JOINT SOLUTION BRIEF

Cloudflare AI Security for Apps

Executive summary of the integration

Wiz and Cloudflare empower security teams with end-to-end protection for applications, AI endpoints, and cloud infrastructure. By connecting Cloudflare's AI Security for Apps into Wiz's Security Graph, security teams gain a single source of truth for discovering which AI application endpoints are actively secured by Cloudflare and which remain exposed. Together, Cloudflare's AI guardrails and Wiz's contextual analysis combine to map sensitive data flows and AI security rules, including those that block prompt injections, PII leaks, and unsafe content, to help security teams identify unprotected and misconfigured attack surfaces and validate their security posture.

Market challenge

As AI adoption accelerates, organizations face a rapidly expanding attack surface characterized by "shadow AI," exposed sensitive data, and unreviewed code shipped at 100x speed. Traditional security lacks the context to identify these new AI-specific risks - like prompt injections, PII leaks, and "rogue" agents - across complex cloud stacks. This integration bridges the gap by providing full visibility and real-time protection across the entire AI lifecycle.

Key benefits of the integration

- **Unified AI Attack Surface Posture and Visibility:** Organizations gain a single view of every application endpoint in their environment with Wiz and discover which are secured by Cloudflare's AI Security for Apps, and which are still exposed so they can close gaps.
- **Sensitive Data Mapping with Full Context:** Security teams can monitor sensitive data flows to prevent private information from reaching LLMs or leaking in responses. With these data flows mapped onto the Security Graph, Wiz provides full visibility into the AI application ecosystem, mapping models, agents, and accessible data sources, helping teams prioritize for remediation.
- **Proactive AI Guardrail Validation:** Wiz verifies that AI deployments are protected by Cloudflare's AI Security for Apps; if guardrails are misconfigured, Wiz alerts teams to allow for direct remediation within the Cloudflare platform.

The better together story

Wiz and Cloudflare unite to provide end-to-end protection for AI and cloud infrastructure with a single source of truth. Cloudflare's AI Security for Apps delivers real-time edge guardrails, such as prompt injection and unsafe topic prevention, to secure AI endpoints. Wiz adds deep context by leveraging its full capabilities, from code to cloud to runtime, to map the complete AI application and surface security gaps.

By integrating Cloudflare's security rules into the Wiz Security Graph, teams can prioritize risks based on exploitability and ensure that their AI Applications, including those with access to sensitive data, production systems and cloud creds, are protected. This automated visibility eliminates blind spots across managed agents and self-hosted models while reducing manual validation.

Together, Wiz and Cloudflare empower security teams to safely accelerate AI adoption and reduce the attack surface across modern cloud environments.

Use Case Overview

Use Case: Unified AI Endpoint Protection: Security teams must secure a rapidly expanding AI footprint—including LLMs, managed agents, and self-hosted models—against new threats like prompt injection, PII leaks, and shadow AI. Effective defense requires bridging the gap between infrastructure visibility and real-time edge protection to ensure every AI application is governed by consistent security guardrails.

Challenge: Visibility Gaps and Unprotected AI Surfaces: The speed of AI development leads to "shadow AI" and unprotected endpoints being deployed without proper security controls. AI applications are often granted access to sensitive data, internal APIs, and critical systems—yet traditional security tools lack the context to identify which AI services are publicly exposed, what data they can access, or which guardrails, such as prompt injection filters or PII blocking, are missing. This lack of visibility creates a fragmented security posture where critical AI vulnerabilities remain unprioritized and unmitigated.

Solution: Wiz and Cloudflare AI Security for Apps: Wiz and Cloudflare empower security teams by merging deep cloud visibility with active edge protection. Wiz maps Cloudflare AI Security for Apps directly onto the Security Graph, providing a single view to see exactly what is protected and where the gaps are. By identifying AI security rules, including those blocking prompt injections, PII leaks, unsafe and custom-defined topics, Wiz helps teams validate their security posture and identify unprotected attack surfaces.

With this unified context, teams can prioritize risks by exploitability and continuously verify that customer-facing bots have active toxicity filters. Together, Wiz and Cloudflare offer an automated solution to monitor sensitive data flows and ensure that AI applications remain behind a verified safety net.

Integration Architecture Diagram

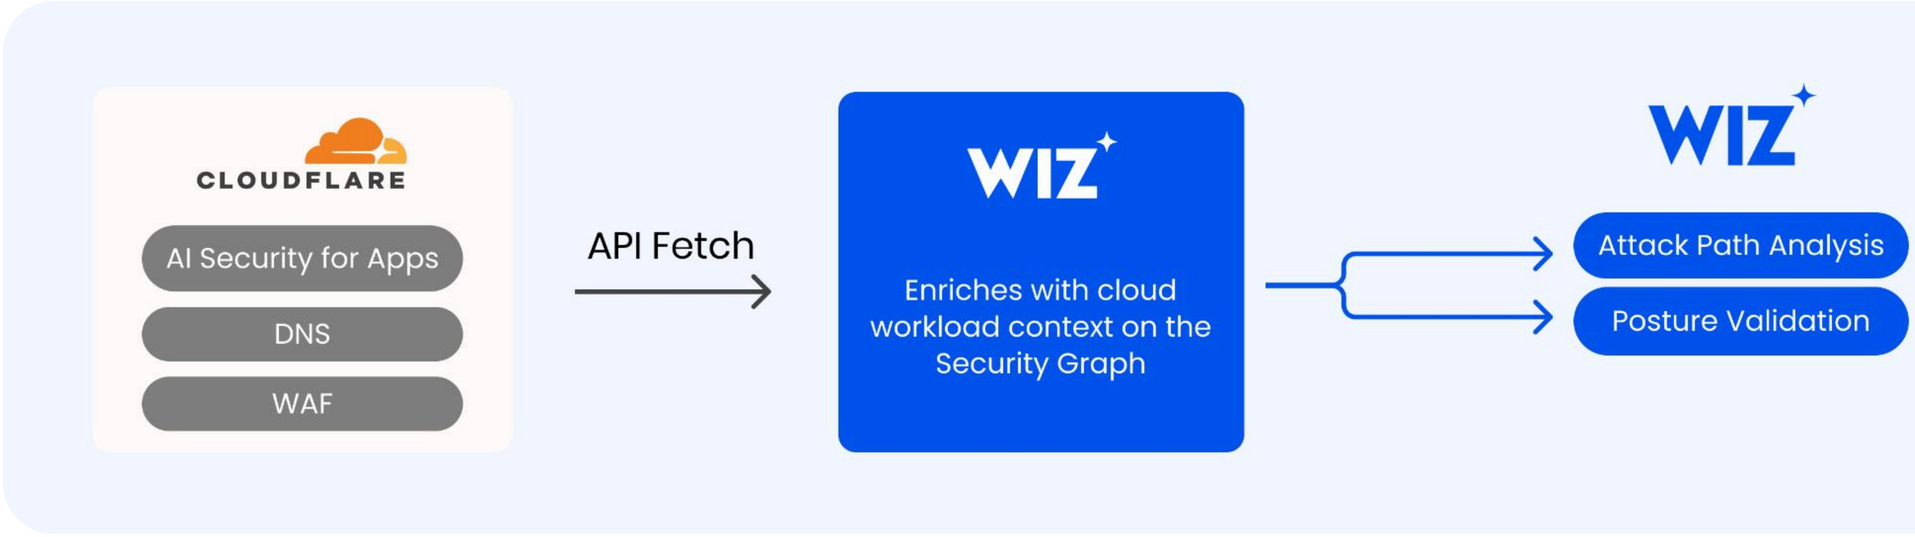

Integration Architecture Diagram

About Wiz

Wiz is on a mission to transform cloud security for customers – which include 50% of the Fortune 100 – by empowering them to embrace a new cloud operating model. Its Cloud Native Application Protection Platform (CNAPP) delivers full-stack visibility, accurate risk prioritization, and enhanced business agility. The result? More context with less noise, so that security teams can focus their time on what matters most.

About Cloudflare

Cloudflare, Inc. (NYSE: NET) is the leading connectivity cloud company. It empowers organizations to make their employees, applications, and networks faster and more secure everywhere, while reducing complexity and cost. Powered by one of the world's largest and most interconnected networks, Cloudflare blocks billions of threats online for its customers every day.