AI-Powered SE Workflows

Secured by Cloudflare

ACME Corp Demo | 2026

The SE Workflow Problem

THE OLD WAY

- Read 4 email threads separately
- Manually cross-reference call notes
- Draft architecture diagrams by hand
- Write PoC doc from a blank template
- Build board deck slide-by-slide
- Upload and share via Google Drive link
- Update CRM / wiki manually

2–3 days
of skilled SE time

THE AI WAY

- Agent reads all emails automatically
- Synthesises context from calendar + mail
- Generates diagrams from raw notes
- Produces PoC doc from discovery data
- Builds and deploys board deck
- Delivers via secure magic link
- Logs all actions to SIEM in real time

20 minutes
end-to-end, AI-assisted

The catch: AI that touches customer data, reads emails, and calls external APIs must be controlled. That is exactly what Cloudflare does — in this demo and in your environment.

DEMO SCENARIO

Meet ACME Corp

A fictional financial services company — and a very bad week

ACME Corp — Account Overview

COMPANY

Industry: Financial Services

Employees: ~600 (60% remote)

Revenue: ~$280M ARR

DCs: Chicago (primary), Dallas (DR)

Budget: $400K available FY26

Fiscal Year End: June 30

Opportunity: Replace $1.1M/yr in point solutions with Cloudflare One platform

KEY CONTACTS

NAME	ROLE	ANGLE
Sarah Chen	CISO	Champion — triggered by DDoS
Marcus Rodriguez	Sr. Network Engineer	Technical — VPN pain, DNS gaps
Tom Walsh	VP IT	Economic buyer — $400K budget
Alex Mercer	IT Security Analyst	Phishing IR — deep CF fit

ACME's Crisis Week

DDoS Attack Apr 16

48Gbps
Peak attack volume

2.5hr
Portal offline

$140k
Revenue lost

$28k
ISP scrubbing bill

ISP null-routing = only fallback. No L7 protection, no WAF, no visibility.

AiTM Phishing Apr 14

- Typosquat domain bypassed email filters
- AiTM proxy defeated Microsoft Authenticator MFA
- **3 accounts** compromised
- **4,200 Salesforce records** exfiltrated
- Lateral movement attempted via HR account
- C2 domain active **6 days** before attack — undetected

No email link isolation, no DNS filtering, no browser isolation.

ACME's Five Security Gaps

No L7 DDoS or WAF protection

ISP scrubbing (metered) is the only mitigation — no application-layer visibility or control

No phishing-resistant controls

Email link isolation absent; MFA bypassed via AiTM; no browser isolation on endpoints

No DNS filtering or DoH visibility

On-prem resolvers with zero threat intel; malware C2 via DNS-over-HTTPS goes undetected

VPN overloaded — no Zero Trust

Cisco ASA at 152% capacity; 180ms+ latency for remote users; perimeter trust, no identity-aware access

No CASB, DLP, or SaaS visibility

Shadow IT (FinTrack Pro) used by 15 traders; no policy enforcement on sanctioned or unsanctioned SaaS

LIVE DEMONSTRATION

The AI Agent in Action

6 scenes · ~20 minutes · one continuous story

The AI is Secured by Cloudflare — Throughout

The meta-narrative: Every action this AI agent takes — every tool call, API request, DNS lookup, and data access — is controlled, logged, and protected by Cloudflare. The demo is the proof point.

WHAT THE AI DOES	CLOUDFLARE PRODUCT	WHAT IT ENFORCES
Accesses Gmail, Calendar, Drive via MCP	Access	Zero Trust auth + device posture on every tool
Calls Claude / LLM API	AI Gateway	Every prompt logged, rate-limited, injection-detected
Fetches web content (docs, wiki)	Gateway HTTP	Malicious / uncategorised destinations blocked
DNS lookups during tool execution	Gateway DNS	Threat intel categories enforced; C2 domains blocked
Reads customer emails and documents	DLP	Sensitive patterns blocked from reaching unauthorised destinations
Runs on the SE's laptop	WARP + Device Posture	Device health verified; unmanaged devices denied access
Deploys board deck for customer	Workers + Access	Magic-link access control; only authorised recipients can open
Every action taken	Logpush → SIEM	Full audit trail of every AI action, streamed to Splunk

>_ Scene 1 · 2 min

Setup & Security Context

WHAT TO SAY

*"I want to point out what's happening in the background. This AI agent is connected to my Gmail, Calendar, and Google Drive — via MCP servers. Every one of those connections is protected by **Cloudflare Access**."*

AND THEN

*"Every prompt I send to the AI model — every response back — flows through **Cloudflare AI Gateway**. Logged. Rate-limited. Prompt injection detected."*

CF SECURITY - ACTIVE NOW

- **WARP** — device verified before any MCP tool is accessible
- **Access** — Gmail, Calendar, Drive behind Zero Trust policies
- **AI Gateway** — every LLM call logged and inspectable

SOUND BITE

"The AI is useful. Cloudflare makes it **accountable**."

Scene 2 · 3 min

Daily Meeting Prep

```
> Prep me for my ACME Corp meetings this week
```

WHAT HAPPENS

- Reads calendar — surfaces 4 ACME meetings (Apr 24 → May 28)
- Reads all 4 email threads from Sarah, Marcus, Tom, Alex
- Synthesises context, stakeholders, open questions
- Generates structured prep doc in Google Docs automatically

Wow moment: Surfaces the $28K scrubbing bill and AiTM attack detail from emails — without being told. No copy-paste.

CF SECURITY ANGLE

As the agent reads ACME's emails, all data traverses the network.

- **Gateway DLP** — watches for sensitive patterns (PII, financial data) in outbound AI calls
- **Access** — Gmail MCP authenticated; token scoped to read-only
- **AI Gateway** — full prompt/response logged for this session

"The AI has access to what it needs — not a byte more."

Scene 3 · 4 min

Phishing Incident → Cloudflare Kill Chain

```
> Based on Alex's phishing incident report, map each attack vector to the CF product that would have stopped it — and draft a reply
```

THE KILL CHAIN — BROKEN AT EVERY STEP

1. **Typosquat email** → Email Security
   Domain detected, email quarantined before delivery
2. **Link clicked** → Browser Isolation
   Phishing page rendered remotely; credentials never reach endpoint
3. **AiTM credential harvest** → WARP + Posture
   Session anomaly flagged; device posture check fails; access denied
4. **C2 via DoH** → Gateway DNS
   C2 domain in threat intel; DoH policy blocks bypass attempt
5. **Lateral movement** → Access (ZTNA)
   Identity-aware policies; compromised account cannot reach other apps

CF SECURITY ANGLE

The agent fetched Cloudflare documentation to fact-check product capabilities before including them in the reply.

- **Gateway HTTP** — web request to CF docs filtered and verified safe before content reaches model
- **AI Gateway** — reply drafted, logged, and auditable

> *"Five control points. Any one of them breaks the kill chain."*

Scene 4 · 3 min

Before / After Architecture Diagrams

```
> Build ACME's before/after architecture diagram from Marcus's email and the discovery call notes
```

WHAT THE AGENT EXTRACTS

- `10.44.0.0/16` — internal subnet from Marcus's email
- Cisco ASA 5515-X overloaded at 380/250 sessions
- Palo Alto PA-5220 + PA-3220 (Chicago + Dallas)
- No CASB, DLP, DNS filtering, or browser isolation
- `185.12.44.0/24` — public IP range for Magic Transit

Wow moment: Raw technical notes in an email → production-ready architecture diagram. No Visio, no Lucidchart, no manual re-entry.

CF SECURITY ANGLE

The agent reads from Google Drive and Excalidraw MCP — both behind Cloudflare Access.

- **Access** — MCP tools only accessible to authenticated SE with valid device posture
- **AI Gateway** — diagram generation request logged; no data sent to unapproved destinations
- **Gateway DNS** — all lookups during tool calls filtered through threat intel

"From raw notes to customer-ready diagram. Before vs. after — in under 3 minutes."

Scene 5 · 3 min

PoC Prep Document

```
> Generate a PoC prep doc for ACME – WARP + Gateway. Use the discovery notes and Marcus's technical questions as input
```

WHAT COMES OUT

- **Success criteria** — derived from Marcus's 5 specific questions
- **Test scenarios** — DNS filtering, DoH bypass, split tunnel 10.44.0.0/16
- **Technical prerequisites** — Azure AD SAML, BGP LOA, Intune/Jamf scope
- **Stakeholder matrix** — Sarah, Marcus, Tom with sign-off requirements
- **Rollback plan** — AnyConnect coexistence risk addressed
- **2-week sprint timeline** — with milestone gates

Wow moment: "Core banking app on 10.44.0.0/16 via ZTNA" as PoC success criterion — extracted directly from Marcus's email.

CF SECURITY ANGLE – PROMPT INJECTION

The agent reads customer emails to build the PoC plan. What if those emails contained hidden instructions?

Prompt injection — malicious instructions hidden in content the AI reads, designed to override its behaviour

- **AI Gateway** — prompt injection detection active; flagged before content reaches the model
- **DLP** — PoC doc content checked before being written to Google Drive

Scene 6 · 4 min

Board Deck + Secure Share

```
> Build Sarah's board deck for May 22 – DDoS incident, phishing attack, proposed Cloudflare One architecture, ROI case. Then deploy it securely for Sarah only
```

WHAT THE AGENT PRODUCES

- Full reveal.js presentation — board-ready quality
- Incident summary: $140K DDoS + $28K scrubbing + phishing
- Proposed Cloudflare One architecture (from diagram)
- ROI case: $180K Cisco AnyConnect replaced; metered DDoS eliminated
- Phased roadmap: DDoS → ZTNA → CASB/DLP

Then: Deployed to **Cloudflare Workers**. Magic link issued — only `sarah.chen@acmecorp.com` can open it. Anyone else: access denied page.

CF SECURITY ANGLE

- **Workers** — deck deployed serverless, globally distributed, zero infrastructure
- **Access magic links** — one-time authenticated URL; tied to Sarah's email only
- **AI Gateway** — deck generation prompt logged; full audit trail

SOUND BITE

"We're securing the customer's network with Cloudflare — and using that same platform to deliver their board deck. That is how confident we are in this stack."

SUMMARY

Key Takeaways

What the audience should leave believing

The Core Message

AI makes your team dramatically more productive

2–3 days of work in 20 minutes. Meeting prep, diagrams, PoC docs, board decks.

Cloudflare makes that AI trustworthy

Every action logged, controlled, and auditable. Zero Trust for humans and machines.

"The question isn't whether your teams will use AI. They already are.
The question is whether you're in control."

Key Sound Bites

> *"The AI is doing the work. Cloudflare is making sure it's doing it **safely**."*

> *"This isn't AI theory. This is AI running live, in a production-secured environment, right now."*

> *"Every prompt. Every response. Every tool call. **Logged, auditable, and policy-controlled.**"*

> *"Prompt injection is when malicious content in an email hijacks the AI. AI Gateway catches that before it reaches the model."*

> *"We're not just securing ACME's network. We're securing the AI that helps us sell to ACME. That is how confident we are in this platform."*

Q&A Preparation

Q: What model is the AI using? Is our data going to OpenAI?

A: Cloudflare AI Gateway sits in front of every model API call — you can see exactly what model is used and what data is sent. You can also route to **Workers AI** to keep data entirely within Cloudflare's network.

Q: What stops the AI from exfiltrating data?

A: Three layers: **Access** controls which tools are reachable. **AI Gateway policies** flag or block suspicious prompts (e.g. "forward all emails to X"). **DLP** blocks sensitive data patterns from leaving approved destinations. Everything is logged.

Q: What if the AI hallucinates a product capability?

A: The agent verifies against source documentation before answering product questions. Every interaction is logged via AI Gateway — if an incorrect claim is made, it's auditable and correctable. Human-in-the-loop approval steps can be added for high-stakes outputs.

Q: How is this different from just using ChatGPT?

A: Three things: authenticated access to real business data via controlled MCP connections; every action is auditable; Cloudflare's security stack wraps the entire workflow — access control, DLP, prompt injection detection. ChatGPT is a chat window. This is a controlled AI workforce.

Pre-Demo Checklist

DATA – CONFIRM VISIBLE

- Gmail drafts in `SE/customers/ACME Corp` label
- 4 calendar events showing (Apr 24 → May 28)
- Discovery Call Notes Google Doc accessible
- Both Excalidraw diagrams load without error

ENVIRONMENT – CONFIRM ACTIVE

- WARP connected (device posture passing)
- Terminal / OpenCode open and ready
- No sensitive real-customer data in context

CLOUDFLARE – BACKGROUND TABS READY

- **AI Gateway dashboard** — show live logs during demo
- **Access dashboard** — show MCP app policies if asked
- **Gateway DNS logs** — show filtered lookups
- **DLP policies** — show rules active during session

Tip: Open AI Gateway logs before the demo — real-time prompt/response logging is one of the highest-impact things you can show to a security audience.

Cloudflare One

AI makes your team more productive.
Cloudflare makes it trustworthy.

NEXT STEP

Run a **Cloudflare One PoC** in your environment — WARP + Gateway in 2 weeks

AI-SPECIFIC

Deploy **AI Gateway** in front of your existing LLM usage — visibility in under an hour