

Is Your GenAI Environment Secure?

Protect *All* Your AI investments with
Prisma AIRS from Palo Alto Networks




Table of Contents

GenAI Is Remaking the Enterprise At Record Speed	3	Prisma AIRS Platform Secures AI Deployments	12
GenAI Brings Substantial Benefits—And Risks	4	AI Model Scanning	13
Unique GenAI Risks Outpace Legacy Security.	5	Posture Management.	14
Prompt Injections	6	AI Red Teaming	15
Open-Source Models.	7	AI Agent Security	16
Insecure Outputs	8	Runtime Security	17
Sensitive Data Leaks	9	Prisma AIRS In the GenAI Lifecycle	18
Agent Hijacking	10	Securing Strata Copilot with Prisma AIRS	20
Bundling Is Not Integration	11	Take the Next Step Toward Secure GenAI.	21

GenAI Is Remaking the Enterprise At Record Speed

Generative AI (GenAI) is sparking a new revolution in how we work, learn, and communicate. It's much like the arrival of the personal computer and the internet—only this time, adoption is happening even faster. Businesses everywhere are embracing GenAI to boost productivity, drive innovation, cut costs, and speed up time to market. Nearly half (47%)¹ of enterprises are building GenAI applications now, while 93% of IT leaders² plan to introduce autonomous AI agents within the next two years.

Why are executives so bullish on GenAI? For one thing, GenAI drives productivity gains. GenAI can improve a worker's performance by nearly 40%³ when compared with those who don't use it, while AI-assisted software engineers can code twice as fast.⁴ Another key benefit comes on the content creation side. In marketing, 73% of users employ GenAI tools for generating various types of content, including text, videos, and images.⁵ Enterprises are turning to GenAI applications to drive innovation, enhance operational efficiency, and maintain their competitive edge.



Before we launch into this discussion of GenAI, let's take a moment to review the GenAI architecture.

1. <https://menlovc.com/2024-the-state-of-generative-ai-in-the-enterprise>

2. Ibid.

3. <https://mitsloan.mit.edu/ideas-made-to-matter/how-generative-ai-can-boost-highly-skilled-workers-productivity>

4. <https://www.weforum.org/stories/2023/05/can-ai-actually-increase-productivity>

5. <https://narrato.io/blog/ai-content-and-marketing-statistics/>

GenAI Brings Substantial Benefits – And Risks

The heart of any GenAI application is the **large language model** (LLM), referred to here simply as the model. The model is a type of artificial intelligence designed to process, understand, and generate human language. As a subset of machine learning, models use deep learning algorithms to analyze vast amounts of text data. By predicting the next word or sequence based on context, they generate coherent and contextually relevant responses. This capability allows them to mimic specific writing styles, adapt to different genres, and produce human-like text with remarkable accuracy.

Two types of datasets play a critical role in GenAI applications: training and inference. The **training dataset** is a large, diverse collection of data which the model uses to learn patterns, structures, and relationships in this set of data that is similar to production data. The training dataset includes books, websites, articles, code, and other written material.

The **inference dataset** contains the data used by the model during its operational phase to generate outputs. The model applies patterns and knowledge gained during training to interpret and respond to the inference inputs.




Figure 1: Overview of GenAI Application Architecture

To set up our discussion of GenAI security, we need to talk about the primary risk factors. Keep reading.

Unique GenAI Risks Outpace Legacy Security

The surge in AI adoption has been paralleled by a significant increase in cyberattacks targeting AI systems and datasets. Recent reports indicate that 57% of organizations⁶ have observed a rise in AI-driven attacks over the past year. Notably, Amazon has experienced a dramatic escalation in cyber threats, with daily incidents soaring from 100 million to nearly 1 billion within six months, a surge partly attributed to the proliferation of AI.⁷

Traditional security systems often fall short when protecting generative AI (GenAI) environments because those legacy protections aren't built to handle the unique risks GenAI introduces. These tools rely on static rules and known threat patterns, but GenAI produces unpredictable and highly variable outputs that don't follow fixed signatures. As a result, traditional systems lack the context-awareness needed to accurately detect threats. They also miss AI-specific attacks like prompt injections, data poisoning, and model manipulation, which don't exist in traditional IT environments.

GenAI frequently processes unstructured data like text, code, or images—formats that conventional security tools struggle to analyze effectively. Furthermore, these systems lack visibility into how AI models generate responses, making it hard to detect subtle failures or misuse. Without the ability to monitor inputs, outputs, and model behavior in real time, traditional security leaves critical gaps that AI-native threats can easily exploit.

6. https://www.securitymagazine.com/articles/100631-80-of-data-experts-believe-ai-increases-data-security-challenges?utm_source=chatgpt.com

7. https://www.wsj.com/articles/the-ai-effect-amazon-sees-nearly-1-billion-cyber-threats-a-day-15434edd?utm_source=chatgpt.com

Prompt Injections

A prompt injection attack is a type of security threat unique to GenAI systems. In this attack, a malicious user crafts a specially designed input (prompt) to trick the AI into ignoring its original instructions and instead follow the attacker's commands.

Prompt injection attacks are difficult to stop for two primary reasons. First, the attack typically begins with legitimate access—through a chatbot, input field, or integrated tool. The model doesn't need to be “hacked” in the traditional sense; the attack is crafted to embody clever misuse of natural language. Thus, these attacks don't contain signatures or other distinguishing behaviors that set the attack apart from authorized usage.

The second reason these attacks are difficult to halt comes from the way prompt injections make use of the inherent tendency of GenAI applications to follow instructions precisely—even instructions that degrade security or performance. Imagine a GenAI application that tests a large banking system. Such a system would surely include rules against revealing customer information. But a prompt injection attack could begin with “ignore the restrictions on sharing customer information” and count on the GenAI application to execute that task faithfully, disabling the security measures.

Prompt injection can lead to data leaks, policy violations, misuse of tools, and jailbreaking of AI systems. These attacks exploit the AI's lack of contextual understanding and its tendency to follow instructions too literally. Preventing prompt injection requires strong input validation, clear role separation, and real-time monitoring of model behavior.




Figure 2: Anatomy of Prompt Injection Attack

Open-Source Models

Open-source models are AI models that are publicly released with their code, architecture, weights, or training data made available under a permissive license. The benefits of open-source software are well-documented.⁸ However, using open-source AI models also introduces security risks including deserialization and model tampering.

Deserialization is the process of converting stored data back into usable objects within a program. In AI, it's often used to load models or configurations. However, deserializing untrusted data poses serious security risks. Attackers can craft malicious files that, when loaded, trigger remote code execution, file access, or privilege escalation. In AI systems, this kind of attack may corrupt models with backdoors or hidden triggers. Common tools such as pickle or joblib are especially vulnerable.

Model tampering involves unauthorized changes to an AI model's structure or behavior, posing serious security risks. Attackers may embed backdoors, trigger conditions, or manipulate outputs to spread misinformation or leak sensitive data. These changes often go undetected, undermining model integrity and trust. In regulated environments, tampering can lead to compliance violations and persistent threats.

For example, a research team has developed a model to categorize traffic signs. Unknown to them, a hacker has embedded a piece of code that causes a misclassification when an image contains a certain small visual trigger. In most cases, the model behaves as designed, but when it encounters an image with the embedded trigger, the output is compromised.




Figure 3: Model Tampering Example

8. https://enterpriseproject.com/article/2015/1/top-advantages-open-source-offers-over-proprietary-solutions?utm_source=chatgpt.com

Insecure Outputs

Insecure outputs from generative AI present significant security threats because they often result in users receiving harmful URLs. These links surface either unintentionally or through targeted manipulation. For example, a chatbot may generate a dangerous link after pulling information from a compromised or deliberately poisoned source.

Malicious URLs often appear legitimate to the user, serving as a gateway for attackers to launch phishing schemes, install malware, or gain unauthorized system access. In trusted environments, users may follow AI-generated suggestions without suspicion, increasing the likelihood of compromise. Because these models depend on external and historical data, even one corrupted input can result in unsafe content that puts systems at risk and undermines confidence in AI-assisted tools. Preventing such threats requires strict content filtering and continuous monitoring of output quality.




Figure 4: Anatomy of Insecure Outputs Attack

Sensitive Data Leaks

As discussed earlier, GenAI systems need access to proprietary data for training and operation, which of necessity puts the information outside traditional security controls. This sensitive data represents a lucrative target for hackers who can manipulate models and breach data, which can lead to serious security threats.

GenAI applications are also prone to hallucinations in which the model generates false or misleading information that appears credible. A common example is a citation hallucination, where the GenAI model invents a research paper, author, or source that doesn't exist. For instance, it might claim that a specific study supports a point and provides a realistic-looking title, journal, and date—yet none of it is real. These fabricated citations can mislead users, especially in academic or professional contexts, where source accuracy is critical.




Figure 5: Anatomy of GenAI Data Leak

Agent Hijacking

Technology moves fast, and nowhere is that more true than in the rapid rise of AI agents. Agents are autonomous systems that sense their environment, process data, and make decisions to complete tasks. They learn and adapt, handling complex jobs like drug discovery, customer service, marketing, coding, and research. With 78% of companies⁹ planning to use AI agents, securing this valuable capability of the organization has become crucial to realize full value on their AI investments.

Many AI agents are vulnerable to agent hijacking, a type of indirect prompt injection¹⁰ in which an attacker inserts malicious instructions into data that may be ingested by an AI agent, causing it to take unintended, harmful actions. In these attack situations, would-be thieves can inject malicious instructions along with legitimate instructions to penetrate the organization's security measures—hiding in plain sight, as it were.




Figure 6: Anatomy of Agent Hijacking Attack

9. <https://www.langchain.com/stateofaiagents>

10. For a detailed dive into indirect prompt injections, see the white paper *"Not what you've signed up for: Compromising Real-World LLM-Integrated Applications with Indirect Prompt Injection."*

Bundling Is *Not* Integration

Surveys from leading industry analysts indicate that Chief Information Security Officers (CISOs) have mixed feelings about their organizations' AI security measures. Although most (83%) C-suite leaders who are investing in cybersecurity say their organization is investing the right amount, many (60%) remain worried that the cybersecurity threats their organization face are more advanced than their defenses.¹¹

Their concerns are well-founded. AI systems bring entirely new risks—like prompt injection, data poisoning, model theft, and hallucination—that traditional security tools were never built to handle. In response, vendors rushed to fill the gaps with point solutions focused on specific AI-related threats. While well-intentioned, this approach has led to a fragmented ecosystem of disconnected tools that do not share threat intelligence, are not integrated with one another, and require separate management. As a result, enterprises are forced to stitch together multiple products just to keep pace, while AI threats continue to evolve rapidly. The reality is clear: securing GenAI demands more than a bundle of isolated tools. It requires an integrated, AI-native approach that can adapt as fast as the technology it protects.

Palo Alto Networks has encapsulated decades of expertise and experience into a comprehensive platform for AI security, as the next section explains.

11. https://www.ey.com/en_us/ciso/cybersecurity-study-c-suite-disconnect




Figure 7: Prisma AIRS Provides Comprehensive Integrated Coverage Across the Entire AI Deployment

Prisma AIRS Platform Secures AI Deployments

Palo Alto Networks responded to this chaotic and convoluted approach by developing a complete platform for AI security. Prisma AIRS is the world's most comprehensive AI security platform, offering protection for models, data, applications, and agents. This unified approach not only addresses today's security needs, it can grow with the technology, protecting investments in AI security.

The key innovation in Prisma AIRS is to secure each component of the AI infrastructure with purpose-built security integrated into a single platform. This approach delivers threat protection with the highest efficacy and lowest false positive rates.

The security components include model scanning, AI agent security, posture management, AI red teaming, and runtime security. **These five capabilities together provide a powerful tool set for complete AI security.**

Each one is explored in more detail in the following five sections.




Figure 8: Prisma AIRS Provides Delivers Targeted Security for Each Component of Your AI Environment

Next up: AI Model Scanning

AI Model Scanning



AI models face several threats that can undermine their security. Model tampering involves altering a model's internal logic or parameters to produce biased or unsafe outputs.

Malicious scripts may be introduced during deployment, enabling unauthorized actions or system compromise. Deserialization attacks exploit how models load saved data, allowing attackers to run harmful code. Model poisoning occurs when false or manipulated data is added to the training set, causing the model to learn incorrect behaviors or hidden backdoors.

Prisma AIRS combats these threats with AI Model Scanning, which helps detect hidden threats like malicious code, backdoors, or insecure configurations before an AI model is deployed. This capability ensures the model is safe, trustworthy, and compliant with security policies.

With AI Model Scanning, you can:



Ensure models from open source and managed services are safe and secure



Prevent malware from entering your environments



Stop the execution of malicious code stored in the AI model

Next up: Posture Management

Posture Management



Posture management is essential for AI security because it provides ongoing visibility into how AI systems are configured and used. Without it, teams may overlook misconfigurations, unsafe behavior, or unauthorized access. Since AI systems evolve and handle sensitive data, posture management helps enforce policies, detect risks, and reduce the chance of breaches, ensuring safer, compliant AI operations.

Getting permissions right is critical for AI agents, which often act autonomously and access tools or data without direct oversight. Overly permissive policies can lead to security violations, data leaks, or system damage, while overly restrictive ones can limit agent effectiveness. To reduce the risk of breaches and ensure secure, compliant AI operations, agents must follow the principle of least privilege—having only the minimum access necessary to perform their tasks, and nothing more.

Prisma AIRS gives your security teams the capabilities they need for effective posture management. Now your team can have continuous visibility into the configuration, usage, and risks of AI systems. Armed with this information, organizations can detect vulnerabilities early, enforce security policies, and reduce the chances of misconfigurations or data exposure. With Prisma AIRS Posture Management, you can:



Continuously monitor and
remediate your security
posture



Prevent excessive
permissions, sensitive data
exposure, platform and
access misconfigurations,
and more



Ensure secure and
compliant AI agent and
application use

Next up: AI Red Teaming

AI Red Teaming



Red teaming is important for AI security because it helps organizations find weaknesses before attackers do. By simulating real-world attacks, red teams test how AI systems respond to threats like prompt injection, data poisoning, and model manipulation. This proactive approach uncovers hidden vulnerabilities in models, training data, and system behavior. It also helps improve defenses, validate policies, and strengthen trust in AI applications.

Red teaming plays a critical role in AI security by identifying weaknesses before attackers can exploit them. It simulates real-world threats—like prompt injection, data poisoning, and model manipulation—to reveal hidden vulnerabilities in models, training data, and system behavior. Unlike static red teaming tools that rely on predefined test cases, our solution is dynamic. It understands the context of the application—for example, healthcare or financial services—and intelligently targets the types of data an attacker would try to extract. Unlike other solutions, our adaptive testing engine doesn't stop at failure but rather learns, re-strategizes, and continuously retests until it identifies exploitable paths. This dynamic, context-aware approach not only uncovers deeper risks but also strengthens defenses and builds lasting trust in AI systems.

Prisma AIRS AI Red Teaming empowers your AI security team to:



Test your complete AI ecosystem across applications, agentic systems and LLMs using predefined and custom attack types



Create comprehensive reports detailing which attacks were successful and what sensitive information was extracted



Provide real-time recommendations to improve AI security posture

Next up: Agent Security

AI Agent Security



Securing AI agents is important because these systems can make decisions and take actions without human oversight. If compromised, they could misuse tools, access sensitive data, or cause serious harm. Threats like prompt injection, data poisoning, or over-permissioning can lead to unauthorized behavior. Securing AI agents ensures they operate safely, follow intended goals, and don't expose organizations to hidden risks. As agentic AI adoption grows, strong security controls are critical to prevent misuse and protect trust.

Prisma AIRS AI Agent Security gives your AI security team powerful tools so that they can:



Detect and block harmful
or toxic content in
prompts and responses



Create custom topic
guardrails to define
topics that your apps and
agents should or should
not discuss



Prevent hallucinating by
detecting outputs that
deviate from application
knowledge source

Next up: Runtime Security

Runtime Security

Runtime Security from Palo Alto Networks is a comprehensive solution designed to protect AI applications, models, data, and agents from both AI-specific and traditional cyber threats. Runtime Security offers real-time protection against risks such as prompt injection, malicious code, data leakage, and model tampering. By continuously monitoring AI systems, Runtime Security ensures the integrity and security of AI operations, helping organizations deploy AI technologies with confidence.

For organizations looking to secure their AI deployments, Runtime Security offers a robust and scalable solution that addresses the unique challenges posed by AI technologies. You can deploy Runtime Security at two levels, developer and network.

Network Level	Developer Level
Secure your AI applications with network-level enforcement: • AI and foundational network security in a single virtual firewall • Centralized platform for easier security management • Changes the network without changing the code	Secure your AI applications with code-level enforcement: • Embed rules, controls and policies directly into code (AI security-as-code as in DevOps) • Simple setup in minutes • Changes the code without changing the network

The Prisma AIRS platform integrates seamlessly with the Palo Alto Networks Strata Cloud Manager, providing centralized management and visibility across the AI ecosystem. Prisma AIRS employs advanced threat detection and prevention mechanisms to safeguard AI workloads, ensuring compliance and reducing the risk of data breaches.

Developers are under pressure to secure their GenAI assets, and now they have a powerful set of tools for doing so—Prisma AIRS. See what we mean below.

Prisma AIRS In the GenAI Lifecycle

Organizations can accrue the benefits of Prisma AIRS by integrating these capabilities into the GenAI lifecycle. Prisma AIRS addresses the complete lifecycle from pre-deployment integration and runtime controls to continuous monitoring and red-teaming for testing agent and model security.

Pre-Deployment Integration

Developers integrate Prisma AIRS into CI/CD or MLOps pipelines by scanning models and training data for backdoors, unsafe serialization, and embedded threats before deployment. Using APIs, Prisma AIRS also connects to model registries like MLflow or Hugging Face Spaces to automatically scan and tag approved models, streamlining early-stage security checks.

Runtime Controls

At runtime, developers use Prisma AIRS through APIs, Software Development Kits (SDKs), Model Context Protocol (MCP) or network config files to enforce strict access controls on GenAI agents, defining which tools or APIs each agent can use. These policies are enforced using sidecars or proxies to prevent unauthorized behavior. Prisma AIRS also enables prompt sanitization, input validation, output logging, and defense against prompt injection.




Figure 9: Prisma AIRS in the GenAI Lifecycle

Continuous Monitoring

Prisma AIRS enables continuous monitoring of AI environments by delivering real-time visibility into models, agents, and data activity. It detects abnormal behavior, misconfigurations, and security policy violations as they happen. By monitoring for threats like prompt injection, data leakage, and misuse of AI tools, it helps protect both development and production environments. The platform continuously analyzes activity to uncover emerging risks and adapts to evolving threats through automated detection and testing. This proactive approach ensures AI systems remain secure, compliant, and resilient—without relying on manual oversight or disconnected tools.

Red Teaming For Model and Agent Testing

Developers use Prisma AIRS red teaming tools to simulate adversarial inputs and misuse scenarios, testing how models and GenAI agents respond under potential attack conditions. These simulated attacks help identify vulnerabilities in logic, behavior, or tool access. Developers can use these insights to strengthen the model's defenses and improve agent safety, ensuring a more secure and reliable system before deployment.

When Palo Alto Networks needed security for its own internal agent, there was only one place to turn—Prisma AIRS—as the next section will explain.

Securing Strata Copilot with Prisma AIRS

Strata Copilot is a Palo Alto Networks AI assistant that uses Precision AI™ to simplify network security operations with real-time insights and natural language interaction.

The Prisma AIRS development team at Palo Alto Networks partnered with the Strata Copilot team for the first production deployment of Prisma AIRS. Strata Copilot helped shape the product roadmap by actively using the platform and providing early feedback. Today, every U.S.-based interaction with Strata Copilot runs through the Prisma AIRS API, which scans prompts and model responses for threats such as prompt injection, sensitive data exposure, malicious URLs, and toxic content. This integration gives the team real-time threat detection, visibility, and enforcement, allowing them to build a secure and compliant chatbot. Prisma AIRS also helps them ship features faster by aligning with Secure AI by Design principles.

The collaboration with Strata Copilot played a key role in developing Prisma AIRS into a flexible, production-ready solution. Insights from both Strata and external customers helped refine the product to meet the fast-moving needs of AI-powered apps, models, and agents. Their engineering team considers Prisma AIRS essential to the development lifecycle, allowing fast deployment, simplified security through API intercepts, and safer AI experiences.

Take the Next Step Toward Secure GenAI

This e-book has taken you through the current state of GenAI, the risks associated with GenAI applications, and the Prisma AIRS platform for AI security. As our world goes more and more in the direction of artificial intelligence, managing the risks and protecting against threats requires the old-fashioned kind of intelligence—what's between the ears. AI application security may be a relatively new notion to enterprise security teams, but threat actors are already using GenAI for their own purposes. The concepts and suggestions in this e-book can help you bridge the knowledge gap and start making informed decisions today about investing in AI security via the Prisma AIRS platform.

To learn more about Prisma AIRS, contact us for a [demo](#).



**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.

prisma_eb_secure_ai_environment_07_12_25