ENTERPRISE-GRADE AI GOVERNANCE INFRASTRUCTURE

Govern your AI. Own your infrastructure.

The rules have changed. In April 2026 Anthropic released Claude Cowork with three extension points that let enterprises redirect the client to their own inference, identity and audit infrastructure. Cowork calls no longer have to touch Anthropic's cloud. systemprompt.io is the single Rust binary that takes the other side of that contract: govern every prompt, tool call and MCP execution, and own the gateway, plugin catalogue and audit trail they run through.

AT A GLANCE

The complete governance stack.

1 RUST BINARY · POSTGRES ONLY	100% PROMPTS, TOOLS, MCP AUDITED
4 ENFORCEMENT LAYERS PER CALL	0 OUTBOUND CALLS REQUIRED

01 Core capabilities

> GOVERN

Secrets & tool-use firewall

Every prompt and tool call is scanned and audited before it reaches inference. Secrets stripped, scopes enforced, quotas checked, denials logged.

Pre-inference audit

> PROVE

Identity-bound audit

Every request, tool call, and MCP execution written to `audit_events`, linked by `trace_id` to a JWT-verified identity.

SIEM-ready JSON

> ROUTE

Any-provider gateway

Point Cowork at `/v1/messages` you operate. Route per call to Anthropic direct, Bedrock, Vertex, Azure, OpenAI, or on-prem.

Your negotiated rates

> SCOPE

Signed plugin catalogue

Ed25519-signed manifest. Per-user, per-role resolution. Atomic stage-then-rename. MDM-distributable. Revoke with one row.

Ed25519 · RFC 8032

02 Why enterprises choose systemprompt.io

→ **Your perimeter stays intact**

Prompts, completions, and audit trails live in a PostgreSQL you own. The upstream model sees prompt content only; never your users, sessions, or tenants. Share one Bedrock or Vertex account across thousands of users without leaking organizational structure.

→ **Revocation without collateral damage**

Revoke a user's AI access at the gateway; effective within one TTL window, configurable to 300 seconds. The user keeps mail, chat, code repos, deploys. HR procedures (leavers, role changes, incident holds) become usable for AI access.

→ **Supply chain under your signing key**

Ed25519 signing key held in your HSM/KMS. Every plugin, skill, and MCP server passes signature verification before a byte lands on disk. Per-user manifest resolution means finance and platform see materially different catalogues.

BINARY	Single Rust executable. `PostgreSQL` 14+ only dependency. Signed releases, SHA-256 verifiable.
CLIENT CONFIG	Five managed prefs sealed via Intune · Jamf · Group Policy: `inferenceProvider`, `inferenceGatewayBaseUrl`, `inferenceCredentialHelper`, `...TtlSec`, `...AuthScheme`.
IDENTITY	JWT (audience + issuer). Helper walks mTLS → session+OAuth/PKCE → PAT ladder. Token in OS keystore (Keychain, Credential Manager, Secret Service).
DEPLOYMENT	Docker, bare metal, systemd, air-gapped. Zero outbound calls required. Sovereign-cloud ready.
OBSERVABILITY	OpenTelemetry, structured JSON export to Splunk, ELK, Datadog, Sumo Logic.
COMPLIANCE	SOC 2 (CC6.1, CC7.2), ISO 27001 (A.5.15, A.8.3, A.8.25), HIPAA §164.312, EU AI Act Art. 9/12/15, NIST AI RMF, OWASP Agentic Top 10.
LICENCE	BSL-1.1, source-available. Free to evaluate. Commercial for production. Perpetual licence & SLAs available.

PROVIDER-AGNOSTIC · WORKS WITH

Agents: Claude · Codex · Gemini · custom
Cloud: Bedrock · Vertex AI · Azure Foundry
Direct: Anthropic · OpenAI · Gemini · Groq
Self-hosted: Llama · Qwen · any `/v1/messages`

Evaluate free. Production-ready in days, not months.

Clone the template, run locally, and evaluate the full governance pipeline. No time limit.

`github.com/systempromptio/systemprompt-template`

`ed@systemprompt.io`

03 Three zones, one boundary: the deployment architecture

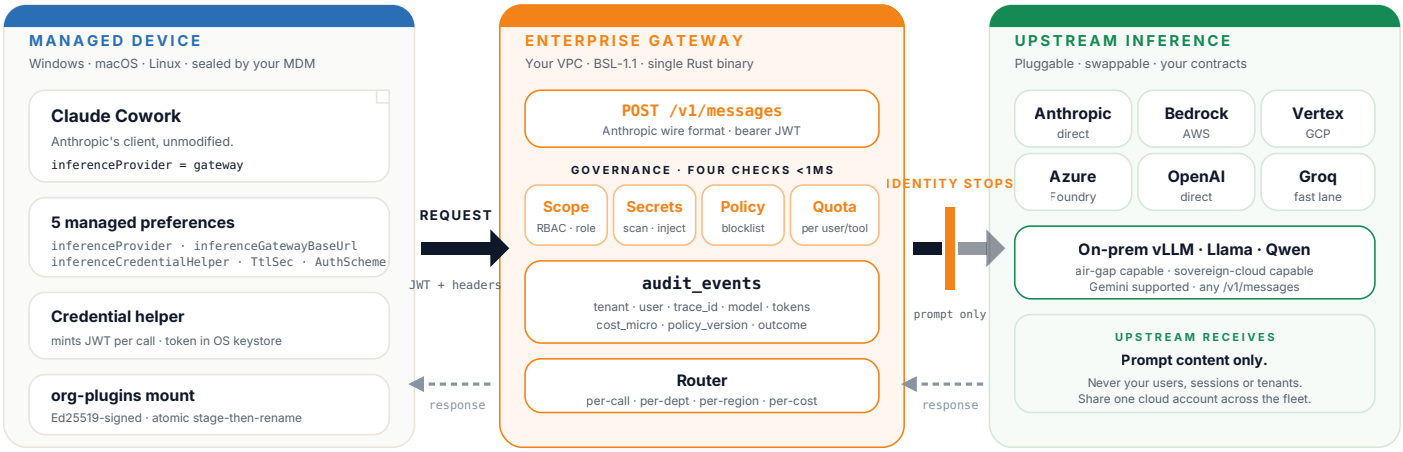

THREE ZONES · ONE AUDIT LOG · ONE BOUNDARY FOR IDENTITY

IDENTITY STOPS AT THE BOUNDARY The gateway strips the seven identity headers before forwarding. The upstream sees prompt content only, which lets you share one Bedrock or Vertex account across thousands of users while preserving per-user audit inside your perimeter.

04 Every Cowork capability, on infrastructure you own

	Claude Enterprise	Claude Custom (cloud)	Claude Custom + systemprompt.io
DATA, AUDIT & DEPLOYMENT			
Data residency	Anthropic US infrastructure	Cloud provider region	Your datacenter, your jurisdiction, air-gap capable
Identity-bound audit	Anthropic-held	OpenTelemetry only	Prompt → tool → MCP → cost, in your DB, by `trace_id`
Air-gapped deployment	Not available	Not available	Single binary, zero outbound calls
IDENTITY, ROUTING & COST			
Revoke a user	Remove from seat / SSO	Remove from cloud IAM	One DB row. SSO untouched. < 300 s.
Inference provider	Anthropic only	Bedrock, Vertex, Azure (Claude only)	Any, per-call routing, your negotiated rates

05 Built for the teams that answer to auditors

FOR THE CISO

Pass the audit in one query.

Full lineage from request to tool call to MCP execution to cost, linked by `trace_id` to a JWT-verified identity. Structured evidence, not policy documents.

Outcome: Auditor asks, you run a SQL query.

FOR THE CTO

Stop building a governance team.

Three teams, narrow interfaces. AI platform owns routing, security owns the signing key and allowlist, endpoint owns the MDM profile. The binary absorbs every vendor-side churn.

Outcome: Own the deployment, not the maintenance.

Read the full enterprise deployment guide. The complete playbook for rolling out Claude Cowork on self-hosted infrastructure.

Architecture, three-team operating model, compliance mappings, MDM profiles, audit schema, revocation runbooks.

`systemprompt.io/guides/claude-cowork-plugins-enterprise`