Solution Brief

Securing AI in AWS with Netskope

Netskope and AWS deliver the industry's most robust AI security fabric. By embedding real-time governance directly into the AWS environment, we transform security from a bottleneck to a competitive advantage, enabling secure, autonomous, agentic workflows at scale.

At a Glance

- **Safeguard AI and agentic workflows:** Protect sensitive data, training sets, and autonomous agents from leaks, prompt injection, and malicious exfiltration.
- **Amazon Bedrock and shadow AI governance:** Accelerate experimentation with visibility into unmanaged AI tools while ensuring Bedrock usage remains secure.
- **Unified data protection for AI:** Secure the entire AI data lifecycle across S3, data lakes, and model inference with real-time, inline controls.
- **Unified zero trust for AI:** Centralize security data across AWS and distributed environments to automate detection and response against AI-driven threats.

"Use AWS to build out your AI - and agentic AI - at scale. Run it safely with Netskope."

The challenge: AI outpaces governance

As organizations accelerate AI adoption on AWS, governance is struggling to keep pace. Traditional security models are ill-equipped to address modern AI risks, including the inability to monitor dynamic, machine-to-machine interactions in real time. Organizations also face significant hurdles such as uncontrolled shadow AI sprawl, the risk of data leakage, and complex compliance requirements.

Unmanaged LLMs and unsanctioned tools expose sensitive information, while prompt injection, jailbreaking, and excessive agent permissions introduce critical new security gaps. Autonomous AI systems further amplify these risks, turning unauthorized access into high-impact breaches. Without continuous, inline enforcement, securing these complex AI environments becomes difficult, stalling innovation and creating operational overhead.

The solution

Netskope embeds real-time security into AWS AI interactions, empowering teams to scale securely. AWS AgentCore Gateway and Netskope One AI Guardrails provide a layered defense: AgentCore handles deterministic, identity-aware authorization, while Netskope inspects prompts and responses to block semantic threats like prompt injection and data exfiltration. By combining the Netskope Zero Trust Engine with AWS infrastructure, organizations transform security into a competitive advantage, ensuring data, prompts, and autonomous actions remain protected as AI ecosystems evolve.

Securing autonomous agentic workflows

As organizations move from AI assistants to autonomous agentic workflows, the security challenges shift from human-driven interactions to machine-to-machine (M2M) activity. AI agents on AWS access data and execute actions independently, creating a new security frontier where traditional tools fail. Dedicated controls are required.

Netskope One distinguishes between Agent Identity, verifying the agent's legitimacy, and Agentic Permission: defining what the agent is authorized to do. This distinction is vital for minimizing the impact of unauthorized autonomous actions. To achieve this, the joint solution pairs deterministic policy enforcement with runtime security. AWS AgentCore Gateway acts as the deterministic policy engine, using Cedar policies to authorize every agent-to-tool request based on identity, tool, and parameter context. This ensures that only authorized actions are executed, while Netskope One AI Guardrails provides the semantic gate, inspecting the content and intent of every prompt and response to block threats like indirect injection or unauthorized data exfiltration.

The Netskope One Agentic Broker provides visibility for Model Context Protocol (MCP) communications, preventing unauthorized data access during tool execution. Automated actions are continuously monitored using context-aware controls and least-privilege principles.

Organizations gain capability to instantly terminate unauthorized autonomous actions and maintain comprehensive, policy-driven audit logs, providing the operational assurance required to trust autonomous systems. By combining real-time visibility and inline enforcement, Netskope One empowers organizations to scale agentic automation securely, protecting sensitive data and agents from leaks, prompt injection, and malicious exfiltration.

Move fast and stay protected. We provide the AI fast lane to ensure your teams innovate safely with continuous governance.

Build safe generative AI on Amazon Bedrock

Organizations accelerate generative AI applications on AWS using Amazon Bedrock and SageMaker. It's critical that the focus is on facilitating secure experimentation, instead of blocking activity wholesale. With Netskope's instance awareness, organizations gain clear, auditable visibility into every model interaction.

Netskope AI Guardrails inspects prompts and responses in real time, safely guiding usage. Using Layer 8 intelligence (deep contextual awareness), the platform identifies and blocks malicious interactions before they impact AWS environments.

Bedrock Guardrails only apply to models invoked through Bedrock. Netskope's guardrails apply to any AI traffic that crosses the network — ChatGPT, Gemini, Copilot, private models, and thousands of other apps — regardless of what cloud or model is behind them.

Easy-to-implement policies guide employees toward approved AI services while providing visibility into shadow AI activity. By combining continuous governance and zero trust controls, teams innovate freely, ensuring every AWS deployment is resilient, compliant, and optimized for success.

Secure the entire data lifecycle. Enforce inline data protection that moves with your information everywhere.

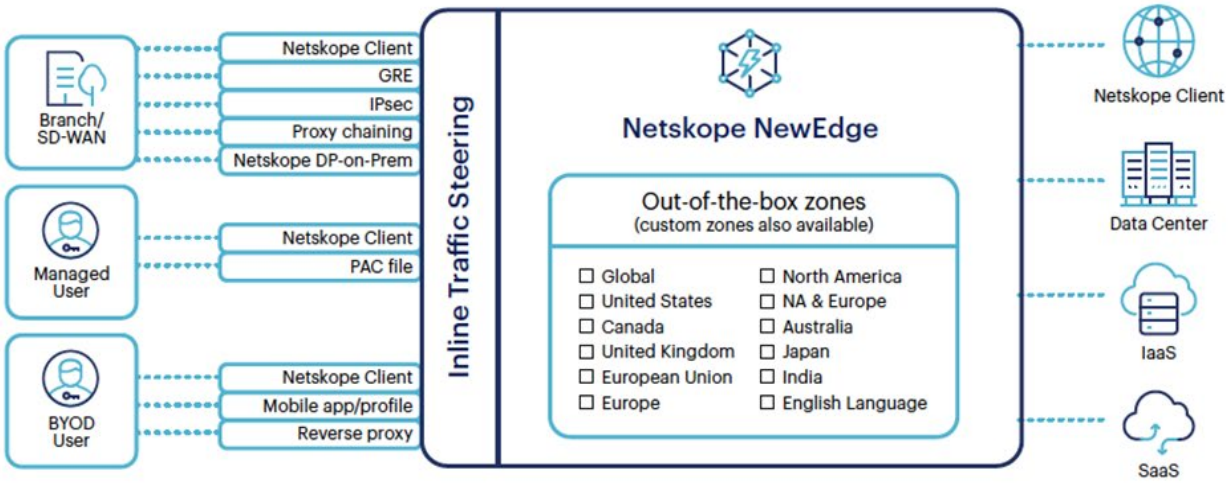

Extensivew steering of branch offices, managed devices, and unmanaged devices to Netskope's security fabric.

Comprehensive data protection for AI workloads

Data is the cornerstone of every AI initiative. As models ingest more and more information for training and inference, sensitive intellectual property faces new exposure risks. Traditional security approaches often fail to govern data dynamically as it flows between cloud services and AI models.

Netskope One provides unified data protection across the entire AI data lifecycle. Data is continuously discovered, classified, and governed through real-time, inline controls that follow data everywhere. Advanced data loss prevention (DLP) policies monitor activity across AWS environments, including Amazon S3 and data lakes, ensuring only approved data is used within AI workflows. This shifts organizations from reactive monitoring to proactive, continuous enforcement, preventing unauthorized exposure.

Netskope One AI Red Teaming identifies vulnerabilities within private models through automated adversarial testing prior to deployment, protecting critical data and models.

Emerging technologies like 5G, augmented reality, and edge computing are set to push boundaries even further, delivering unprecedented speed, responsiveness, and interactivity.

Unified zero trust for AI

Securing modern AI environments requires more than isolated security controls. As organizations scale AI across AWS, SaaS, and distributed applications, fragmented tools create visibility gaps and inconsistent enforcement.

Netskope One provides a unified control plane delivering real-time visibility and zero trust enforcement across users, applications, data, and autonomous AI workflows. Interactions are continuously evaluated based on identity, device posture, behavior, and context before access is granted. This approach applies consistent controls across cloud workloads, SaaS, and external traffic.

Human and non-human activity are monitored through a single operational framework, reducing blind spots. Netskope user and entity behavior analytics (UEBA) assesses risk and identifies anomalous activity within AI interactions, with inline enforcement containing risk immediately.

By combining AWS with Netskope's real-time zero trust architecture, organizations simplify operations and scale AI initiatives with full alignment to business objectives.

BENEFITS	DESCRIPTION
Enable generative AI securely	Confidently adopt AI with real-time inspection of prompts and responses. Prevent sensitive data exposure and mitigate critical risks, such as prompt injection and jailbreak attempts, directly within your AWS environment.
Protect training and inference data	Apply consistent data protection across S3, data lakes, and AI workflows. Reduce the risk of unauthorized access, misuse, or exfiltration during every stage of the model training and inference lifecycle.
Secure autonomous agentic workflows	Govern autonomous AI actions with visibility across MCP communications and AI-related API traffic. Ensure agents operate strictly within approved security boundaries to minimize the blast radius of autonomous actions
Automated AI model hardening	Identify vulnerabilities before production deployment. Utilize automated adversarial testing and red teaming to uncover security weaknesses in private AI models and custom workflows, ensuring resilience against sophisticated threats.
Govern shadow AI usage	Discover unmanaged AI applications and services across AWS. Improve visibility into the environment and guide users toward approved, secure enterprise AI platforms to maintain compliance and reduce shadow IT.
Consistent zero trust access	Continuously verify human and non-human identities, device posture, and context. Enforce secure, granular access policies across all AWS-hosted AI applications, data, and services through a single unified control plane.

About Amazon

Amazon Web Services (AWS) is the world's most comprehensive and broadly adopted cloud, offering more than 200 full-featured services from data centers globally.

AWS's suite of cloud computing products enables developers to create and host applications in a secure, scalable way. That's why millions of customers, including everyone from the fastest-growing startups, to the largest enterprises and leading government agencies who use AWS to lower costs, become more agile, and innovate faster.