Datasheet

Prisma Cloud AI Security Posture Management (AI-SPM)

The Challenge: Ensuring Security and Compliance Amid the AI Gold Rush

Artificial intelligence (AI) and large language models (LLMs) are part of the strategic focus of modern organizations. At the same time, many new AI and LLM models and tools – including managed model APIs, open source foundation models, and an emerging ecosystem of operational tooling, plug-ins, and applications – have become available in recent years.

However, while organizations enjoy unprecedented opportunities to build and deploy AI-powered applications rapidly, there's no shortage of new accompanying security challenges.

Development far outpacing security: In the rush to realize AI's value, security considerations can fall by the wayside. The pressure to deploy AI quickly is immense, and can lead to inadequate security reviews and hasty implementations.

Black box systems: The inner workings of large AI models are often opaque, even to their creators, making it difficult to anticipate potential security and compliance issues. Models may exhibit unexpected behaviors or vulnerabilities that are not easily detectable through traditional testing methods.

New attack vectors challenging existing approaches: Existing security measures such as firewalls and posture analysis tools (e.g., CSPM, DSPM) do not address AI-specific attacks like data poisoning, model inversion, and adversarial attempts.

Evolving compliance risk: The EU AI Act, which is expected to come into force later this year, imposes new requirements around data privacy, algorithmic bias, and explainable AI. It also raises the stakes for non-compliance, with penalties nearly double those of GDPR. Similar legislation is expected in the US and elsewhere.

The Solution: Visibility, Control, and Governance with AI-SPM

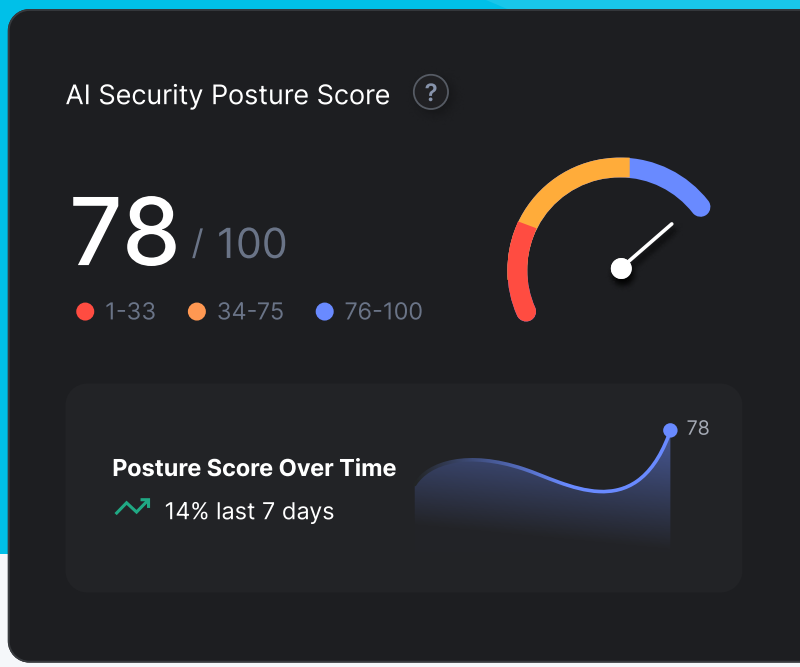

Prisma Cloud AI-SPM is a set of capabilities designed to protect organizations against the unique risks associated with AI, machine learning (ML), and Generative AI (GenAI) models, including data exposure, misuse, and model vulnerabilities. As part of our broader Code-to-Cloud™ approach, we have seamlessly integrated AI-SPM capabilities with the Prisma Cloud security platform, while building on existing data security posture management (DSPM), cloud security posture management (CSPM), and cloud-native application protection (CNAPP) capabilities.

With Prisma Cloud AI-SPM, organizations gain visibility into the AI model life cycle – from data ingestion and training to deployment. By analyzing model behavior, data flows, and system interactions, Prisma Cloud identifies potential security and compliance risks that may not be detected through traditional risk analysis and detection tools.

Organizations can use AI-SPM insights to:

- Enforce policies and best practices ensuring that AI systems are deployed in a secure and compliant manner.
- Uncover AI-specific threats like data poisoning, model inversion, and adversarial attacks.
- Stay ahead of evolving compliance requirements by embedding privacy and acceptable use considerations into the AI development life cycle.

Prisma Cloud AI-SPM integrates with AI Runtime Security, empowering security teams to protect their AI application ecosystem at run time effectively.

Features & Benefits

AI Model Discovery and Inventory

The proliferation of managed, semi-managed, and unmanaged AI models can make oversight difficult. Security teams need visibility into AI deployment to monitor usage effectively and prevent downstream risk.

Control model sprawl and shadow AI: See an inventory of model APIs, open source models, and models deployed on virtual machines.

Prevent model misuse: Identify who is using which model to prevent unsanctioned model use and unauthorized use cases.

Improve governance: Receive alerts for new model deployments to ensure that appropriate controls are in place.

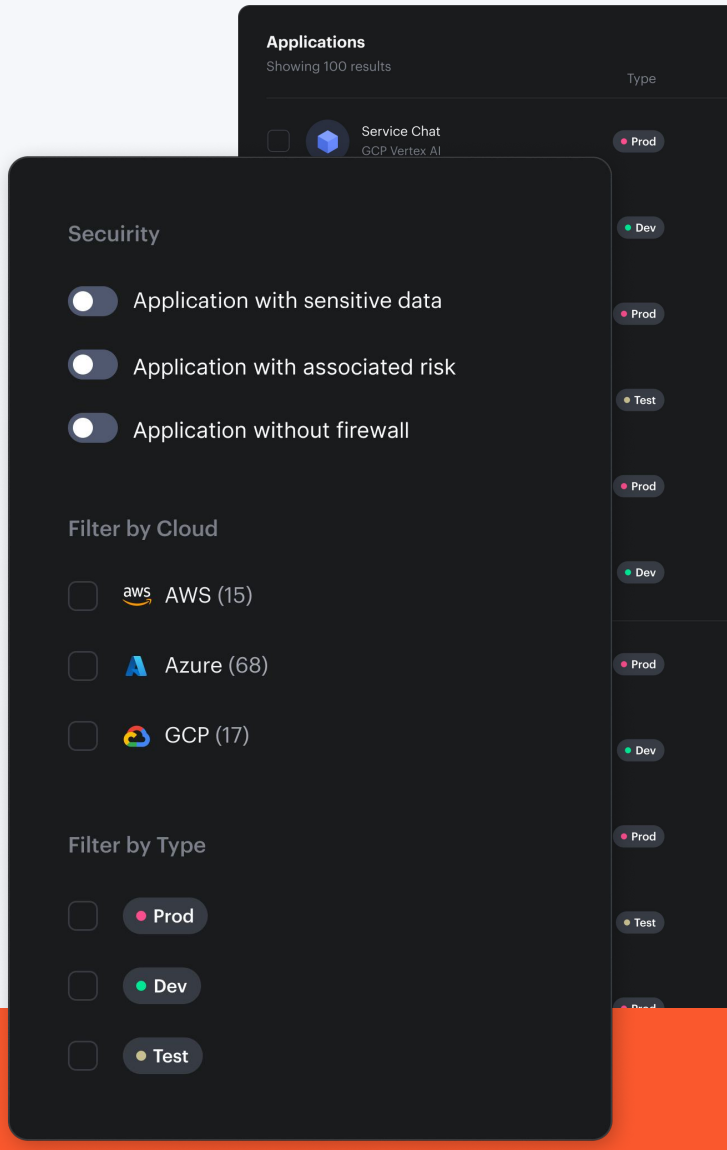

Data Exposure Prevention

AI models are trained on vast amounts of data that may contain sensitive or regulated data such as personally identifiable information (PII) or trade secrets. In addition, they can be exposed inadvertently or via adversarial attacks. Prisma Cloud AI-SPM helps you understand what internal data is accessible through each deployed model.

Discover and classify training datasets: Prevent data poisoning and find out if models are being trained or fine-tuned on sensitive data – before they are deployed.

Carry out retrieval-augmented generation (RAG) and inference data monitoring: See which datasets and data flows are used for retrieval, and understand how they impact effective access to data.

Analyze model interactions: Scan prompt and output logs for evidence of model misuse or data exposure.

Posture and Risk Analysis

Misconfigurations or weak access controls in data pipelines, training environments, and deployment infrastructure can introduce significant security and compliance risks. Prisma Cloud AI-SPM scans your end-to-end AI deployment to find weaknesses and prioritize the most critical remediations.

Prioritize and address misconfigurations: Analyze the full attack path and find vulnerabilities across AI applications, data, and pipelines – building on contextual insights delivered by Prisma Cloud's DSPM and CSPM capabilities.

Model access governance: Get a visual mapping of who has access to deployed AI models and associated resources such as compute, data, and applications.

Rightsize permissions to applications and data: Combat overgenerous permissions to internal models or unwanted public access to models.

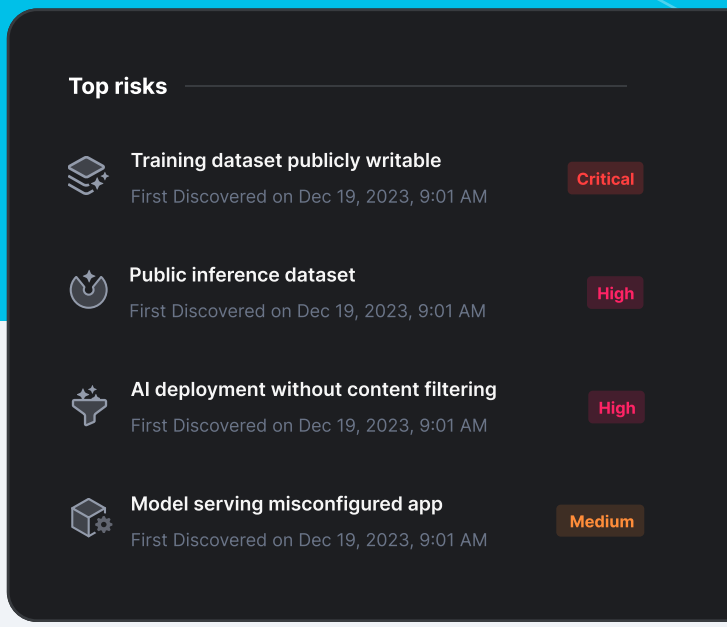

Risk Scenario Examples

Prisma Cloud AI-SPM not only gives you a laundry list of issues to fix. It also enables you to see the full context of your data and applications, providing your busy security teams with the tools they need to prioritize issues for effective remediation.

Below is a partial list of risks that Prisma Cloud AI-SPM can detect, alongside additional context made available based on agentless scanning of your cloud environment.

Risk	Added context for prioritization
AI deployment enabling anonymous access	Which model is being used in the deployment and whether its training or inference datasets contain sensitive data
Publicly writable training dataset	Which models are trained on the dataset and who can access them
Misconfigured asset containing prompt history	Severity of misconfiguration and scope of prompt history (public/internal)
Lack of content filtering for production AI deployment	Who has access to the deployment and on which internal data was the model trained
AI deployment accessible from out of production	Potential compliance implications and their impact on environments
Cross-region AI pipeline	Which sensitive data can be exposed and which relevant compliance policies are being violated

Quick and Easy Agentless Deployment

Installs in minutes and does not require agents or connectors to access your cloud data.

Data is scanned out of band to minimize the impact on production.

Sensitive data never leaves your cloud environment.

About Prisma Cloud

Prisma Cloud is the industry's most comprehensive cloud-native application protection platform (CNAPP). It provides 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 at prismacloud.io.

paloalto NETWORKS | PRISMA BY PALO ALTO NETWORKS