paloalto NETWORKS | PRISMA CLOUD

Defending Your AI Future with Prisma Cloud

How to Protect Your Critical Cloud AI Workloads and Data Without Slowing the Pace of Innovation

Table of **Contents**

PART 01

AI Innovation Is Taking Place in the Cloud

The integration of AI into hybrid and cloud environments is rapidly gaining momentum as companies across all industries recognize AI's potential in driving digital transformation and future growth. To handle the colossal volumes of data and execute energy-intensive AI workloads, organizations are turning to cloud services, increasingly adopting multicloud and hybrid cloud deployments to support and scale their AI initiatives. Virtual machines, cloud container services like Kubernetes®, and extensive cloud datastores form the backbone of large language models (LLMs), retrieval augmented generation (RAG), and AI agents.

The allure of the cloud lies in its provision of a scalable, adaptable, and cost-efficient foundation for AI workloads with the required compute power. This allows organizations to seamlessly scale up or down without incurring exorbitant physical infrastructure costs.

There's Power in the Cloud

Gartner estimates that by 2025, 95% of new digital workloads will be deployed on cloud-native platforms—up from 31% in 2021.[1] Forrester predicts that from 2023 to 2030, the compound annual growth rate of the off-the-shelf AI software market will be 36%, reflecting the growing adoption of AI technologies as part of strategic growth initiatives.[2]

95%

of new digital workloads will be deployed on cloud-native platforms—up from 31% in 2021.

Cloud-powered AI enables organizations to:

Train AI models.

Implement ChatGPT capabilities.

Automate workflows.

Analyze massive datasets.

Experiment with innovations.

Run large-scale AI deployments.

1. "Gartner Says Cloud Will Be the Centerpiece of New Digital Experiences," Gartner, November 10, 2021.
2. Michael O'Grady and Mike Gualtieri, *Global AI Software Forecast, 2023 To 2030*, Forrester, September 4, 2023.

Multifaceted Risks of Multicloud

To run the most demanding AI/ML workloads, organizations rely on multiple cloud service providers to develop, deploy, and consume AI products. Virtual machines, containers, and APIs are essential for AI innovation, and dedicated cloud environments for testing, staging, and production are essential to manage workloads, compute resources, and the separation of sensitive assets.

Multicloud infrastructures may be necessary for AI projects. However, the complexities of managing and securing each cloud environment become more difficult in our era of advanced threats. The growing attack surfaces created by multicloud deployments introduce visibility gaps, policy inconsistencies, and disparate access controls, which in turn open the door to bad actors.

One of the key hacker targets within the infrastructure of cloud-centric organizations are the sensitive and proprietary datasets that form the DNA of their AI workloads.

$300 billion
AI tech spend

expected in 2026[3]

3. *Worldwide AI and Generative AI Spending Guide*, IDC, 2023.

AI Datastores Are Essential—and at Risk

We're all familiar with the risks of a data breach. Losing employee or customer information, trade secrets, intellectual property, source code, operational information, business forecasts, and other sensitive data opens a company up to tremendous financial, reputational, and legal risks.

We can all agree that data breaches are bad for business.

Breaches have increasingly severe repercussions. In 2023, the global average cost of a data breach was US$4.5 million, a 15% increase since 2020.[4]

But in the data-reliant world of AI, data breaches can be catastrophic, existential events.

AI feeds on data, and it's hungry. To train LLM and other AI models, organizations must create massive datastores of their most sensitive information. They leverage terabytes of data, usually hosted on public cloud services that contain:

- Intellectual property
- Business intelligence
- Supplier data
- Personal identification data
- Health information

4. *Cost of a Data Breach Report 2023*, IBM, July 2023.

Types of Attacks on AI Data

The type of data used in AI/ML datastores has been referred to as "radioactive gold." It's extremely valuable, but it must be handled, stored, and controlled very carefully because in the wrong hands it can be dangerous.

In addition to conventional security threats like SQL injection, AI datastores are subject to constantly evolving, unique attacks that target them specifically. Here are just five threats to look out for:

1 Data poisoning

Attackers can modify the core data/code on which an AI model is trained, ensuring corrupted results before training even takes place. This can lead to wildly biased data outputs and security breaches.

2 Prompt injection

Just like it sounds, threat attackers inject prompts that substantially change the intended behavior of the model. For instance, malicious prompts could cause your AI chatbot to expose sensitive data from the training dataset.

3 AI DDoS attacks

Much like the classic denial-of-service attacks, AI DDoS threats overwhelm LLMs to deny access to data or apps, which forces your system to consume excess resources just to keep up.

4 Membership inference

In this attack, a target machine's learning model is observed with the goal of being able to create their own "shadow model." The shadow model is then used to predict what data records were part or not part of the target model's training data.

5 Model inversion

Similar to membership inference, this attack targets specific data records such as PII, but in this case attackers are attempting to predict the input data based on the outputs of your AI models.

The Challenge of Securing AI Data

The cloud delivers the scalable reservoir of data storage and processing power needed to run today's and tomorrow's AI workloads, but running AI workloads across different cloud environments can pose challenges.

The pace of innovation is stretching many organizations beyond their cloud security comfort zone, endangering the very AI projects that will drive growth.

Mission-critical, cloud-stored data is at risk. Attackers can exfiltrate and sell your proprietary data or leverage it for ransomware. Bad actors and malicious insiders can inject tainted and malicious data into your datastore, compromising your AI outputs.

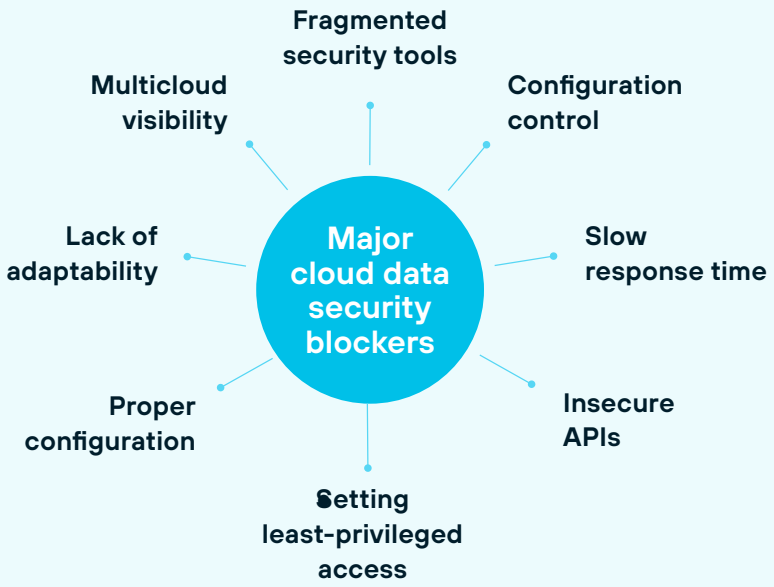

More Tools, More Problems

So, how do you protect your vital AI tools and datasets in this rapidly changing technology landscape? The answer isn't more disconnected security tools and siloed solutions. The more tools you have, the less secure your AI datastores and cloud workloads really are.

More tools mean:

- More maintenance, more training, and more processes.
- More work with fewer results and an increased number of alerts.
- More effort to implement, maintain, and integrate with other solutions in your cloud security stack—and that means investigation and remediation slows to a crawl.
- More risk in visibility overload and the inability to correlate attack paths.

In fact, organizations with 16+ security tools experienced 2.8X more data security incidents in 2023 than orgs with fewer tools.[5]

The answer is an integrated, comprehensive CNAPP that makes cloud security less complex—a solution that leverages AI/ML to secure your cloud-based AI workloads without slowing down the pace of innovation.

The answer is Prisma Cloud.

Organizations with
16+ security tools experienced
2.8X more data security incidents
in 2023 than orgs with fewer tools.

5. *Cost of a Data Breach Report 2023*, IBM, July 2023.

PART 02

How Prisma Cloud Helps Defend Your AI Future

Reducing security risks in the cloud is a challenge in the best of circumstances. Today, though, given the rapid acceleration of AI-assisted development—not to mention AI-driven attacks—it can feel almost impossible. There's simply too much to do and not enough resources to manage it all.

To that end, Prisma® Cloud takes a multipronged approach to help customers address this new attack vector and securely adopt AI internally to deliver better business outcomes.

1. **Countering AI with AI**: Prisma Cloud utilizes AI to graph model all the possible pathways an attacker could take from an initial vulnerable asset so you can easily understand how extensively a risk could spread.
2. **Securing AI by design**: Prisma Cloud secures your usage of AI to help your organization protect itself against the unique risks associated with AI, ML, and GenAI models—including data exposure, misuse, and model vulnerabilities.
3. **Security simplicity with AI**: Prisma Cloud enables anyone to effectively solve security issues with a simple conversation. Copilot understands natural language and not only surfaces the most pressing risks but also recommends the best remediation workflow, eliminating guesswork.

Prisma Cloud Harnesses the Power of Precision AI™

Precision AI is high-fidelity automation built on a rich security dataset. Scale and automate cyber defense—detect, prevent, and resolve alerts in real time.

Machine Learning

Machine learning to help our security applications to become more accurate at preventing, predicting, and remediating security problems by using precisely defined historical and current data.

Deep Learning

Continuously building predictive models to anticipate issues before they can happen.

Generative AI

Controls that speak "human" so you don't have to speak machine.

Precision AI

Combining AI in all its forms to predict and block AI attacks even as they escalate.

Countering AI with AI

The sophistication and scale of AI-powered attacks overwhelm current tools. What's needed is automated intelligence that can continuously analyze your cloud stack, precisely connect insights from across your environment, and rapidly identify the greatest risks to your organization.

Prisma Cloud combines its existing ML-powered threat detection capabilities (UEBA, network anomaly detection) with new AI risk modeling and blast radius analysis. With AI-driven analytics, see the impact of compromise along with the most efficient remediation workflow recommendations. Always take the most efficient remediation steps with action plans.

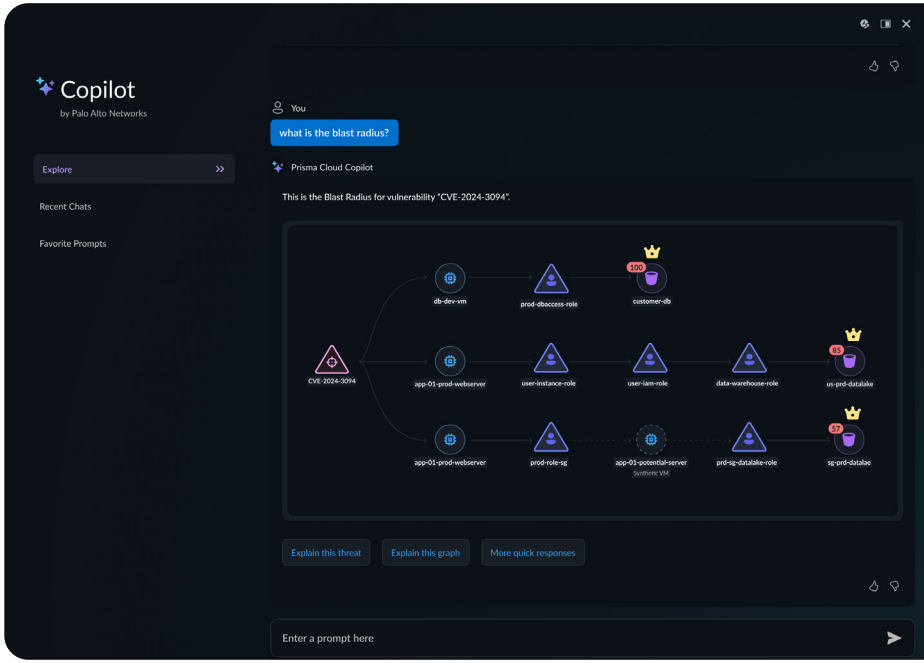

Figure 1: Prisma Cloud Copilot helps you quickly understand the broader risk impact to your crown jewels by just having a conversation

Learn more about how to secure with AI →

Securing AI by Design

To address the unique challenges of deploying AI and GenAI at scale while helping reduce security and compliance risks, Prisma Cloud AI-SPM delivers a new set of capabilities that addresses those challenges. It provides visibility into the AI model lifecycle, from data ingestion and training to deployment. By analyzing model behavior, data flows, and system interactions, AI-SPM helps identify potential security and compliance risks that may not be apparent through traditional risk analysis and detection tools. Organizations can use these insights to enforce policies and best practices, ensuring that AI systems are deployed in a secure and compliant manner.

Additionally, Prisma Cloud AI-SPM monitors for AI-specific threats like data poisoning, model theft, and proper output handling, alerting security teams to potential incidents and providing guidance on remediation steps. As regulations around AI continue to evolve, AI-SPM can also help organizations stay ahead of compliance requirements by embedding privacy and acceptable use considerations into the AI development process.

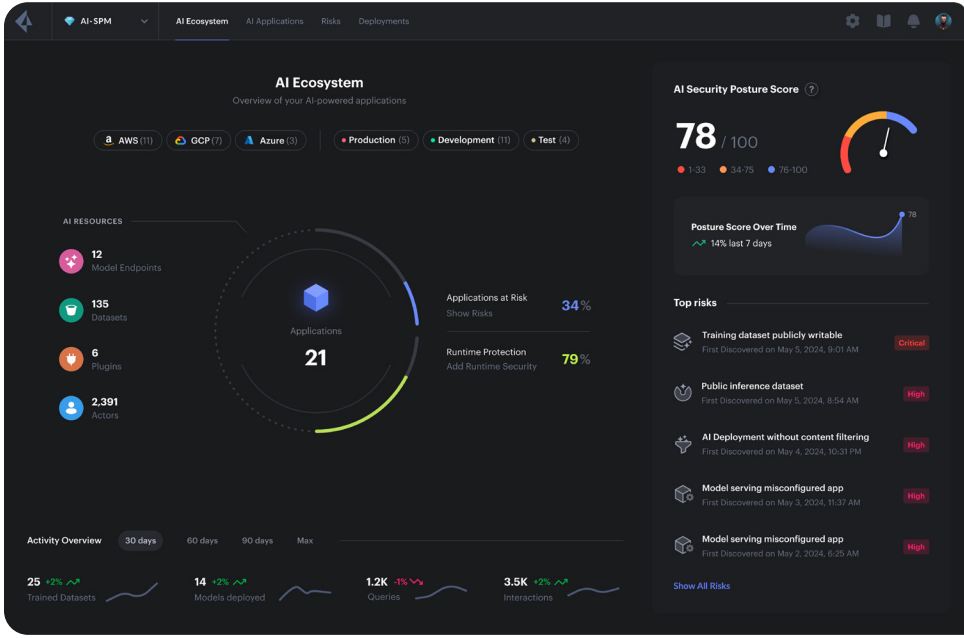

Figure 2: AI-SPM dashboard highlights the ecosystem, resources, risks, and the security posture score

Check out the AI-SPM datasheet →

Security Simplicity with AI

Democratize cloud security and empower your teams to effectively reduce cloud risk by simply having a conversation. Prisma Cloud Copilot eliminates the uncertainty of what to focus on and how best to fix it. It details the code-to-cloud impact and provides additional context, such as why a risk is critical and how an attacker could exploit it. Copilot also helps you find what you're looking for—without wasting time sifting through documentation or navigating multiple dashboards.

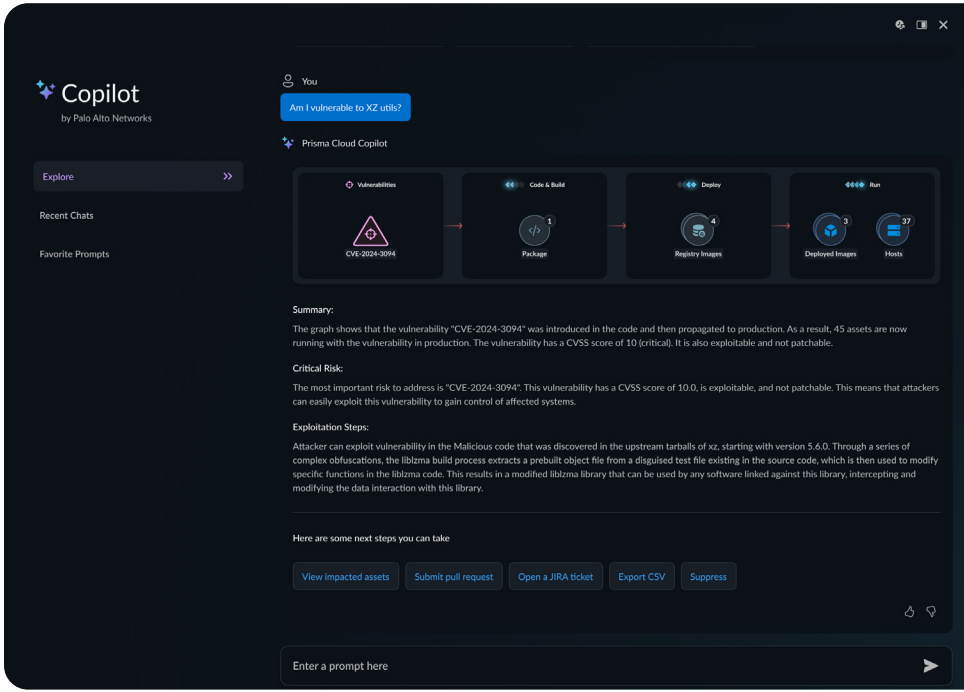

Figure 3: Prisma Cloud Copilot shows the code to cloud impact, details the risk, and offers relevant next steps

Watch the Prisma Cloud Copilot demo →

Cloud Security Health Check

How safe are your most sensitive cloud workloads?

Our Cloud Security Health Check is a complimentary service that helps identify gaps and vulnerabilities in your cloud security posture.

The service includes:

- 10-day free trial of Prisma Cloud.
- Onboarding guidance customized for your tech stack.
- A detailed security risk assessment report.
- A report walkthrough with a Prisma Cloud solutions architect.

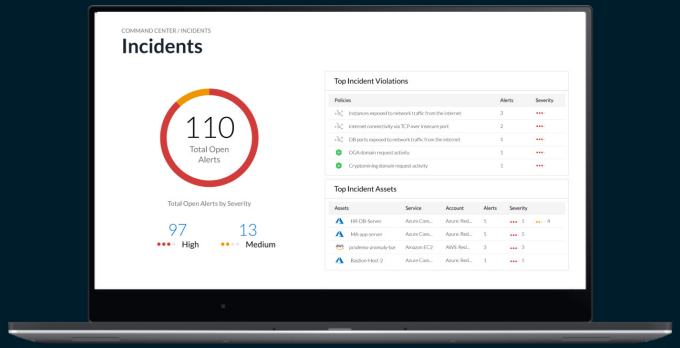

Approx. **$50K value**

You'll get a customized report that includes your top policy violations and misconfigurations, your most exposed assets, and your most critical alerts.

Get your free Cloud Security Health Check.

PRISMA® CLOUD
BY PALO ALTO NETWORKS

About Prisma Cloud

Prisma® Cloud is the industry's most comprehensive cloud-native application protection platform (CNAPP) with the broadest security and compliance coverage—for applications, data, and the entire cloud-native technology stack—throughout the development lifecycle and across multicloud and hybrid environments. Our integrated approach enables security operations and DevOps teams to stay agile, collaborate effectively, and accelerate secure cloud-native application development.

To learn more, visit us online.

paloalto® NETWORKS

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

parent_eb_defending-your-ai-future_062424