Strangers in your *home*

A NORTHLIGHT PICTURES PRODUCTION A FILM BY JONAS EIKEN
"STRANGERS IN YOUR HOME" LARS EKLUND ELLA NORDSTRÖM ADAM LUNDGREN
DIRECTOR OF PHOTOGRAPHY MATTIAS LINDGREN PRODUCTION DESIGN SOFIA KARLSSON
EDITED BY ERIK DAHL MUSIC BY JOHAN SÖDERQVIST
WRITTEN BY JONAS EIKEN PRODUCED BY MIKAEL BERG DIRECTED BY JONAS EIKEN

Controlling access of Claude & Co

Runtime governance for AI agents

Johannes Keienburg · Co-Founder & CEO Cakewalk Security

Claude Code Munich · June 10, 2026

What’s your guess?

List of all 51 MCP tools available (as of June 23 2025) on GitHub official Remote MCP (via streamable http) server at https://api.githubcopilot.com/mcp/

github-mcp-tools.json

Raw


```
[
    {
        "inputSchema": {
            "json": {
                "properties": {
                    "body": {
                        "description": "Comment content",
                        "type": "string"
                    },
                    "issue_number": {
                        "description": "Issue number to comment on",
                        "type": "number"
                    },
                    "owner": {
                        "description": "Repository owner",
                        "type": "string"
                    },
                    "repo": {
                        "description": "Repository name",
                        "type": "string"
                    }
                },
                "required": [
                    "owner",
                    "repo",
                    "issue_number",
                    "body"
                ],
                "type": "object"
            }
        },
        "name": "add_issue_comment",
        "description": "Add a comment to a specific issue in a GitHub repository."
    },
    {
        "inputSchema": {
```


cakewalk

GitHub:
51 tools

cakewalk

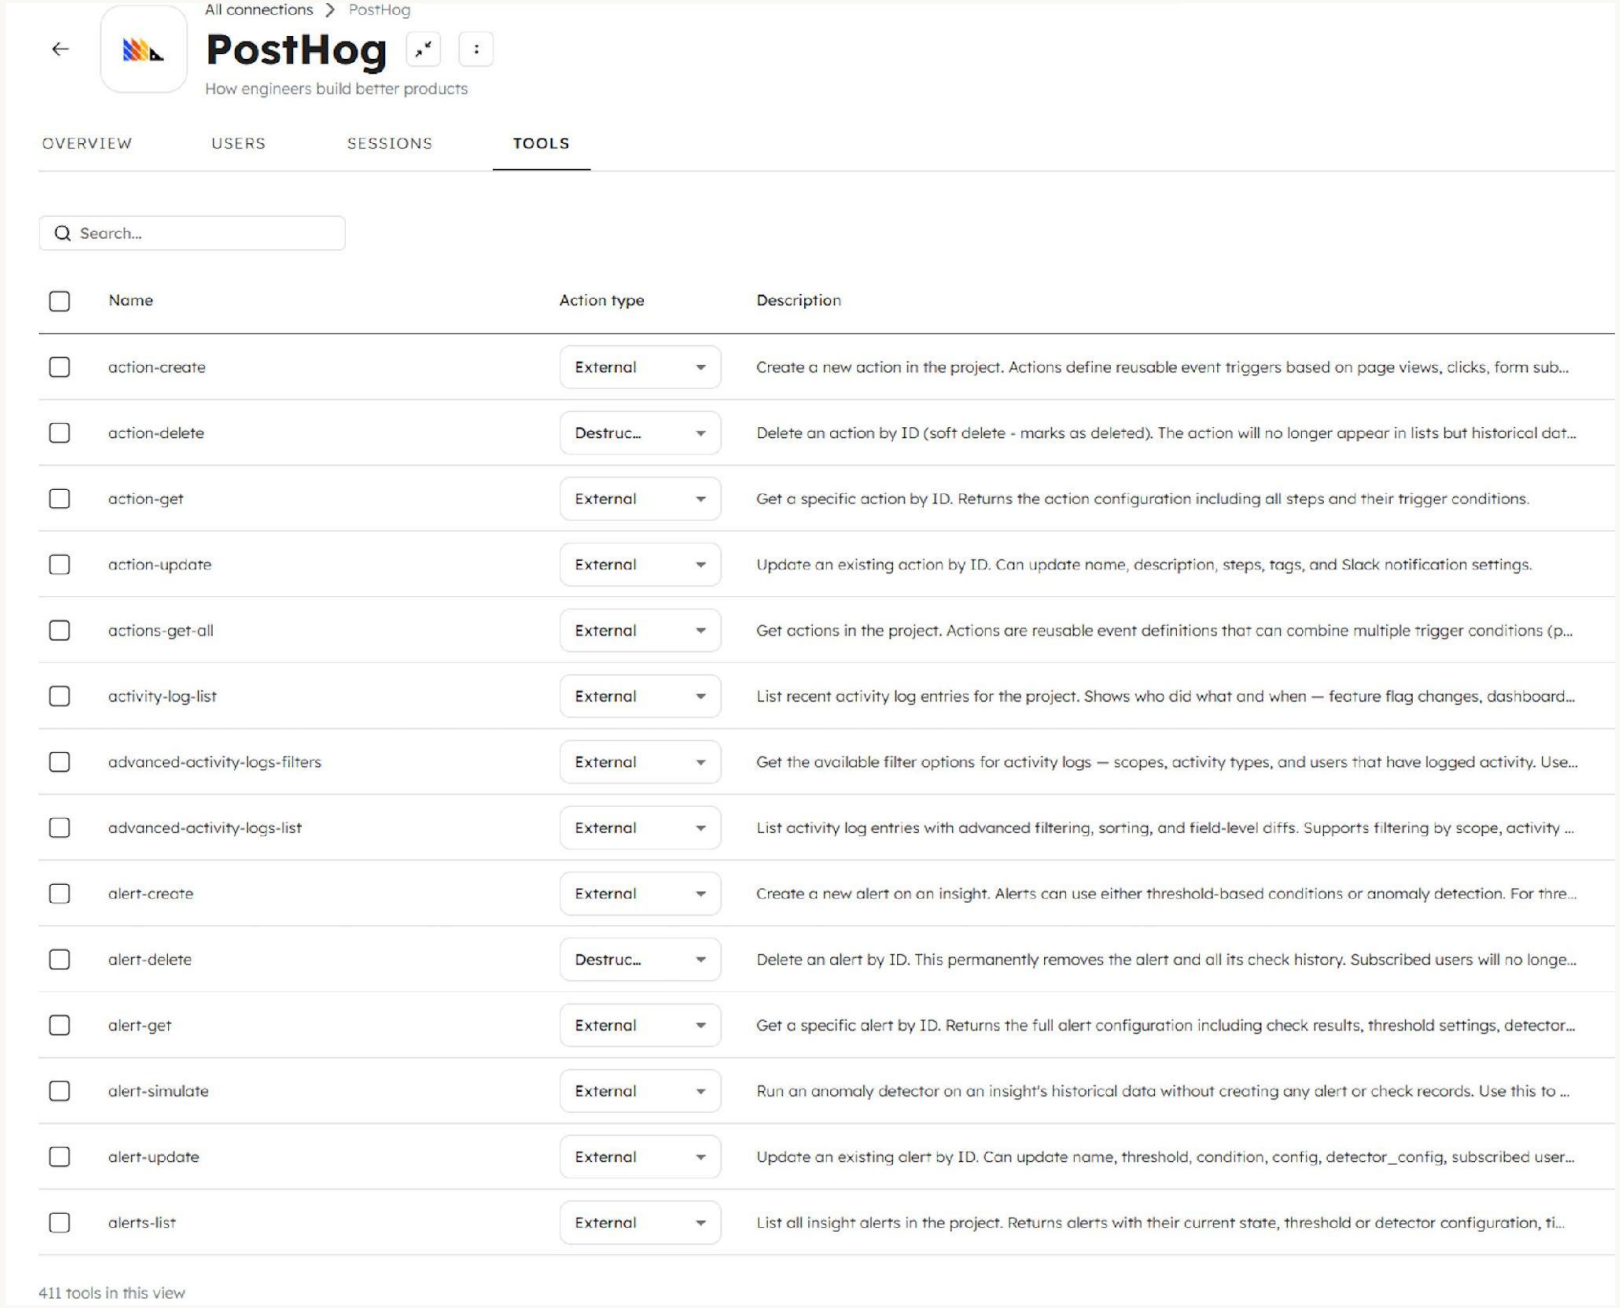

PostHog:
411 tools

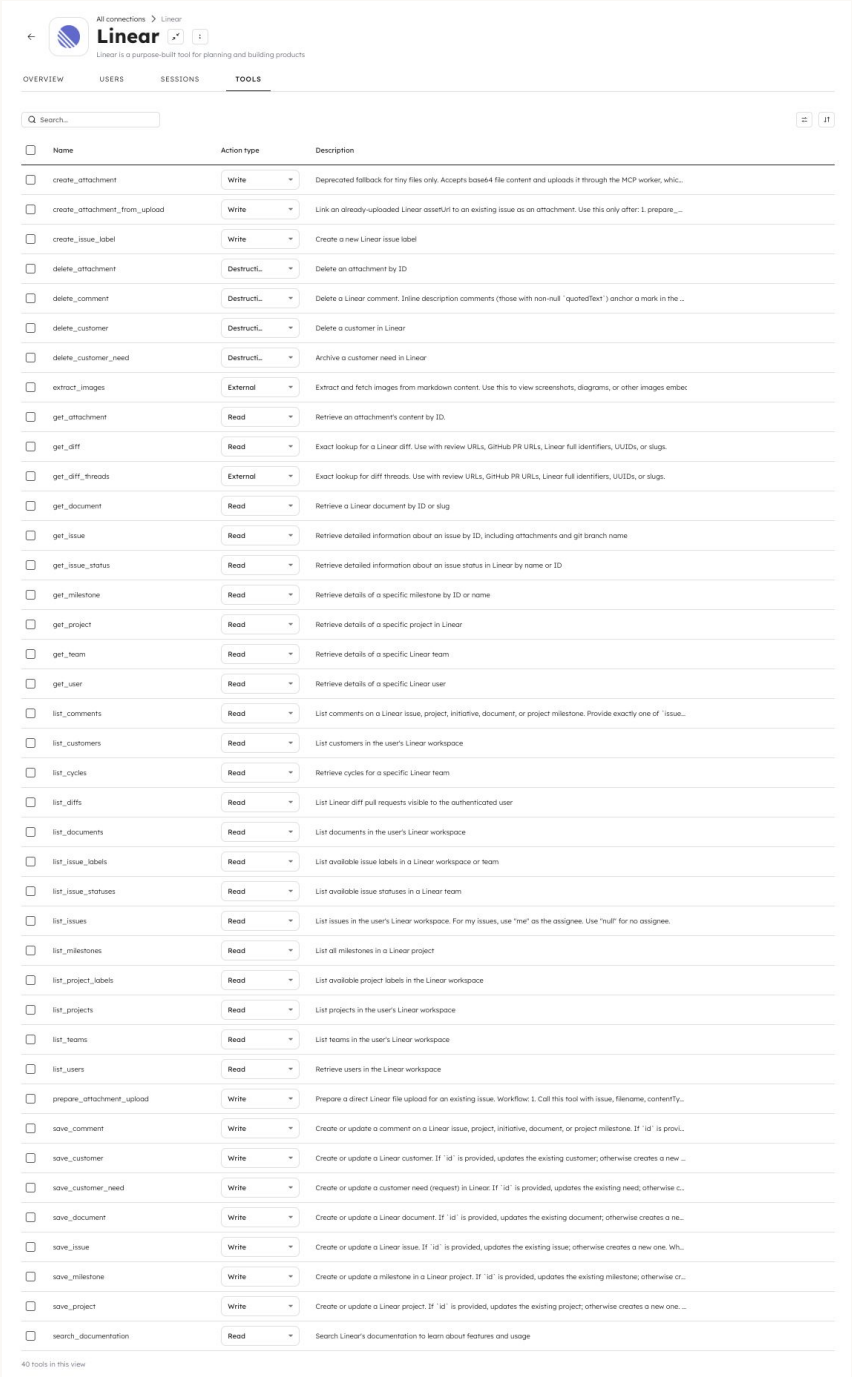

cakewalk

Linear:
40 tools

Zero access governance

cakewalk

Static access

Broad permissions

No central visibility


Agentic Infrastructure.
▲Vercel
Framer
Enterprise needs.
Startup speeds.
OUTFRONT
14'x48'
In Memoriam
The Billable Hour
In lieu of flowers, send contracts to crosby.ai
CROSBY
The First AI Law Firm
001008
Agents trace.
You draw.
Graphite
LangChain
Track your agents like your macros
Monitor agents with
LangSmith
Agents work where you don't.
Retool
ONE WAY
ON RED
Pay here
Pay here
SAN FRANCISCO · MAY 2026

How Vercel got breached via an agent

cakewalk

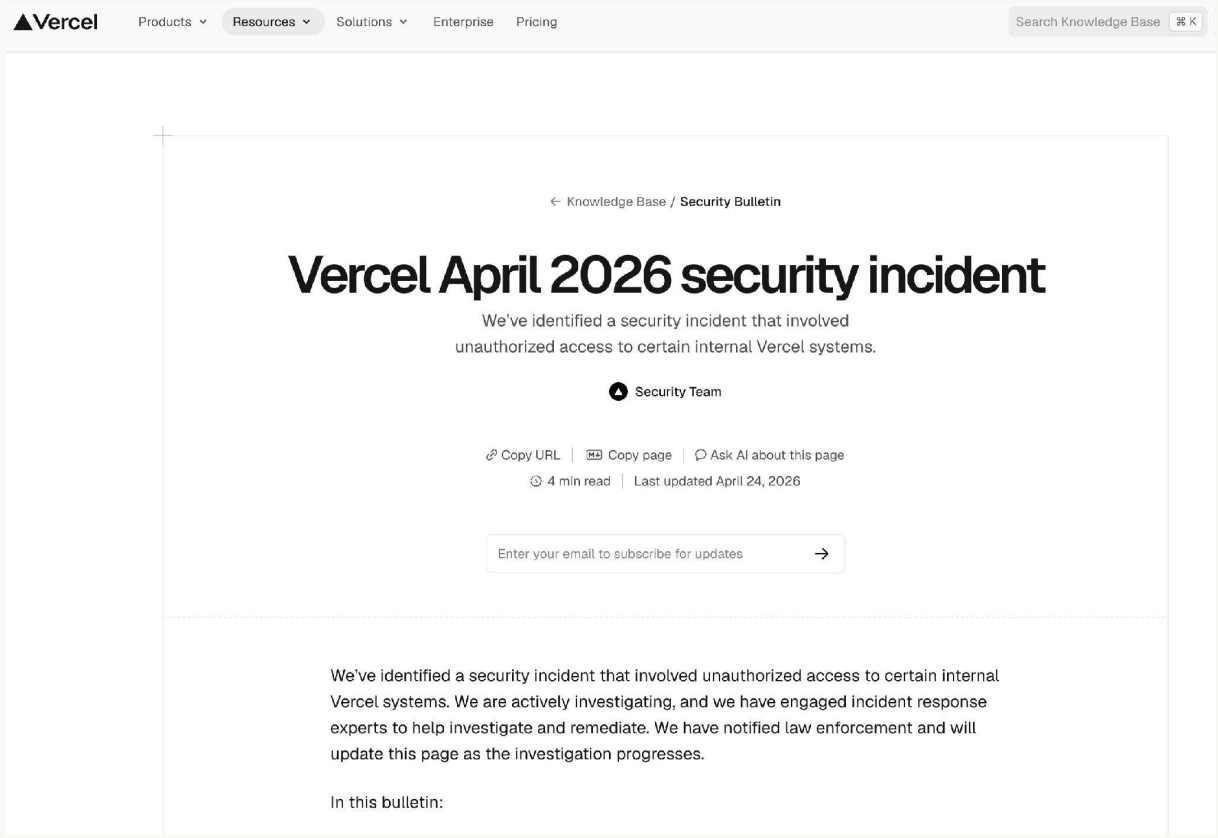

Employee gives Context.ai "Allow All" on Google Workspace

Infostealer at Context.ai exposes stored OAuth grants

Attacker rides the grant into Vercel → internal env → keys, tokens, credentials

Four runtime governance approaches

cakewalk

Protocol Gateway

NETWORK LAYER

Sees every tool call. Cloud + self-hosted.

SDK Instrumentation

AGENT RUNTIME

Deepest context, but per-framework and bypassable.

Kernel

OPERATING SYSTEM

A backstop for forbidden actions. Limited context.

Vendor App Integration

INSIDE EACH APP

Only where the vendor builds the hooks.

At a Glance

cakewalk

Same problem, four layers, different tradeoffs

	Protocol Gateway	SDK Instrumentation	Kernel	Vendor App Integration
Where it sits	Protocol layer	Runtime	Operating system	Vendor app
How hard to circumvent	High	Medium	High	Depends on vendor
Context awareness	Medium	High	Low	Depends on vendor
Integration effort	Low	Medium	None	Low (if available)
Self-hosted agents	✓	✓	✓	✓
Cloud agents	✓	×	×	✓
Meets the AARM standard	✓	✓	×	Only if vendor provides full hooks
Third-party dependency	Not needed	Required	Not needed	Required
Best for	*Any agent that connects to tools via MCP or standard protocols*	*Self-hosted agents where you control the code*	*Backstop layer for forbidden actions*	*Cloud agents where the vendor provides governance hooks*

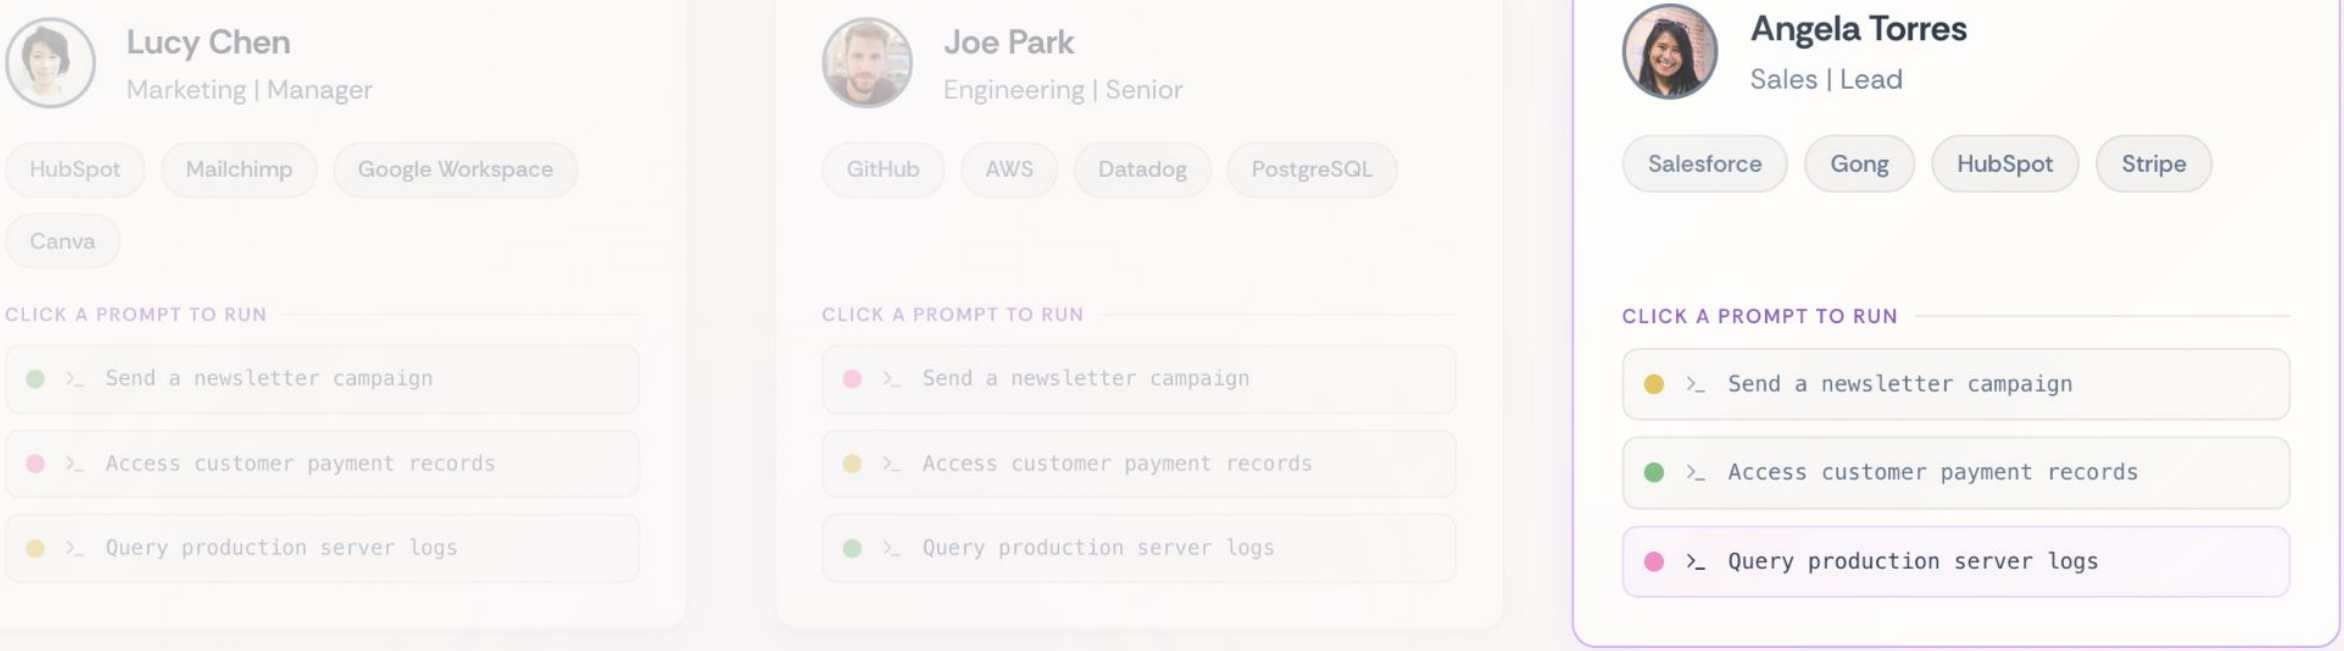
Lucy Chen
Marketing | Manager
HubSpot
Mailchimp
Google Workspace
Canva
CLICK A PROMPT TO RUN
>_ Send a newsletter campaign
>_ Access customer payment records
>_ Query production server logs
Joe Park
Engineering | Senior
GitHub
AWS
Datadog
PostgreSQL
CLICK A PROMPT TO RUN
>_ Send a newsletter campaign
>_ Access customer payment records
>_ Query production server logs
Angela Torres
Sales | Lead
Salesforce
Gong
HubSpot
Stripe
CLICK A PROMPT TO RUN
>_ Send a newsletter campaign
>_ Access customer payment records
>_ Query production server logs

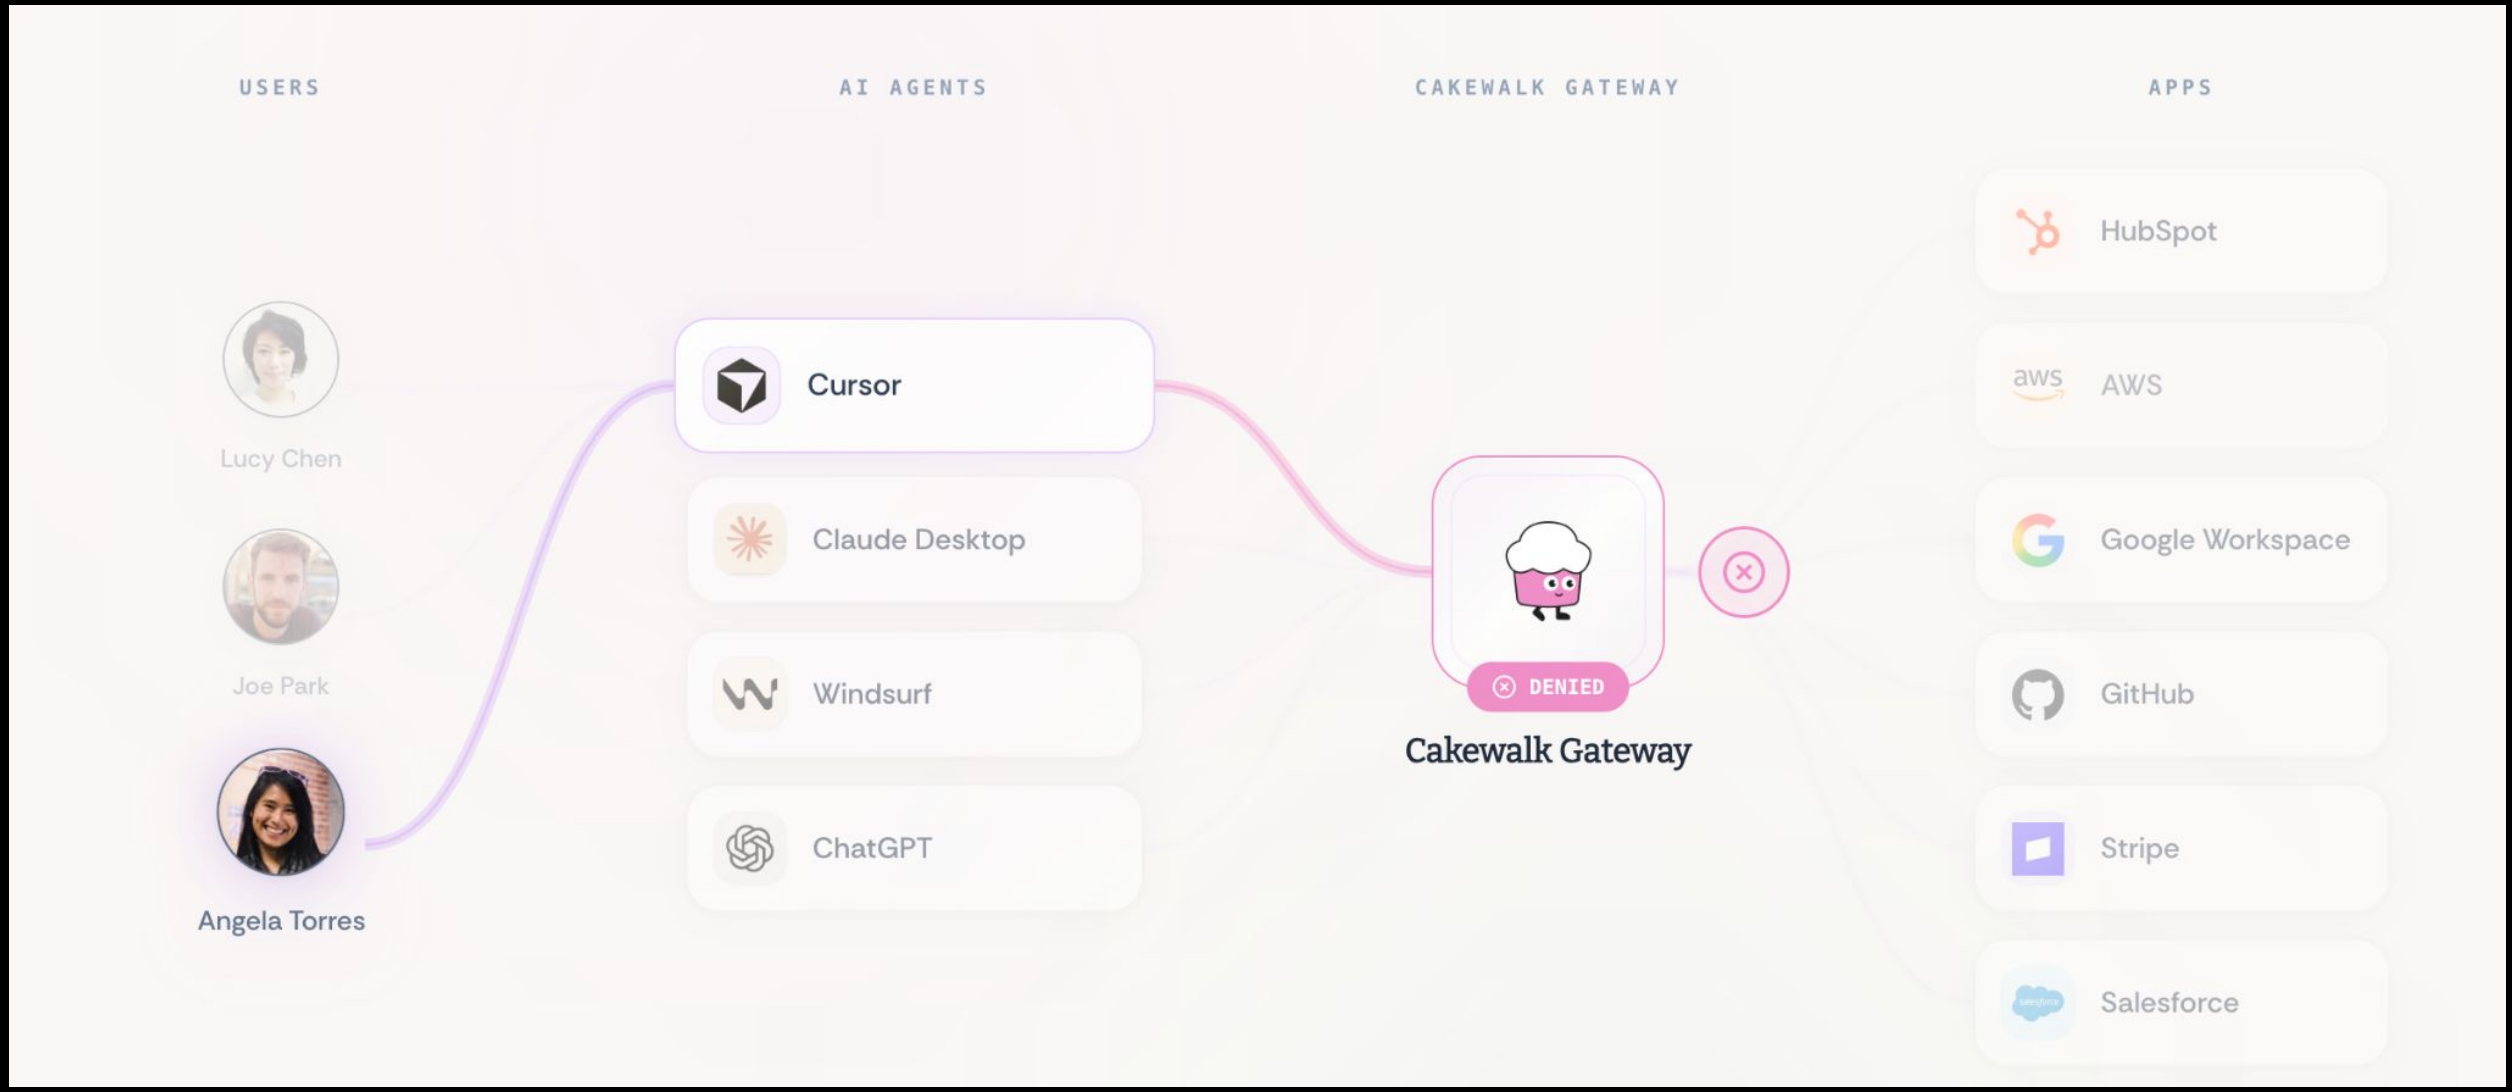
USERS
AI AGENTS
CAKEWALK GATEWAY
APPS
Lucy Chen
Joe Park
Angela Torres
Cursor
Claude Desktop
Windsurf
ChatGPT
DENIED
Cakewalk Gateway
HubSpot
AWS
Google Workspace
GitHub
Stripe
Salesforce


policy_evaluation_engine
ACTION TYPE
READ
USER CONTEXT
Angela Torres (Sales / Lead)
APP CONTEXT
AWS | Production | High Risk
POLICY MATCH
“Sales has no access to production infrastructure.”
OUTCOME
DENY

The policy layer between agents & apps

cakewalk

Policy Engine

Per-action rules in real time. Reads approve, writes escalate, destructive deny.

Credential Mediation

Agents never see real tokens. Held in the vault, injected per call.

Least Privilege

Just-in-time access, scoped to the task, revoked when the work is done.

Identity-Linked Lifecycle

Offboard the human and the agents go too. Every call traces to a person.

Agent Discovery

Every agent, one catalog, one kill-switch. Shadow AI becomes sanctioned AI.

Works across cloud, local, self-hosted and multi-agent setups - one gateway for all of them.

Governance and autonomy.

Not either / or.