks, unlike the aforementioned use cases that have already become an integral part of daily workflows.

It's important to note that many GenAI apps are incredibly versatile and can address more than one use case. For instance, enterprise search and conversational agents share core technologies, such as natural language processing and large language models (LLMs), making them ideal for both enhanced information retrieval and interactive communications.

Introducing the Top 10 GenAI Apps

A top 10 list of the most used GenAI apps provides a different lens into how AI is being leveraged across organizations today.




Figure 3. Top 10 GenAI apps by transactions as a percentage of total SaaS transactions

Grammarly Takes the Top Spot Among Writing Assistants—and All Apps for That Matter

Writing assistants help users with real-time grammar, spelling, punctuation, and style suggestions, making communication clearer and more polished. These AI tools are designed for convenience, often providing desktop versions, browser extensions, and integration with widely used platforms like Microsoft Word, Google Docs, and various email clients.

Grammarly takes the top spot with 39% of AI transactions, showing strong demand and universal appeal with enhancing communications. Whether in emails, reports, or presentations, clear writing is essential for most users across any industry. With digital communications and hybrid work now the norm, we should expect to see AI tools like Grammarly more tightly integrated and continue to become a mainstay for businesses.

While it's clear that writing assistants are helpful, they also come with risks. Many of these risks also apply to GenAI apps in general and include:


- **Data leakage and privacy concerns:** Writing assistants require access to the content you're writing, which can include sensitive or confidential information. Without the proper safeguards, intellectual property and personal data could be exposed, stored, or mishandled by third-party AI providers, leading to potential privacy concerns, sensitive data loss, and noncompliance.
- **Security vulnerabilities:** If an AI writing assistant is integrated into an organization's systems without proper security controls, it could become a vector for cyberattacks. Hackers could exploit weaknesses in the GenAI app to gain access to internal systems or sensitive data.
- **Inappropriate responses:** Writing assistants may inadvertently generate or suggest inappropriate content that goes against organizational policies or ethical guidelines. If such content is used or shared within the organization, it could lead to organizational liabilities, reputational damage, and other potential harms.

Vulnerabilities with AI Writing Assistants

Our research team conducted a curated study on apps that are also commonly used as writing assistants. We found that over 70% of all tested apps were vulnerable to single- and multi-turn jailbreaking techniques that produced responses that had instructions for self-harm, contained discriminatory or offensive content, or provided information on building weapons, and other illegal activities.

As GenAI technology advances, it's highly important for organizations to control access to these powerful AI apps and ensure they have established security guardrails to avoid safety violations in their responses.

Read [Investigating LLM Jailbreaking of Popular Generative AI Web Products](#) for the complete red-teaming results.



Microsoft Leads Enterprise Adoption of AI Agents with MS Copilot

Conversational AI platforms like OpenAI ChatGPT, Google Gemini, and Microsoft Copilot are rapidly evolving beyond their origins as natural language interfaces. While these systems began as tools that could generate content, answer questions, and assist with basic tasks through conversation, they are now transforming into AI agent platforms that allow customers to build their own custom agents.

Sometimes called agentic apps, AI agents are software programs that use AI models to understand instructions and perform tasks autonomously. They often exhibit all or some of the following behaviors:

- Proactively plan and execute multistep tasks.
- Maintain context and memory across multiple interactions.
- Access and use external tools and data sources.
- Make decisions and optimize for goals.
- Perform autonomous actions with minimal, if any, human input.

The combination of nondeterministic outputs with the autonomy to pursue goals, and take actions using high-privileged access to tools, can exacerbate the systemic risks associated with LLMs.

Our research shows that almost half (49%) of organizations use Microsoft Copilot or Copilot Studio. Integrating these copilots across Microsoft 365 allows the copilots to access and use context from various sources within the Microsoft ecosystem—along with connectors to external systems—to provide more relevant and personalized assistance. Copilot works across LLMs, Microsoft 365 apps, and user data—including calendars, emails, chats, documents, meetings, and contacts—to perform tasks that are otherwise difficult or time-consuming. Organizations also use Copilot Studio to build AI agents that can make real-time and autonomous decisions.⁶

Generative AI apps and AI agents inherently carry operational risks—some of which are unique to AI agents, such as:

- **Data exposure and exfiltration:** A copilot's access to sensitive data within enterprise ecosystems can lead to unintended exposure or exfiltration of confidential and proprietary information through attacks such as:
 - › **Tool misuse:** Occurs when attackers exploit an AI agent's capabilities by crafting deceptive prompts, causing unauthorized data access or system manipulation.
 - › **Privilege compromise:** Happens when malicious actors exploit an AI agent's elevated permissions to perform unauthorized actions, making unauthorized activities appear legitimate.
- **Overreliance on AI outputs:** Users might trust GenAI responses without proper oversight, potentially leading to errors or security breaches if the AI outputs are misleading or malicious.

We recently collaborated with Open Worldwide Application Security Project (OWASP) in pioneering AI security research and defining best practices for securing agentic systems.

Read "[Palo Alto Networks & OWASP Collaborate to Secure AI Agents](#)" to learn more about the top AI agent security threats.

49%

of organizations use
Microsoft Copilot or
Copilot Studio

6. Sangya Singh, "Introducing agent flows: Transforming automation with AI-first workflows," Microsoft, April 2, 2025.

Vulnerabilities with AI Agents

Our research team conducted an offensive security test on AI agents built on open-source agent frameworks, as well as out-of-box AI agents. By leveraging techniques, such as prompt injection, tool misuse, intent breaking and goal manipulation, identity spoofing, and agent communication poisoning, our security researchers successfully:

- Accessed and exfiltrated cloud-service account tokens.
- Altered the AI agent's memory, causing it to make flawed decisions and bypass security measures.
- Exploited AI agent tools to gain unauthorized access to sensitive data.

Read the Unit 42® [AI Agents Are Here. So Are the Threats.](#) report on AI agents.



Technology and Manufacturing Use AI Coding Tools a Lot

AI developer tools are revolutionizing software development processes with their ability to enhance coding efficiency, improve software quality, and assist with complex developer tasks. They automate repetitive workflows by offering real-time suggestions and personalized support during the coding process. This enables developers to focus on innovation without taking away from speed to market or quality.

We observe AI coding tools, including Microsoft Power Apps, GitHub Copilot, Hugging Face, Tabnine, and Codeium, used across industries, with predominant adoption in the High Technology and Manufacturing sectors. Approximately 39% of coding transactions stem from these two industries alone. For technology, AI tools can help improve code quality, speed up software delivery for competitive advantage, and free up engineers for more creative tasks. Manufacturing can use GenAI to quickly create and iterate through designs and prototypes, optimize supply chains by predicting demand, and automate quality inspections.

While AI developer platforms and code generators may improve efficiency and productivity, they can introduce vulnerabilities and risk. Some of these risks include:

- **Data exposure and breaches:** Third-party GenAI developer tools can store sensitive inputs—such as proprietary source code—without adequate security or compliance safeguards.
- **Malicious code execution:** AI-generated source code can potentially include malicious code and security vulnerabilities, such as insecure patterns and libraries, which can lead to exploitation.
- **Legal and ethical concerns:** AI developer tools can raise legal concerns, including potential copyright infringement, unclear ownership of generated code, and liability for vulnerabilities or errors.





Figure 4. GenAI coding transactions distributed by industry

Vulnerabilities with AI Developer Tools

In 2024, high-tech companies downloaded ~53 GB and uploaded ~14 GB of data per company to apps belonging to the “Code Assistant and Generator” use case. This is particularly concerning because our researchers were able to exfiltrate data from conversational apps that have the ability to generate code with a 9.9% success rate. What’s more, GenAI apps capable of code generation were found to be vulnerable to malware generation and malicious links within responses.

While there is a significant uptick in apps that employ improved security guardrails against jailbreak attacks, it’s important for security teams to rigorously vet apps and apply comprehensive content filtering wherever applicable.

Read [Investigating LLM Jailbreaking of Popular Generative AI Web Products](#) for the complete red-teaming results.



GenAI Is Reshaping the SaaS Landscape

Modern GenAI apps use dynamic and adaptive systems that integrate predictive analytics, real-time data synthesis, and automation. These advancements enable hyperpersonalized solutions that have massive potential to disrupt traditional SaaS. We should anticipate an overall reshaping of the SaaS landscape as AI-native platforms start to emerge as category leaders.

In 2024, enterprise use of GenAI across our customers averaged 32% growth month over month, compared with 20% for SaaS. What's more, the ratio of GenAI transactions as a percentage of SaaS also increased from 1% to 2% on average, suggesting a shift in how businesses may prioritize digital tools with AI becoming mission-critical for productivity gains.

This trend also implies that enterprises may begin reallocating IT budgets to GenAI solutions that deliver a faster return on investment.

32%

Average monthly
growth in
GenAI transactions

Data and Security Observations: A New Era of Risk

As GenAI adoption grows, so do its risks. Without visibility into GenAI apps—and their broader AI ecosystems—businesses can risk exposing sensitive data, violating regulations, and losing control of intellectual property. Monitoring AI interactions is no longer optional. It’s critical for helping prevent shadow AI adoption, enforcing security policies, and enabling responsible AI use.

Shadow AI Is an Emerging Source of Risk

Shadow AI refers to the unauthorized use of AI tools by employees without the knowledge, approval, or governance of its IT departments. Operating outside IT, shadow AI creates blind spots where sensitive data might be leaked or even used to train AI models. As unauthorized AI tools—from writing assistants to coding apps—become more prevalent, organizations face escalating risk exposure.

Expect at Least Six High-Risk GenAI Apps

GenAI apps are considered high risk when they pose significant threats across security, compliance, and operational dimensions. Key factors that increase risk include attributes like data input/output types, whether user-submitted data can train its AI models, adherence to various regulatory requirements, and other elements like authentication, data privacy, and encryption.

Table 1. High-Risk GenAI Apps by Industry and Block Rate			
Top 10 Industries	Average Number of High-Risk Apps	Top 10 GenAI Apps Blocked	Block Ratio
1. Mining	10	1. TinyWow	36.52%
2. Insurance	8.5	2. Amazon Bedrock	30.06%
3. Professional & Legal Services	8.3	3. Anyword	29.57%
4. Financial Services	8.2	4. Sembly AI	25.40%
5. Manufacturing	8.2	5. CapCut	24.81%
6. Healthcare	8.1	6. Cohere Coral	24.81%
7. Utilities & Energy	7.9	7. Amazon Transcribe	21.79%
8. Aerospace & Defense	7.8	8. Jasper	21.77%
9. Transportation & Logistics	7.7	9. Lightning AI	20.45%
10. State & Local Government	7.4	10. Writesonic Chatsonic	19.58%

We observe an average 6.6 high-risk GenAI apps per company. The level of exposure to high-risk apps can vary according to industry, digital maturity, and existing security controls. For instance, we see that the mining industry has the highest exposure to high-risk GenAI apps perhaps due in part to unique operational demands, as well as insufficient training and awareness.

It's also notable that, among all GenAI apps, TinyWow had the highest block ratio at 36%—the percentage of usage incidents where access was restricted or blocked. Free GenAI apps like TinyWow may be popular for its unrestricted accessibility and ease-of-use, but they lack the advanced security features organizations need to help protect data and enable compliance.

What's certain is that IT security teams face an urgent need for greater visibility and control over how GenAI tools are being used.

GenAI-Related DLP Incidents Are Rising

Implementing data loss prevention to help identify and block sensitive data from being transferred to GenAI apps is crucial for maintaining organizational security. You need the ability to catch those moments when someone might inadvertently paste confidential company data or proprietary source code into third-party GenAI apps.




Figure 5. GenAI-related DLP incidents as a percentage of all DLP incidents

So far in 2025, we're seeing the total number of GenAI-related DLP incidents increase more than 2.5X, now comprising 14% of all DLP incidents across our SaaS traffic.

For enterprise IT teams, this trend demands proactive measures, such as deploying advanced DLP tools, implementing secure AI solutions, and educating employees on safe AI use. Continuous monitoring of GenAI traffic and enforcing strict access controls are essential to help prevent sensitive data from being inadvertently exposed or misused.

6.6

Average number
of high-risk GenAI
apps per company

Key GenAI Challenges and Risks: AI's Growing Pains

Let's face it—AI is like that shiny new toy everyone wants to play with. However, we need to be aware of the potential pitfalls and risks that come along with it. As more companies experiment with third-party GenAI apps, it's crucial to understand its evolving risk landscape:⁷

- **Lack of visibility into AI usage:** Shadow AI makes it difficult for security teams to monitor and control how GenAI tools are being used across the organization.
- **Unauthorized access and data exposure:** Difficulty restricting access to GenAI tools via personal accounts or detecting (and blocking) when end users upload sensitive data.
- **Insecure AI interactions:** Jailbroken or manipulated AI models can respond with malicious links and malware, or allow its use for unintended purposes.
- **Proliferation of plugins, copilots, and AI agents:** Complex AI ecosystems with browser plugins, AI agents, bots, and copilots create an overlooked "side door."

Navigating the Sensitive Data Minefield

Traditional cloud access security broker (CASB) and DLP tools are not designed to help mitigate the unique data risks that stem from GenAI apps. Traditional methods struggle to keep pace with the rapid proliferation of GenAI apps, along with their unique characteristics and evolving AI ecosystems. What's more, unstructured AI interactions require contextual understanding to accurately identify and classify sensitive data.

Enterprises must, therefore, implement robust safeguards around GenAI app classification, user access controls, and AI-specific DLP capabilities to help mitigate its risks while enabling secure AI adoption. Table 2 shows a few of these data security risks.

Table 2. GenAI Security Risks		
Data Exposure	Privacy Breaches	Unauthorized Access
GenAI apps require vast amounts of data to function effectively. Sensitive data, if not properly controlled, can be inadvertently leaked to GenAI apps.	Leaking sensitive customer or employee data about individuals or groups to third-party GenAI apps can result in privacy violations if not carefully managed.	Hackers could potentially gain access to sensitive data that was inadvertently used to train third-party AI models, leading to severe consequences.

7. For actionable insights into GenAI risks, see [The C-Suite Guide to GenAI Risk Management](#).

The Evolving AI Regulatory Landscape

As AI becomes more prevalent, governments and regulatory bodies are scrambling to keep up. This creates regulatory uncertainty for businesses looking to develop or deploy these tools for their employees. Regulations around the world will accelerate with more investment pouring into AI. The EU is applying provisions of the AI Act in phases.⁸ China already requires labeling of AI-created content⁹ and is carefully looking to augment the country's regulations with focus on GenAI.¹⁰

Some AI-related compliance concerns include:

- **Evolving regulations:** With AI regulations, what's compliant today might not be tomorrow. For example, the EU's proposed AI Act could significantly impact how companies develop and use AI systems.
- **Data protection laws:** Data protection regulations, including but not limited to the GDPR and CCPA, require organizations to be extremely careful on how they handle and share personal data with respect to GenAI apps used in particular contexts.

Staying Ahead of AI Security Blind Spots

In today's AI-first world, gaining visibility and control of every GenAI app used by an organization isn't just important—it's nonnegotiable. Why? Because these tools can easily become the biggest security blind spot for your organization. Imagine your workforce sharing sensitive trade secrets or source code to unapproved AI platforms, or imagine employees using unvetted GenAI tools that are vulnerable to poisoned outputs, phishing scams, or malware disguised as legitimate AI responses.

Enterprise organizations face mounting challenges in defending against AI-based threats, especially with the rapid, uncontrolled rate of GenAI adoption and evolving attack methods. The problem is intensifying as GenAI adoption accelerates and new vulnerabilities—like susceptibility to prompt injection attacks—emerge.¹¹

In this era of AI, solutions with zero trust architectures, real-time data monitoring, and AI capabilities are now critical to closing the gap. Without proactive governance and advanced threat detection, organizations risk falling behind.

8. "Implementation of the EU AI Act," EU Artificial Intelligence Act, 2025.

9. Yan Luo and Xuezi Dan, "China Releases New Labeling Requirements for AI-Generated Content," Covington, March 18, 2025.

10. Mark Greeven, "China And AI In 2025: What Global Executives Must Know To Stay Ahead," Forbes, December 23, 2024.

11. *Securing GenAI: A Comprehensive Report on Prompt Attacks, Taxonomy, Risks, and Solutions*, Palo Alto Networks, April 9, 2025.

Recommendations and Best Practices: The AI Game Plan

Your organization must adopt a proactive, multilayered approach to GenAI governance to effectively help mitigate AI risks. First, you need comprehensive visibility, which is essential to identify shadow AI usage and track data flows in and out of third-party GenAI apps. You must pair this with strict access controls to prevent or limit the unauthorized sharing of sensitive information. In addition, you must implement real-time monitoring capabilities to continuously analyze both inputs and outputs—scanning for malware in AI-generated files, detecting potential data leaks in employee prompts, and flagging suspicious user activities.

Some of the GenAI security best practices include:

- **Understand GenAI usage and control what is allowed.** Implement conditional access management to limit access to GenAI platforms, apps, and plugins based on users and/or groups, location, application risk, compliant devices, and legitimate business rationale.
- **Guard sensitive data from unauthorized access and leakage.** Enable real-time content inspection with centralized policy enforcement across the infrastructure and within data security workflows to help prevent unauthorized access and sensitive data leakage.
- **Defend against modern AI-based cyberthreats.** Implement a zero trust security framework to identify and block highly sophisticated, evasive, and stealthy malware and threats within GenAI responses.

Mitigating risk from GenAI apps is crucial due to its potential to expose sensitive data and enable adversarial attacks. Without proper IT security oversight, unauthorized use of GenAI apps can ultimately lead to an inadequate security posture and sensitive data loss.



Summary:

Innovate Fast, but Secure Faster

The explosive growth of GenAI in 2024 has fundamentally altered the digital landscape for enterprise organizations. On one hand, GenAI unlocks innovation, accelerates competition, and opens new opportunities; but on the other hand, the proliferation of unauthorized AI tools are exposing organizations to greater risk of data leakage, compliance failures, and security blind spots.

As organizations race to integrate GenAI tools, many will struggle with a lack of governance, inconsistent security standards, and unmanaged risks. Simply put, don't let the speed of innovation outpace your safeguards. To harness the full potential of GenAI while mitigating its risks, prioritize strong data controls, access management, and ongoing employee training to lay the foundation for secure GenAI adoption.

Palo Alto Networks Secure AI by Design portfolio helps organizations address the critical need for robust security in the face of rapid AI adoption. [AI Access Security™](#) enables organizations to safely access and use third-party GenAI apps by reducing data and security risks. It provides real-time visibility into GenAI adoption and usage, enforces access controls, provides advanced threat protection, and prevents sensitive, proprietary data from being inadvertently leaked to unauthorized GenAI apps. [Prisma AIRS™](#) empowers organizations to protect their own AI ecosystems, including AI apps, agents, models, and data.

Join us in shaping a future where the transformative power of AI can be harnessed safely and responsibly. Now, you can secure your company's GenAI usage and empower your business to capitalize on its benefits—without compromising on security.

Connect with an expert to get a [personalized demo of AI Access Security](#) or try our [interactive product tour](#).

Authors

Charles Choe

Gurpreet Kaur Khalsa

Jiangnan Li

Yongzhe Huang

About Palo Alto Networks

As the global cybersecurity leader, Palo Alto Networks (NASDAQ: PANW) is dedicated to protecting our digital way of life via continuous innovation. Trusted by more than 70,000 organizations worldwide, we provide comprehensive AI-powered security solutions across network, cloud, security operations and AI, enhanced by the expertise and threat intelligence of Unit 42. Our focus on platformization allows enterprises to streamline security at scale, ensuring protection fuels innovation.

Explore more at www.paloaltonetworks.com.

3000 Tannery Way
Santa Clara, CA 95054


Main: +1.408.753.4000

Sales: +1.866.320.4788

Support: +1.866.898.9087

www.paloaltonetworks.com

© 2025 Palo Alto Networks, Inc. A list of our trademarks in the United States and other jurisdictions can be found at <https://www.paloaltonetworks.com/company/trademarks.html>. All other marks mentioned herein may be trademarks of their respective companies.



CYBERSECURITY
PARTNER OF CHOICE

