netskope

The AI Security Checklist

A practical guide for CISOs and security teams securing AI usage

The AI Security Checklist

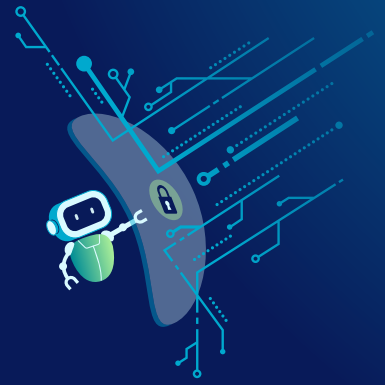

Enterprises are rapidly adopting AI, putting security teams under pressure to make sure they have the right foundations in place for deploying it safely. But in the race to harness AI's benefits, it's vital that organizations don't miss a step.

This checklist provides the most important questions you need to ask as you examine your environment, policies, and controls. It's designed to help security leaders quickly assess current status and identify what to do next—making it easier for you to spot gaps, prioritize actions, and feel more confident about where to focus activity.

Discover

Identify all AI usage across the organization

There's been an explosion of AI tools within enterprises recently, not all of which is evident to security professionals. Shadow AI usage remains a persistent problem: 47% of AI users still use personal AI apps at work (Netskope, 2026 Cloud and Threat Report). There's also the possibility that your existing portfolio of SaaS applications might add AI features that increase data risk. Security teams need to scrutinize every tool for its AI impact.

Consider

Do we have visibility into both managed and unmanaged AI tools?

Can we identify AI features embedded in SaaS applications?

Do we understand how users and agents are interacting with AI models today?

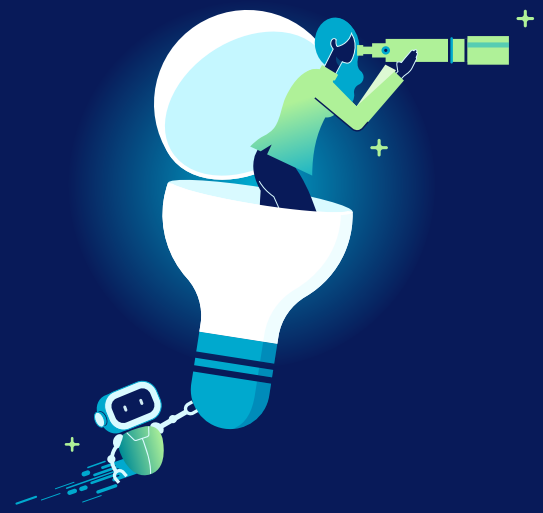

Analyze

Understand your AI-related data exposure and risk patterns

Not all data is valued equally. That principle has long guided enterprise security policies. But it's never been more important than in the AI era, when security teams are trying to strike the right balance between speeding up innovation and controlling risk. A more granular understanding of how AI tools are being used, and what type of data is flowing through them, can help organizations make smarter decisions and focus their resources more effectively.

Consider

Do we know what types of data are being shared with AI tools?

Can we distinguish between low-risk and high-risk AI interactions?

Are we able to spot patterns that indicate emerging risk?

Enforce

Apply safe usage controls across all AI interactions

At the heart of the zero trust approach to security is the understanding that access needs are contextual. Granting appropriate access becomes more complicated with AI, when tools might inadvertently serve confidential information to a colleague who queries a corporate AI tool. Adding to this complexity is the ongoing deployment of AI agents, which introduce new attack vectors, including unsafe autonomous actions and expanded pathways for data exfiltration.

Consider

Can we control AI usage and data access based on user/agent identity, data sensitivity, or risk?

Are controls consistent across different AI tools and environments?

Are our private AI models protected from jailbreaking or malicious prompt engineering attempts?

Govern

Ensure AI usage is compliant, traceable, and responsibly managed

Security practitioners know that compliance with industry standards is essential. But AI regulations are regularly evolving to keep pace with the technology, so it's important to adapt security policies in alignment and maintain clear audit trails at all times. This is a particular imperative in sectors like financial services and healthcare where regulation is business-critical.

Consider

Can we demonstrate how AI usage is governed today?

Do we have auditability and traceability across AI interactions?

Are responsibilities for AI usage clearly defined?

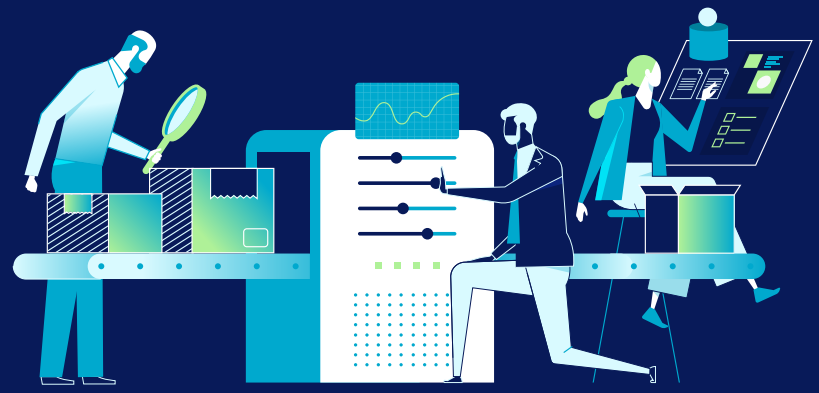

Ready to find out more? Learn more about Netskope's approach to securing AI here.

About Netskope

Netskope, a leader in modern security and networking for the cloud and AI era, addresses the needs of both security and networking teams by providing optimized access and real-time, context-based security for people, devices, and data anywhere they go. Thousands of customers, including more than 30% of the Fortune 100, trust the Netskope One platform, its Zero Trust Engine, and its powerful NewEdge network to reduce risk and gain full visibility and control over cloud, AI, SaaS, web, and private applications—providing security and accelerating performance without trade-offs.

Interested in learning more?

Request a demo