CLOUDFLARE

Accelerate AI adoption with security by design

AI is here, and traditional security is being left behind

98%
of organizations have unsanctioned apps, including shadow AI. These blind spots risk data exposure and compliance violations.[1]

50%
success rate of prompt injection attacks, the top-ranked security threat on the OWASP Top 10 for LLM apps.[2]

97%
of organizations experiencing AI-related security incidents lacked proper AI access controls.[3]

93%
of employees admit to putting info into AI tools without approval.[4]

Secure the AI lifecycle with Cloudflare

across generative and agentic AI

Human-to-AI security

Agentic AI security

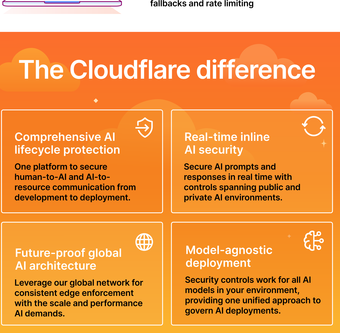

Extend visibility
across AI apps, API endpoints, and AI agent and MCP connections

Manage risk
with guardrails, posture control, and real-time threat mitigation

Protect data
with prompt protections and access controls for employees and AI agents

Security is built-in
when developing AI on Cloudflare

Secure workforce use of GenAI

- Discover shadow AI
- Evaluate risks of AI apps
- Manage AI app posture
- Restrict sensitive data inputs
- Enforce guardrails in prompts and responses

Protect AI-enabled apps and workloads

- Discover shadow AI endpoints
- Protect AI models and training data from abuse
- Prevent PII from being exposed in prompts and responses
- Enforce content moderation

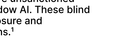

Secure what you build

- Build MCP servers with authentication / authorization built-in
- Log and restrict AI model requests / responses, token use, and costs
- Prevent service disruptions with model fallbacks and rate limiting

The Cloudflare difference

Comprehensive AI lifecycle protection
One platform to secure human-to-AI and AI-to-resource communication from development to deployment.

Real-time inline AI security
Secure AI prompts and responses in real time with controls spanning public and private AI environments.

Future-proof global AI architecture
Leverage our global network for consistent edge enforcement with the scale and performance AI demands.

Model-agnostic deployment
Security controls work for all AI models in your environment, providing one unified approach to govern AI deployments.

Customer case studies

indeed
World's #1 job website

Identifies and controls shadow AI
in parallel with remote access modernization project

newfold digital
Web hosting provider

Discover, label, and protect
AI experiments and production apps across the enterprise for a unified AI-native safeguard layer

Cyndx
AI-driven fintech company

Reduced inference costs by 95%
by adopting Cloudflare to cache and run responses from AI model providers

Discover how Cloudflare can secure your AI adoption

Explore use cases

1. Varonis, 2025 State of Data Security
2. OWASP Top 10 for LLM Apps, v1.1, Vu, Y., & Liu, Y. (2024). Automatic and Universal Prompt Injection Attacks against Large Language Models.
3. IBM, 2025 Cost of a Data Breach
4. ManageEngine, The Shadow AI Surge in Enterprises

1 888 99 FLARE | enterprise@cloudflare.com | Cloudflare.com

REV:BDES-6274.20251013

CLOUDFLARE