C1.ai
A Workday Company

SOLUTION BRIEF

C1 LLM Gateway

Route every model call through one governed endpoint.



ramp · instacart · zscaler · DigitalOcean · Brex · klaviyo · qualtrics XM

Uncontrolled model spend often starts with a shared model key.

A team connects its first application to a model provider. Then another application needs a model, and it goes to the same provider. Before long, internal tools, customer-facing products, and AI agents all send their requests to the same place, whatever the request contains or what it costs.

At that point, basic questions become surprisingly difficult to answer. Who or what made this call, and what did it send? Is the data in it allowed to leave your environment? And which team should own the cost?

An agent pulling finance data and a routine email summary take the same path, even though one belongs within your own cloud boundary, like AWS Bedrock in your account, and the other could go to an approved lower-cost provider. The invoice arrives later as one large number, with no clear way to tie spend to the teams, applications, or agents behind it. What began as a shortcut is now part of the company's infrastructure.

C1 LLM Gateway gives every application and agent one governed path to approved models, so every request goes where it should, based on who is asking, what it contains, and what it costs.

Route inference without hard-wiring the provider.

One endpoint for supported providers.

First, connect apps and agents to one C1.ai policy-controlled endpoint across approved public, private, and customer-controlled model deployments.

Policy-based smart routing.

Second, auto-route inference based on identity, workflow, data-handling requirements, region, and cost, giving organizations control over where the data goes.

Cost accounting.

Third, route inference based on workload needs to manage cost and produce cost accounting and reporting.

Identity is the starting point.

Many gateways rely on an application to identify itself with a tag. But a tag is only as trustworthy as the application supplying it. C1.ai anchors every request to the identity of the person, application, or agent that actually made it.

That creates a consistent link between model access, usage, and spend. The same identity layer a company already uses to decide who can access an application can also determine which models they can reach, what they are allowed to spend, and when that access should end.

One gateway for every approved model.

C1 LLM Gateway gives tools and AI agents a single endpoint for accessing public, private, and customer-managed models. Instead of hard-coding a provider into every application, teams define the routing decision as policy.

Routing also accounts for which model handles the request and who hosts it. The same model can carry different data-handling implications depending on where and how it is deployed. C1.ai gives security and platform teams a consistent place to govern both. Sensitive workloads can stay on infrastructure the company controls, while lower-risk requests can use the provider best suited to the job.

Use one endpoint across providers.
Add or change models without rewriting every application that depends on them.

Route by sensitivity.
Keep designated data on approved infrastructure and block it from providers that should not receive it.

Route by performance and cost.
Use policy to select the right approved option for each workload rather than relying on one default.

Reduce dependence on a single provider.
Applications integrate with the gateway instead of becoming tightly coupled to one model vendor.

Make model spend attributable.

C1 LLM Gateway attributes usage per application and per user. Platform and finance teams can see where spend is coming from and act. Budgets can follow the same request-and-approval process as other governed access. When an application needs more capacity, the owner can review the request, approve an appropriate amount, and time-box the grant. Alerts can surface a potential overage before it happens, and access can be revoked when it is no longer needed.

See who is spending what.
Connect model usage to the relevant person, team, application, or agent.

Set budgets around real owners.
Give application owners a clear role in reviewing and approving additional capacity.

Respond before an overage.
Alert the right people as usage approaches its limit instead of waiting for the provider invoice.

Keep a useful audit trail.
Record who made each call, which model handled it, and what it cost.

You decide what leaves your environment.

C1 LLM Gateway gives teams a common control point for deciding which models can see which data. It also creates a record of the identity, routing decision, usage, and cost behind every request.

The result is a simpler path for developers and a more governable one for security, platform, and finance teams: one endpoint to build against, clear policies for where requests go, centralized provider credentials, and accountability for every model call.

GET STARTED

Talk to a product expert

Book a demo
c1.ai/demo

About C1.ai

C1.ai
A Workday Company

C1.ai, a Workday company, is the identity platform for the AI era. Our platform connects humans and agents to enterprise systems with scoped access, enforced policies, human approval, and a record of every action. With C1.ai, the fastest path to AI adoption at scale is the governed path — giving organizations visibility and control without sacrificing security or slowing innovation. Companies such as Ramp, Zscaler, Qualtrics, and DoorDash trust C1.ai to govern identity, secure agents, and empower their teams to put AI to work.

Put AI to work.

Copyright © 2026 ConductorOne, LLC. All rights reserved.