CLOUDFLARE

The Connectivity Cloud & AI Security

Dan Mitchell
Sr Solutions Engineer, US Southeast Majors
6/17/26

One network, one control plane on a global scale

330+ cities
in 125+ countries, including mainland China

w/190+ cities
for AI inference powered by GPUs

~20%
of web properties sit behind Cloudflare

190 billion
cyber threats blocked every day

500 Tbps
of network capacity (and growing)

Unmatched network resilience

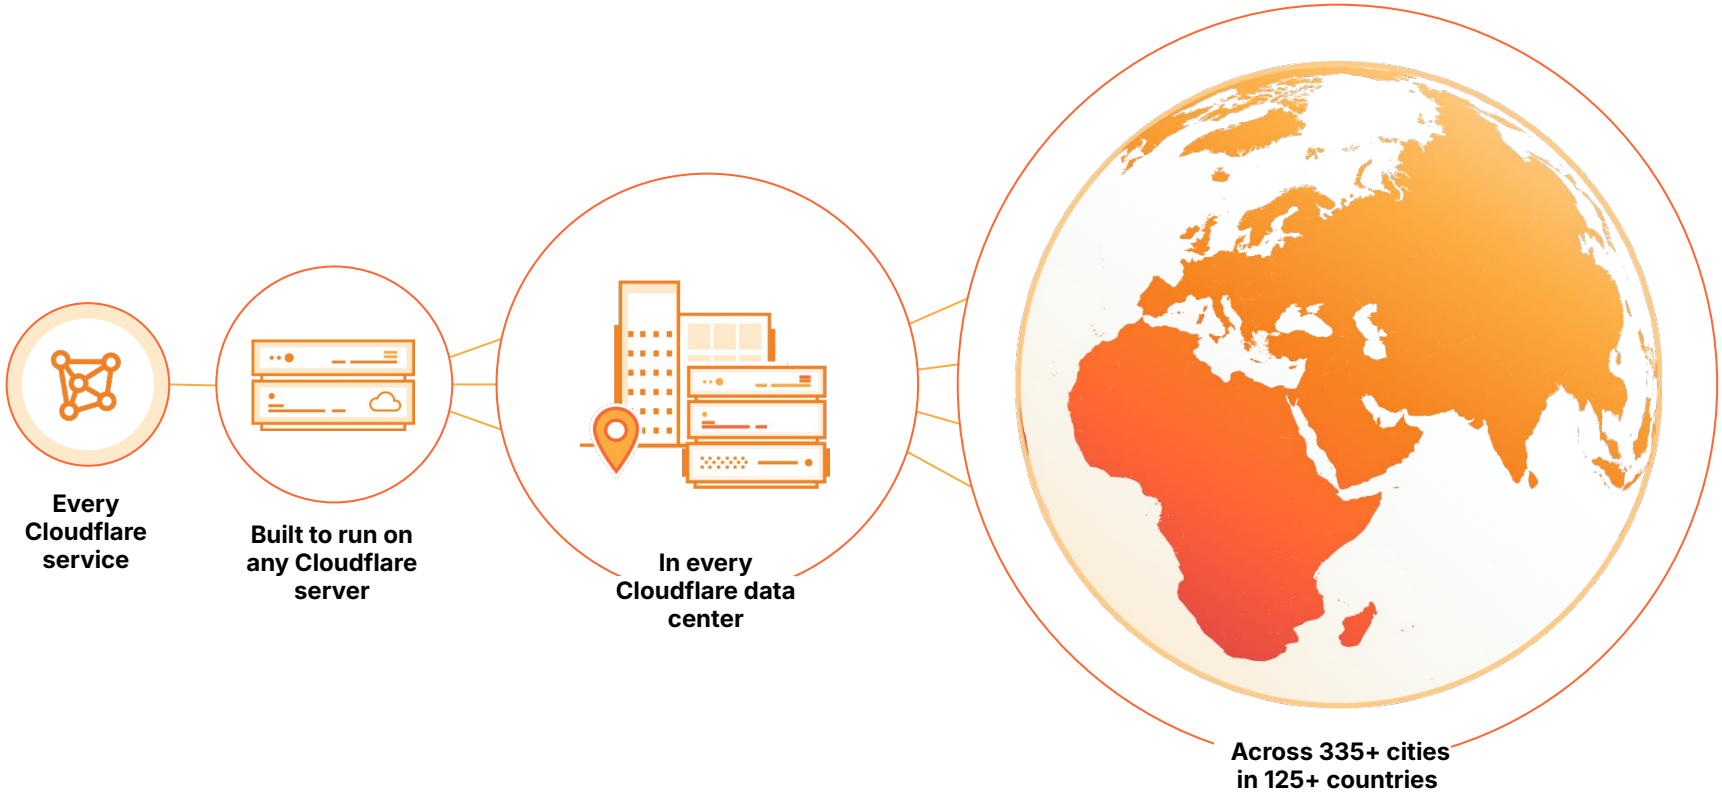

Cloudflare innovation is continuously accelerating

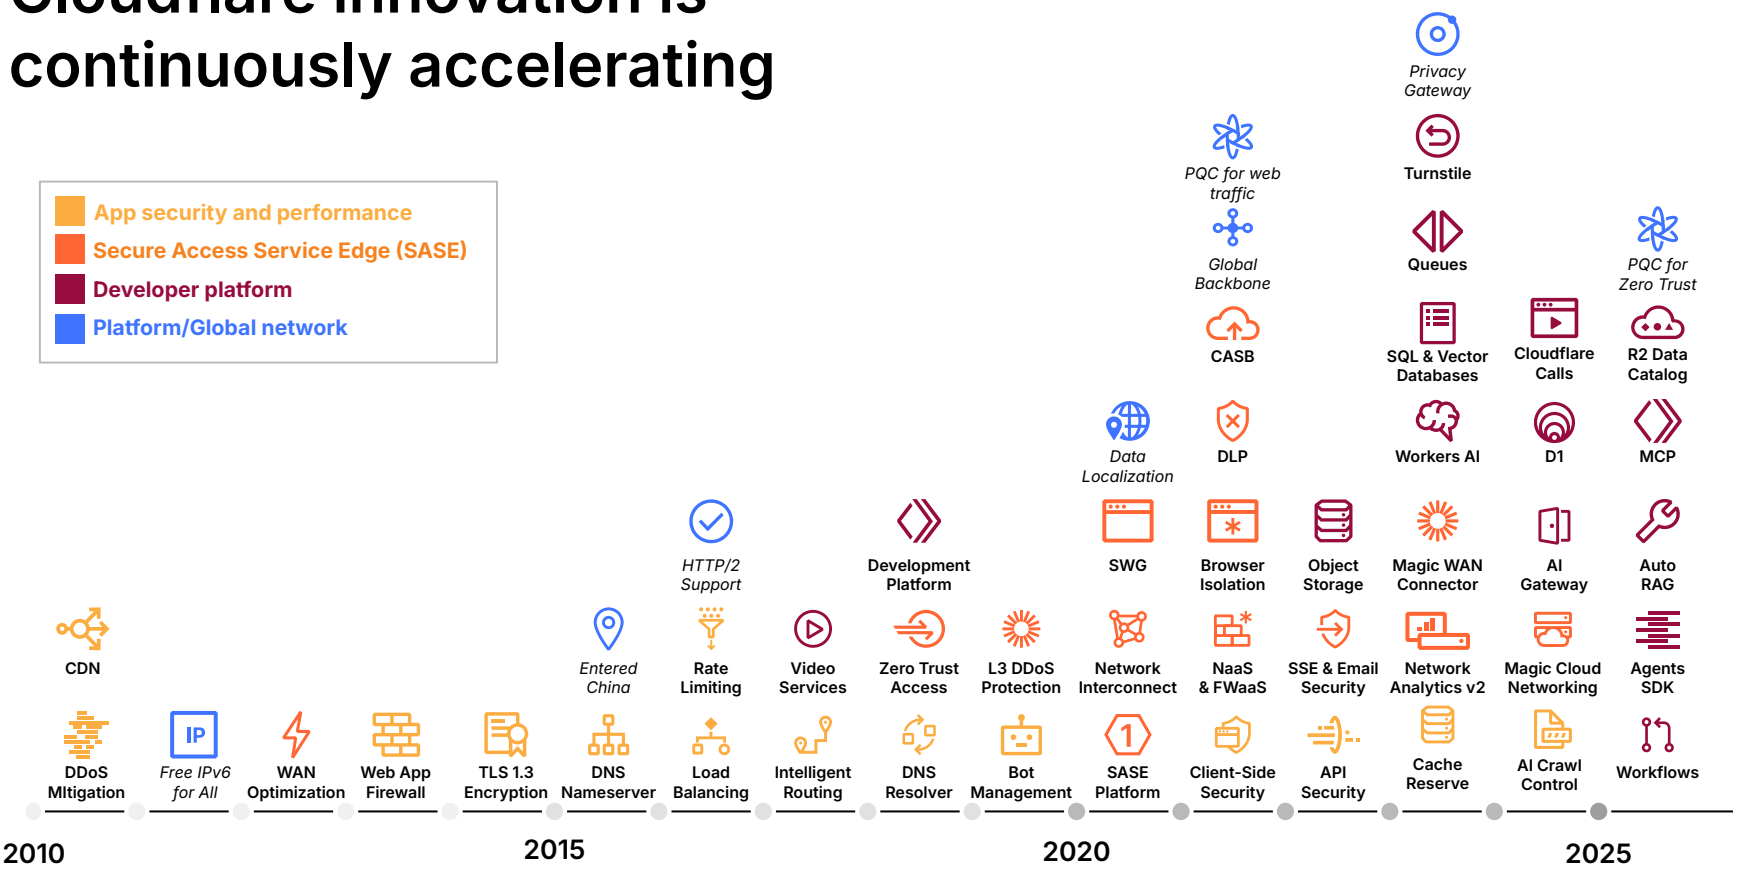

CLOUDFLARE

Integrated Global Cloud Platform

Cloudflare Platform

Cloudflare One

Cloudflare Zero Trust Services

- Zero Trust Network Access
- Secure Web Gateway
- Cloud Access Security Broker
- Cloud Email Security
- Remote Browser Isolation
- Data Loss Prevention

Cloudflare Network Services

- WAN-as-a-Service
- Firewall-as-a-Service
- L3 & L4 DDoS Protection
- Network Interconnect
- Smart Routing
- IDS/IPS

Cloudflare Application Services

- WAF with API Protection
- Rate Limiting
- Load Balancing
- Bot Management
- L7 DDoS Protection
- CDN and DNS

Cloudflare Developer Platform

Workers | Pages | R2 | Workers KV | Durable Objects | Images | Stream

Cloudflare Global Network

Compliance/Privacy: FedRAMP, ISO, SOC, PCI, GDPR compliant, Logs & Analytics, Data Localization Suite

CLOUDFLARE

AI Security

Cloudflare AI Security Architecture

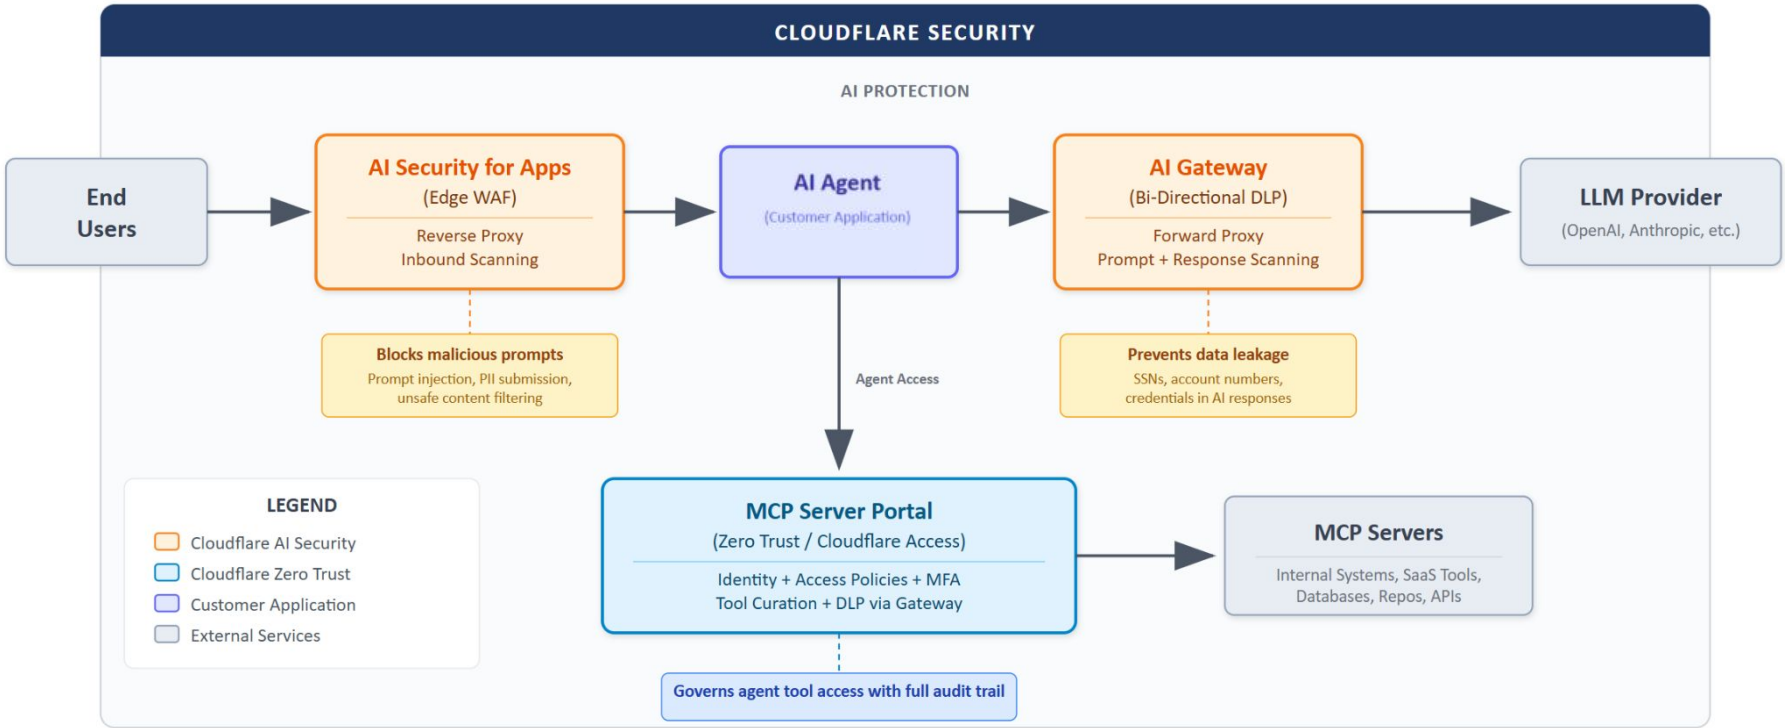

Securing workforce use of AI apps (human →AI) with Cloudflare's SASE platform

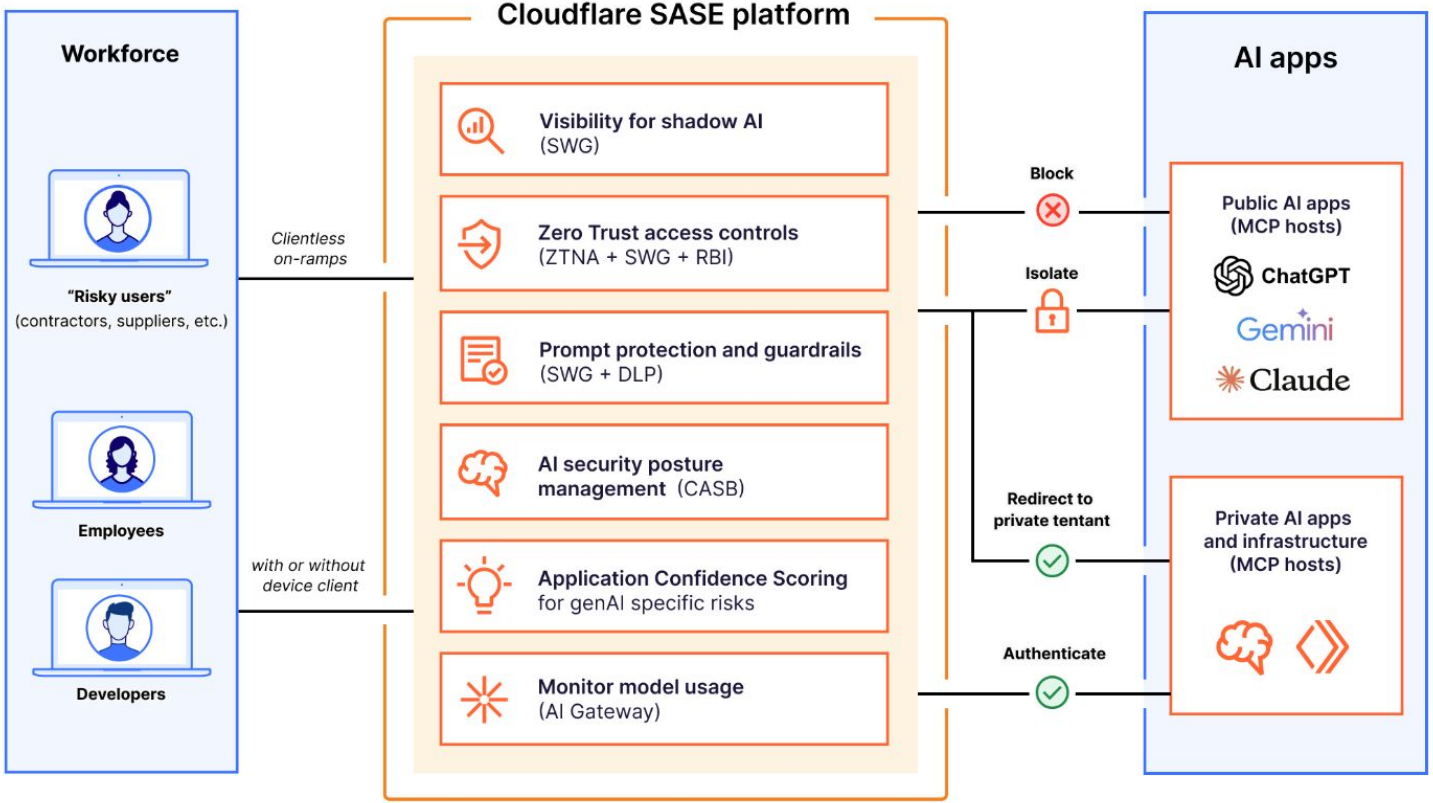

CLOUDFLARE

Frontier Models At Cloudflare

CLOUDFLARE

Cloudflare's focus on Agentic AI + Frontier Models

Codex Controls & Governance

Consistent control integrated into CI/CD pipeline (& agent deployment) leveraging our Codex

Agentic Security Automation

Streamline operations and automation of foundational security controls. Core Focus is on Validation and Response

Target Operating Model

Update our security operating model to a "Platform" model. Strategic Functions + Engineering + Ops.

Red Teaming & Vulnerability Scanning (Mythos)

Mythos integrated into our code scanning harness and red team workflows. Building model-agnostic Agents across Security.

Cloudflare Agentic Security Approach

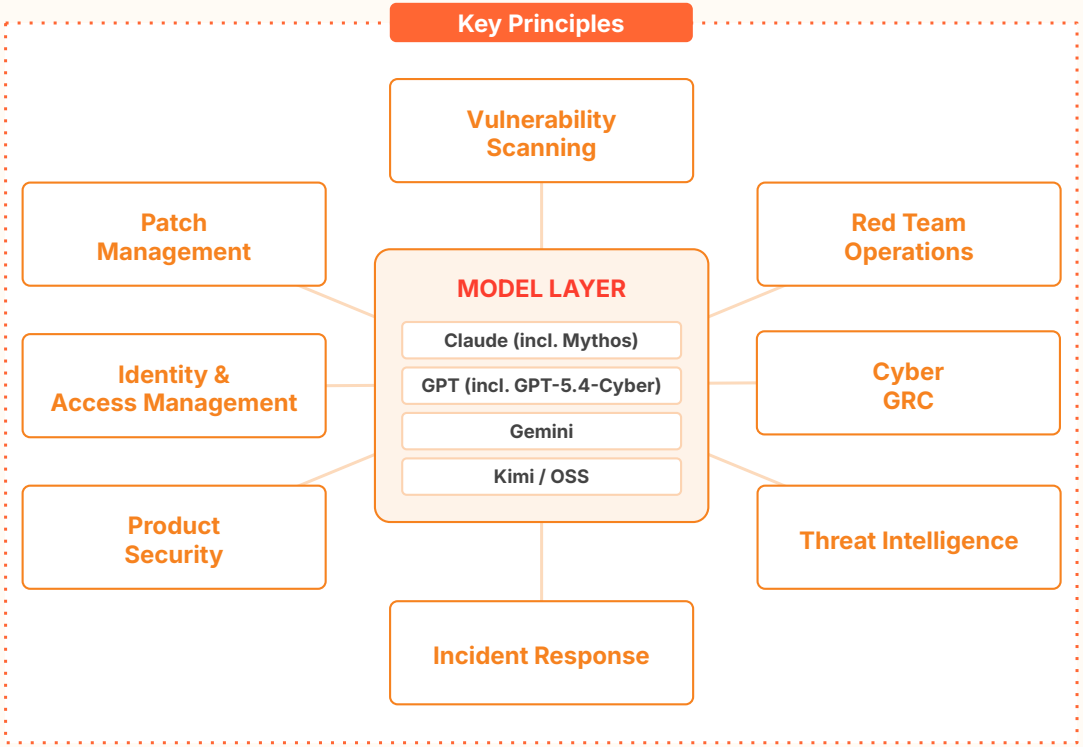

MODELS WE USE TODAY

Anthropic Mythos

+ Chains primitives into full exploit paths
– Guardrails refuse PoCs & suppress findings

OpenAI GPT-5.4-Cyber

+ Fewer guardrails — probes live targets
– Less nuanced exploit chaining

Kimi / Open-Source

+ Cheaper per token — volume scanning
– More noise on complex chains

Cloudflare's Reference Architecture for AI Usage

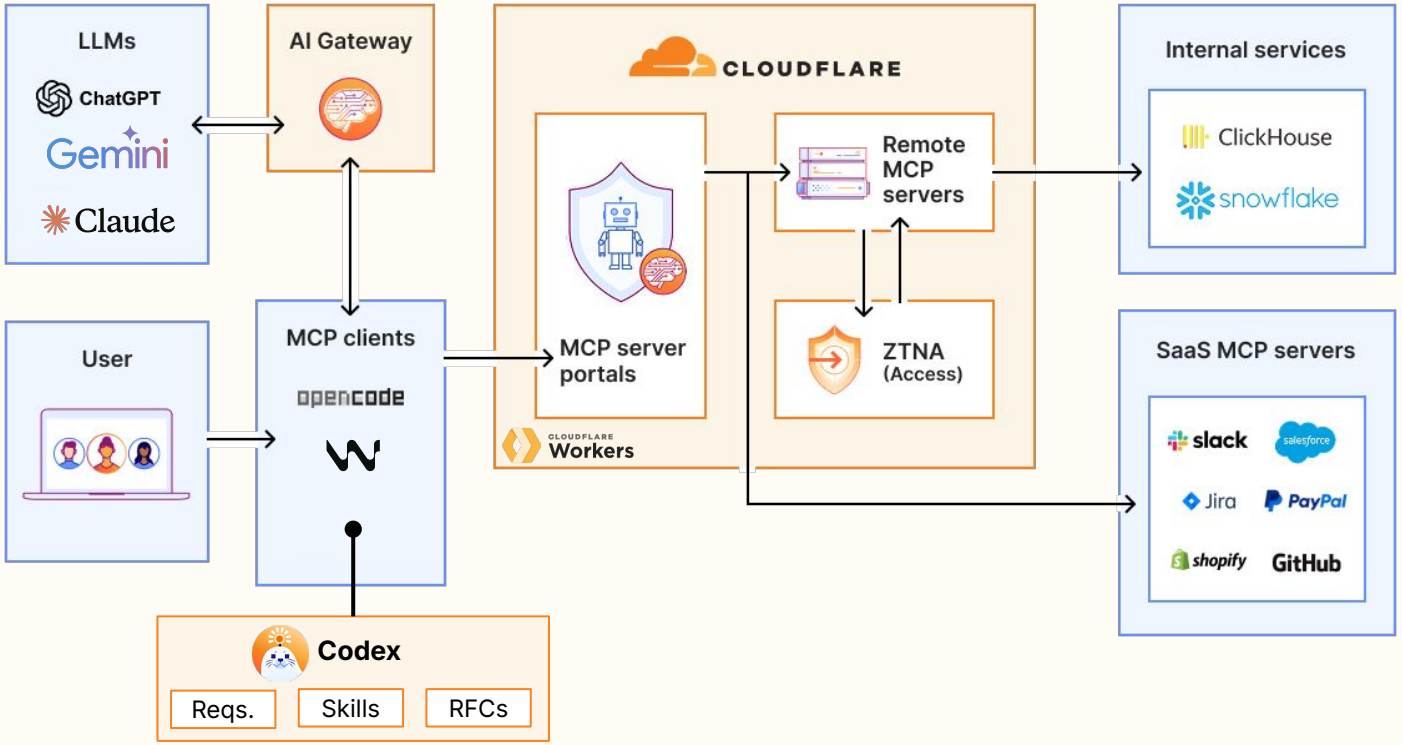

Our Security AI Harness Approach

How are we approaching Harnesses?

A harness is the brain behind the AI.

- It decides what the model looks at, in what order, and what to do with the results.
- Without it, you have a capable model with no direction. With it, you have a security pipeline.

Model-agnostic

- The harness outlasts any model. Swap Mythos for the next frontier model — the pipeline stays the same.
- Preparing for more capable models that may come in the future.

Red Team AI Harness Example

- Shown on the right
- Stages:
  Recon → Hunt → 2nd Opinon → Validate → Dedup → Report

Red Team Ops Harness

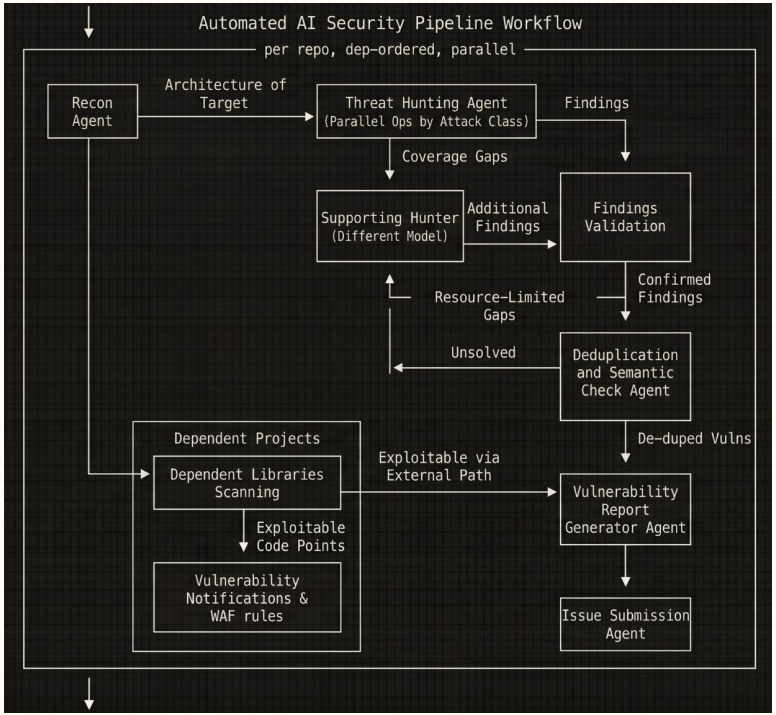

Harness used for AI-enabled Scans

Target Architecture Context

- Gives AI agents full repo context (dependencies, exposure paths, service relationships, etc.) before hunting begins.

Model-Agnostic Hunting Agents

- Hunts across attack classes in parallel. Swap the model anytime; the pipeline logic stays constant.

Supporting Hunter

- Independent second model challenges primary findings, reducing noise & providing different perspective.

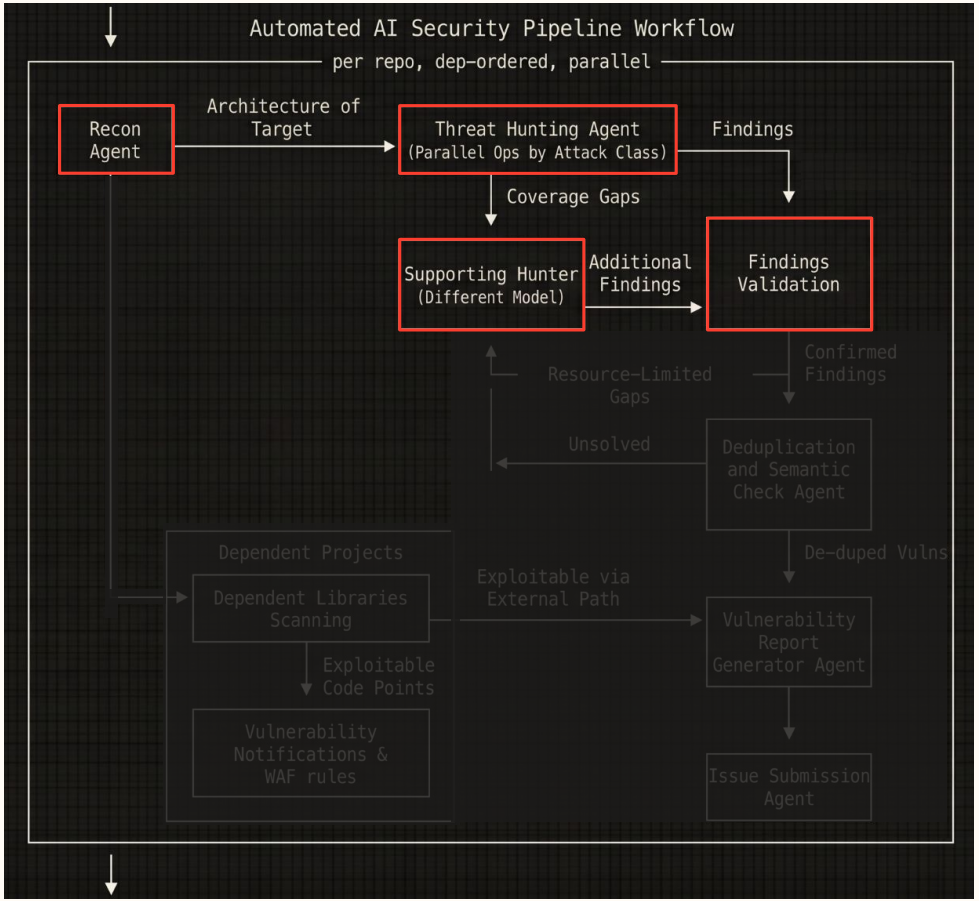

Harness used for AI-enabled Scans

Dependent Libraries / OSS Scans

- Identifies exploitable code paths in third-party dependencies / OSS libraries that are used by target repo.

WAF Protections

- Confirmed vulnerabilities automatically generate WAF rule candidates for rapid WAF rule creation.

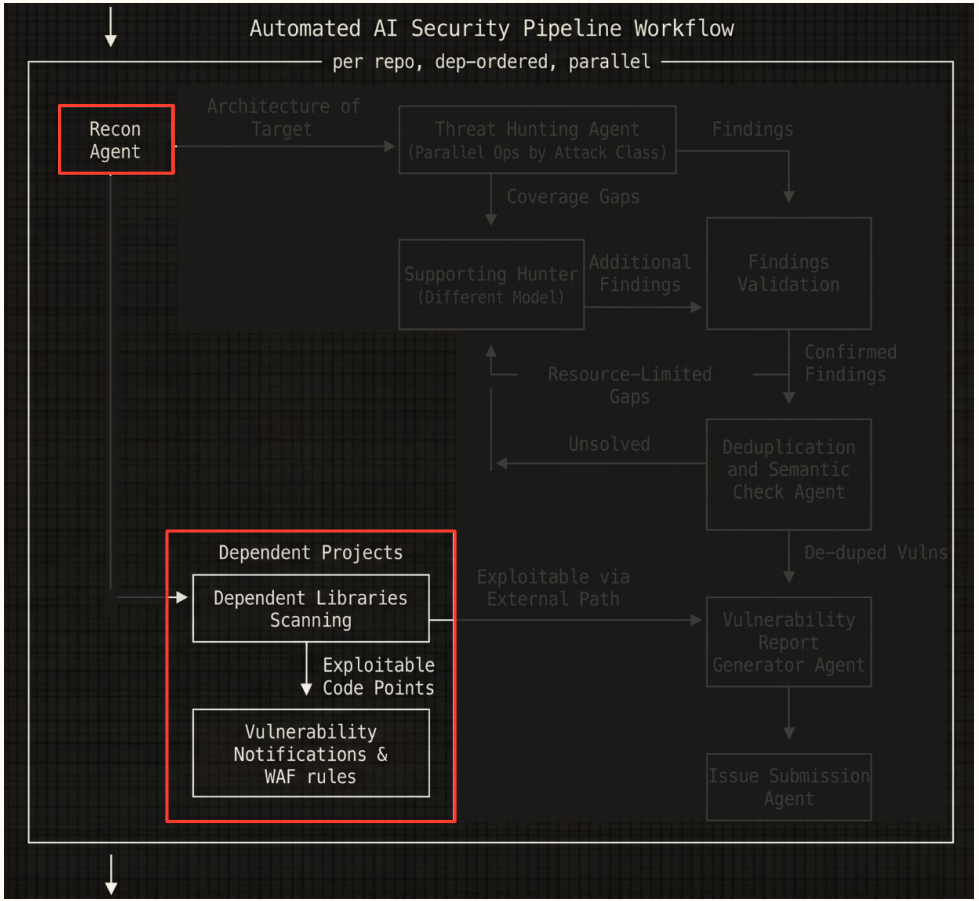

Harness used for AI-enabled Scans

Deduplication & Exploitability Check

- Agent semantically de-duplicates findings across hunters, surfacing only confirmed, distinct vulnerabilities.

Exploit Proof-of-Concept

- Generates working PoC code for confirmed findings - helps defensive teams prioritize based on real exploitability v/s theoretical risk.