CLOUDFLARE

AI-Powered Operations

Secured by Cloudflare

ACME Corp · AI Demo Track · 2026

ACME Corp — Who They Are

THE ORGANISATION

~600

employees · 60% remote

Industry: Financial Services

Revenue: ~$280M ARR

Locations: Chicago (HQ), Dallas, Sydney, London, Singapore

Regulated by: APRA (CPS 234)

AI TOOLS IN USE TODAY

M365 Copilot — 22 pilot licences (IT-approved)

GitHub Copilot Enterprise — engineering team (IT-approved)

ChatGPT (personal) — 31 employees (not approved)

Gemini, Perplexity — 16 employees (not approved)

KEY AI CONCERNS

- No visibility into what employees submit to AI tools
- Client financial data potentially in personal AI sessions
- No audit trail for APRA CP 009-2026 compliance
- Board asking hard questions after phishing incident
- M365 Copilot summarised a board pack containing M&A discussions

SPONSORS

David Kim (CTO) · **Sarah Chen** (CISO)

Download PDF

AI Is Already In Your Organisation

The uncontrolled reality

- **47 ACME employees** using personal AI tools — no policy, no visibility
- Client portfolio summary uploaded to personal ChatGPT by Finance
- Board pack — including M&A discussions — summarised in M365 Copilot
- C2 domain active 6 days before phishing attack — DNS had no way to block it
- Zero APRA-ready audit trail of any AI interaction

Risk: AI is already operating — without controls, without accountability.

The controlled alternative

- AI agent reads emails, calendars, and docs — prepares your team in seconds
- Maps security incidents to Cloudflare controls in real time
- Generates board decks from live incident data
- Blocks unapproved AI domains. Enterprise platforms (M365, Google) restricted to company accounts via tenant restrictions
- Every prompt logged. Every action auditable. APRA-ready.

Cloudflare: The AI is doing real work. We make sure it's doing it safely.

This demo runs on a real AI agent — and **every action it takes is secured by Cloudflare in real time**. The demo is the proof point.

CLOUDFLARE

ACME CORP

The AI Reality

Three incidents that demand action

Download PDF

ACME Corp — What the AI Scan Found

HIGH AI-001 — ChatGPT

Who: Finance team member

What: Q1 Client Portfolio Summary (47 pages, CONFIDENTIAL)

Uploaded to: Personal ChatGPT account

Risk: 4,200 client holdings exposed. APRA reportable.

HIGH AI-002 — Copilot

Who: VP IT (Tom Walsh)

What: Board pack for May 22 meeting

Used: M365 Copilot to summarise

Risk: M&A discussions, exec compensation — price-sensitive.

MED AI-003 — Shadow AI

Who: 47 employees

Tools: ChatGPT (personal), Gemini, Perplexity

Using for: Document summarisation, drafting, analysis

Risk: Unknown data exposure. No audit trail.

APRA is requesting AI governance evidence under **CP 009-2026**. The board has been notified. David Kim (CTO) needs a solution before June 12.

CLOUDFLARE

LIVE DEMONSTRATION

Six AI Scenarios

Each powered by AI. Each secured by Cloudflare.

Download PDF

CLOUDFLARE

ARCHITECTURE

Platform Architecture — How Every Component Connects

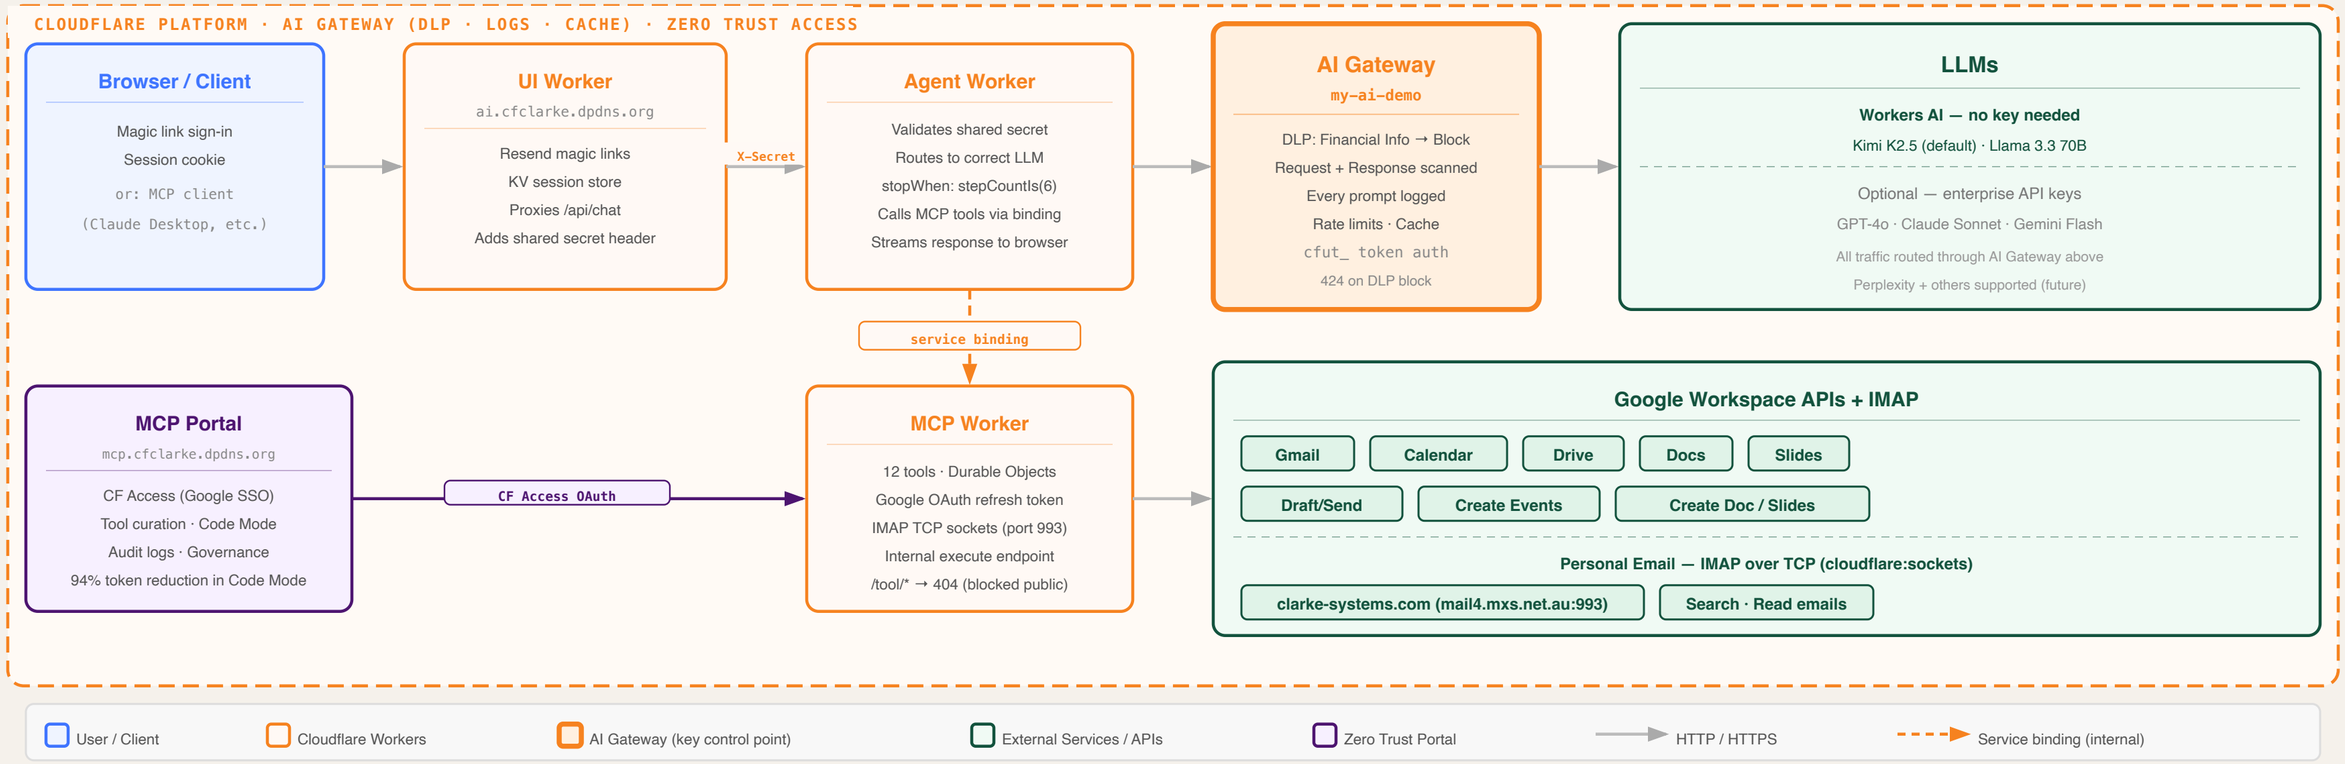

Download PDF

1 Instant Account Intelligence

THE CHALLENGE

45 minutes manually reading 4 email threads, calendar events, and call notes before every customer call

Context scattered — Gmail, calendar, Excalidraw diagrams, Google Docs. No consolidated view.

Critical detail missed: Sarah's DDoS email mentions the $28K scrubbing bill — buried in thread 3

THE CLOUDFLARE OUTCOME

- **Speed:** 45 minutes → 30 seconds. Every stakeholder, open question, and talking point surfaced automatically.
- **Security: Cloudflare DLP** watches all ACME data flowing through the AI. **Access** ensures only authenticated users reach AI tools.
- **Audit:** Every read, every synthesis, every output — logged in **AI Gateway**. Full record. APRA-ready.

Access | AI Gateway | DLP

Download PDF

CLOUDFLARE

2 Phishing Killed at Every Layer

WHAT HAPPENED

AiTM proxy defeated **Microsoft Authenticator MFA** — credentials harvested, session hijacked

3 accounts compromised, **4,200 Salesforce records** exfiltrated, lateral movement attempted

C2 domain active in threat intel for **6 days** — on-prem DNS had no way to act on it

FIVE CONTROLS. KILL CHAIN BROKEN.

- **Email Security** — typosquat domain caught before delivery
- **Browser Isolation** — phishing page rendered remotely; credentials never reach the device
- **WARP + Device Posture** — session anomaly detected; access denied before lateral movement
- **Gateway DNS** — C2 domain blocked at resolution, Day 1. Before a packet leaves.

Email Security | Browser Isolation | WARP | Gateway DNS | Access

Download PDF

3 Security Architecture — Built from an Email

THE CHALLENGE

Architecture diagrams take hours — Visio, Lucidchart, or whiteboard sessions with multiple engineers

Marcus Rodriguez sent 5 detailed technical questions about ACME's environment. Each answer needs a diagram.

Current state diagram is 8 months out of date. No one owns it.

THE CLOUDFLARE OUTCOME

- **Speed:** Agent reads Marcus's email and generates a before/after architecture diagram. Subnet `10.44.0.0/16`, PA-5220, VPN gaps — all extracted automatically.
- **Accuracy:** The proposed architecture is verified against Cloudflare documentation before being committed to the diagram.
- **Security:** Diagram tools (Excalidraw, Drive) are behind **Cloudflare Access**. The AI can only reach what you authorise.

Access | AI Gateway | Gateway DNS

4 PoC Planning — Minutes, Not a Morning

THE CHALLENGE

PoC scoping is a half-day exercise — success criteria, test plans, prerequisites, rollback procedures

Marcus raised 5 specific technical concerns: split tunnel for `10.44.0.0/16`, DoH bypass, Splunk integration, AnyConnect coexistence, SAML app registration

If they're not in the PoC doc, they become surprises on Day 1

THE CLOUDFLARE OUTCOME

Complete: Success criteria, 2-week sprint, stakeholder sign-off matrix, rollback plan — all generated from ACME's own discovery notes.

Precise: "Core banking app on `10.44.0.0/16` accessible via ZTNA without VPN" — extracted directly from Marcus's email. Not invented.

Protected: AI Gateway prompt injection detection — if Marcus's email contained hidden instructions, they're caught before reaching the model.

Access (ZTNA) | Gateway | AI Gateway | Logpush → Splunk

Download PDF

5 Board Deck Built and Securely Delivered

THE CHALLENGE

Sarah Chen needs a board-ready security narrative by May 22 — DDoS incident, phishing response, proposed architecture, ROI case

Building a polished, accurate board deck from scratch takes a skilled person 2–3 days

Sharing via Google Drive — no access control. Anyone with the link can view sensitive security content.

THE CLOUDFLARE OUTCOME

Built: Board-ready presentation from ACME's own incident data — $140K DDoS loss, $28K scrubbing eliminated, phased security roadmap. Minutes, not days.

Delivered securely: Deployed on **Cloudflare Workers**. Magic link issued — only sarah.chen@acmecorp.com can open it. Expiring. No shared links.

Logged: Every prompt used to build the deck is recorded in **AI Gateway**. Complete record of what went in and what came out.

Workers | Access Magic Links | AI Gateway

Download PDF

CLOUDFLARE

6 Shadow AI — Two Layers of Control

THE DISCOVERY

47 employees using personal AI tools — ChatGPT, Gemini, Perplexity via browser

Client portfolio in ChatGPT. Board pack in M365 Copilot. API keys in GitHub Copilot.

APRA requesting AI governance evidence under **CP 009-2026**. No audit trail exists.

LAYER 1 – BROWSER AI CONTROL

For employees browsing to ChatGPT, Gemini, Perplexity

- **Gateway DNS** — logs every AI domain lookup. Block or allow by category, Day 1.
- **Gateway HTTP** — AI domains blocked by category (SNI inspection). Enterprise platforms enforce tenant restrictions — company accounts pass, personal accounts denied.
- **DLP** — CONFIDENTIAL files and financial data blocked from all AI web destinations.

Gateway DNS | Gateway HTTP | DLP

LAYER 2 – API AI CONTROL

For enterprise apps, Copilot, and AI agents making API calls

- **AI Gateway** — every prompt logged: user, model, timestamp. Prompt injection detected. Rate limits enforced.
- **APRA audit trail** — per-prompt log with full context. Exportable. Board-ready. CP 009-2026 compliant.
- **This demo** — OpenCode is making API calls to Claude through AI Gateway right now.

AI Gateway | Logpush

Download PDF

THE RESULT

AI You Can Trust

Productive, auditable, and policy-controlled

CLOUDFLARE

Before & After — AI Track

~~45 min~~ 30 sec Meeting prep time	~~Undetected~~ Blocked Shadow AI exposure	~~None~~ Complete APRA AI audit trail	~~3 accounts~~ 0 compromised Phishing outcome

AI GATEWAY

Every LLM API call logged, rate-limited, and auditable. Prompt injection detection active.

DLP + GATEWAY

Financial data, PII, CONFIDENTIAL files blocked from all AI endpoints. Personal AI tools blocked at DNS.

LOGPUSH → SPLUNK

Complete AI interaction audit trail streamed to Splunk. APRA CP 009-2026 evidence. Board-ready.

Download PDF

Two Layers of AI Control

LAYER 1 – BROWSER AI & WEB CONTROL

Employees browsing to ChatGPT, Gemini, Perplexity

- **Gateway DNS** — visibility and blocking of AI domains, Day 1
- **Gateway HTTP** — block AI domains by category; tenant restrictions on M365/Google enforce company-account-only access
- **DLP** — intercept CONFIDENTIAL data before it reaches any AI website

LAYER 2 – API-LEVEL AI CONTROL

Enterprise apps, Copilot integrations, AI agents making API calls

- **AI Gateway** — every API call proxied: prompt logged, injection detected, rate-limited
- **Access** — Zero Trust auth on every AI tool and MCP connection
- **Logpush → Splunk** — full per-prompt audit trail, APRA CP 009-2026 ready

"The question isn't whether your teams will use AI. They already are.

The question is whether you're in control — and whether you can prove that to your board and regulator."

THE META-NARRATIVE

This AI agent is itself secured by Cloudflare right now. Every prompt logged. Every tool access gated. Every output auditable. **You're watching it live.**

Download PDF

Cloudflare One

AI that works. AI you can trust.

IMMEDIATE

Deploy AI Gateway in front of your approved AI tools — visibility in under an hour

JUNE 5

AI Security Strategy Briefing — David Kim, Sarah Chen, board prep

BOARD – JUNE 12

Present APRA CP 009-2026 compliance plan — AI Gateway audit trail as evidence