

Cloudflare AI Security Suite

Secure AI interactions by controlling data and managing risk across your AI lifecycle.

Confidently scale AI

AI risks demand modern security

Your teams are using AI to innovate faster, but this creates critical security risks. The old playbook of blocking AI or adding complex point solutions is failing because it stifles innovation and ignores reality:

- **85% of employees** use AI tools before IT can vet them.¹
- **93% admit** to putting company data into AI without approval.²
- **63% of breached** organizations have no AI governance policies.¹

Learn to empower your teams — and reap the productivity benefits of AI — while maintaining the security and control your business requires.



A unified platform to secure agentic and GenAI

Cloudflare provides a single platform for your organization to embrace AI with confidence. We empower security leaders to manage risk, technology teams to unlock productivity, and platform engineers to build securely, ensuring that everyone can innovate together safely.

Get started with Cloudflare AI Security Suite

Cloudflare AI Security Suite is delivered on a unified platform for securing workspace use of AI tools and public-facing applications. Discover shadow AI, protect models from abuse, secure agent access, and prevent data exposure in prompts — so your enterprise can innovate safely and efficiently, with easier visibility and stronger control.



Extend visibility

across AI apps, API endpoints, and AI agent connections.



Mitigate risk

with guardrails, posture control, and real-time threat mitigation.



Protect data

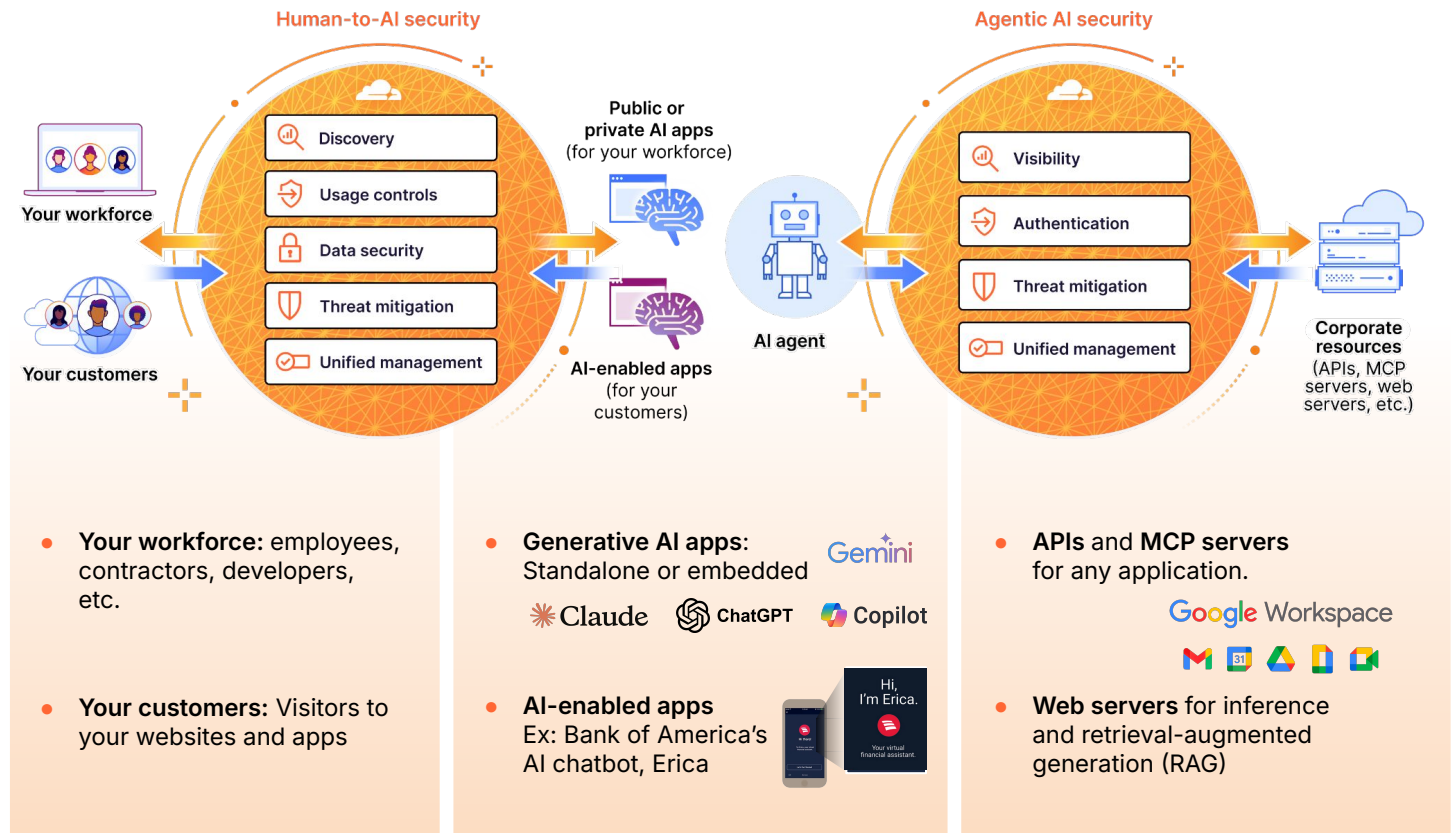
with prompt protections and usage controls for employees and AI agents.



Security is built-in
when developing AI on Cloudflare.

Secure generative and agentic AI communication with Cloudflare AI Security Suite

Protect all AI communications by controlling the data your workforce uses in generative AI and managing the security risks posed by autonomous agents.



Accelerate AI adoption with security by design across the AI lifecycle



Secure workforce AI use

Implement AI usage controls and AI security posture management (AI-SPM) to mitigate risk and protect data.



Protect AI-powered applications


Protect your AI apps and APIs from prompt injection and data leaks in real time.



Build AI securely

Empower developers to secure AI apps with integrated observability, rate limiting, and inline AI guardrails.

Secure workforce use of AI apps and workloads with Cloudflare's SASE platform



SSE to protect human-to-AI communication

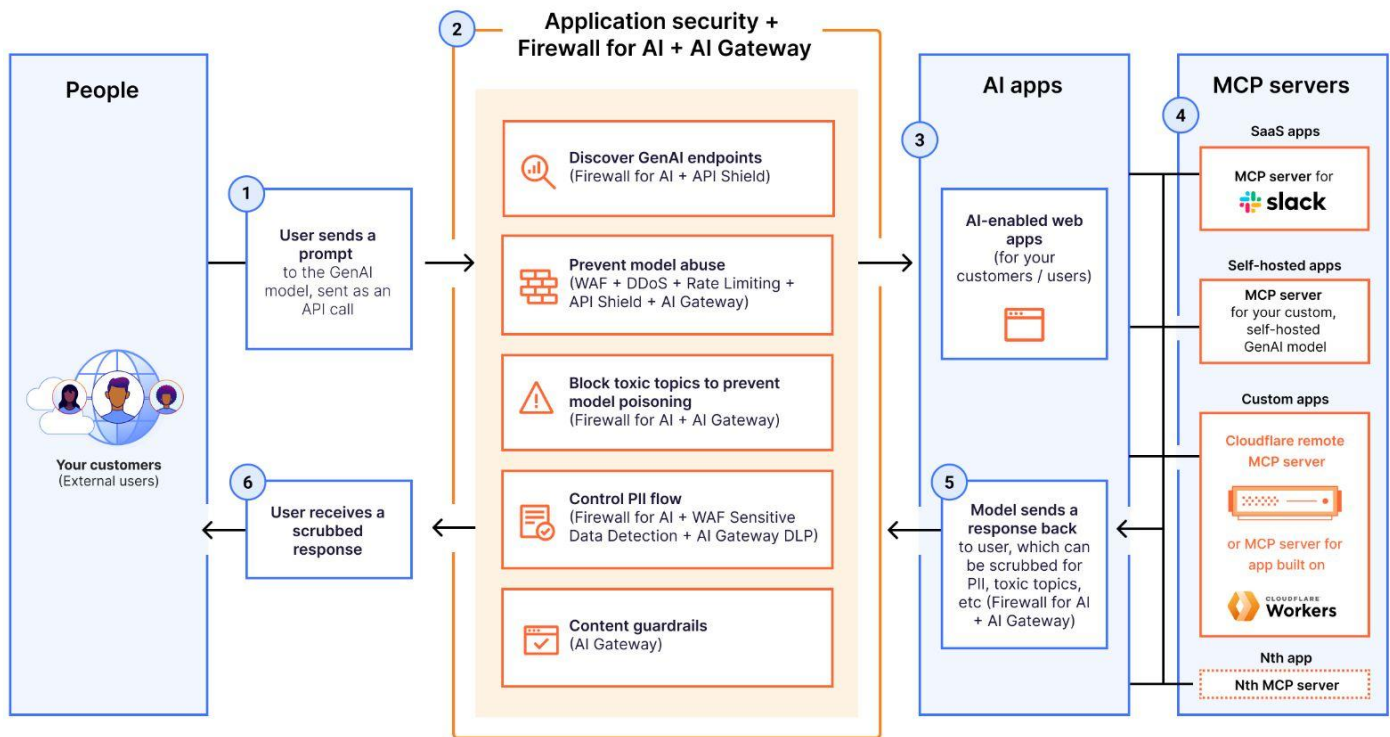
- **Visibility:** Discover and analyze [shadow AI](#) use via inline traffic inspection. Evaluate risks posed by those AI apps with [transparent scoring](#).
- **Access controls:** Block, isolate, redirect, or allow user connections. Enforce identity-based zero trust rules per app.
- **Prompt protection and guardrails:** Detect and block user prompts based on [intent](#) (e.g., jailbreak attempts, code abuse, PII requests).
- **Data security:** Stop sensitive data exposure with AI-powered [data loss prevention \(DLP\)](#) detections for PII, source code, and more.
- **AI security posture management:** Integrate with GenAI tools via API to scan for misconfigurations with our cloud access security broker (CASB). Available now for [ChatGPT](#), [Claude](#), and [Google Gemini](#).

MCP server portals to protect AI-to-resource communication

- **Visibility:** Aggregate all Model Context Protocol (MCP) request logs for audit and analysis. Review and approve each MCP server before adding to portal.
- **Authentication:** Authenticate user access to portal based on identity. Scope access to MCP servers based on least privilege.
- **Connections:** Connect all accessible MCP servers with a single URL, instead of individually configuring each MCP server.
- **Unified management:** Enforce the same granular access policies for your AI connections as you do for your human users.

Note: [MCP server portals](#) supports any MCP server, including (but not limited to) [remote MCP servers you build or deploy](#) on Cloudflare. This capability is available as a [zero trust network access \(ZTNA\)](#) control.

Protect AI-enabled apps and workloads with Cloudflare's model-agnostic inline security



Protect public-facing AI with application security and Firewall for AI

- **Discover GenAI endpoints:** Automatically discover all AI models and APIs across your web properties.
- **Protect AI models from abuse:** Use our purpose-built [Firewall for AI](#) to block prompt injection, model poisoning, excessive usage, and other threats that may bypass traditional security protections.
- **Control PII flow:** Scan user prompts and model responses to [block sensitive data](#) from being exposed, helping you maintain compliance.
- **Content guardrails:** [Block unsafe or toxic prompts](#) using integrated models like Llama Guard. Create custom WAF rules to easily block or log suspicious AI interactions.

Protect AI you build with Developer Platform and AI Gateway

- **Unified [AI control plane](#):** Manage all your AI apps from a single dashboard. Route requests, cache responses, control costs, and monitor performance.
- **Protect credentials at the edge:** Securely store API keys and [secrets](#) at the edge, preventing client-side exposure and simplifying key rotation across providers.
- **Enforce content safety guardrails:** Automatically identify and [block](#) / redact harmful content and PII in prompts and responses.

Cloudflare is the only vendor to secure both your public-facing and private AI environments

Implement the right guardrails to adopt AI with confidence, ensuring security speeds up your innovation, not hinders it.

- **Unified AI ecosystem protection:**
Complex security stacks increase risk. Use one platform to protect data and ensure compliance across the entire AI lifecycle.
- **Future-proof global architecture:**
Prevent tomorrow's challenges today with a post-quantum safe network that scales for any traffic volumes, constantly adapts to new threats, and is programmable for new use cases.
- **AI-powered security:**
Our AI-powered defenses inspect prompts and responses for real-time threat detection.
- **Proven AI leadership:**
Innovate with confidence on a platform trusted by 80% of the top 50 GenAI companies.
- **Model-agnostic deployment:**
Security controls work for all AI models in your environment, providing one unified approach to govern AI deployments.

What customers are saying



Identify & control shadow AI

World's #1 job website
[Read case study](#)

In parallel with VPN replacement project



Isolate public genAI tools like ChatGPT

Insurance technology
[Read case study](#)

to block copy-paste of sensitive data



AI-enabled SaaS company

Protect PII

by preventing customers from submitting sensitive information to public-facing GenAI endpoints



AI-driven fintech

95% reduced inference costs

by adopting Cloudflare to cache and run responses from AI model providers

Ready to discuss your AI security needs?

Talk to an expert

1. 2025 Manage Engine research: [Source](#)
2. 2025 IBM, Cost of a Data Breach report: [Source](#)