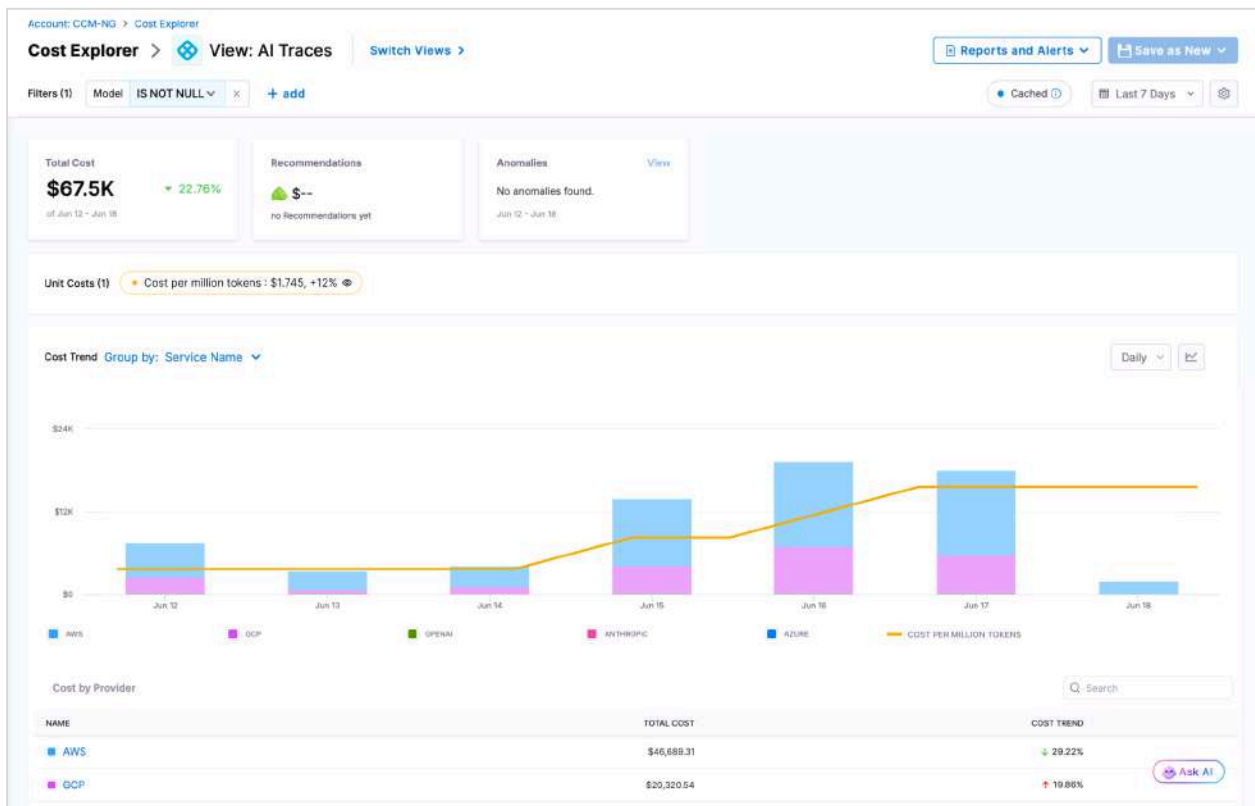


AI Cost Management and Unit Cost Economics on a reimagined Cost Explorer

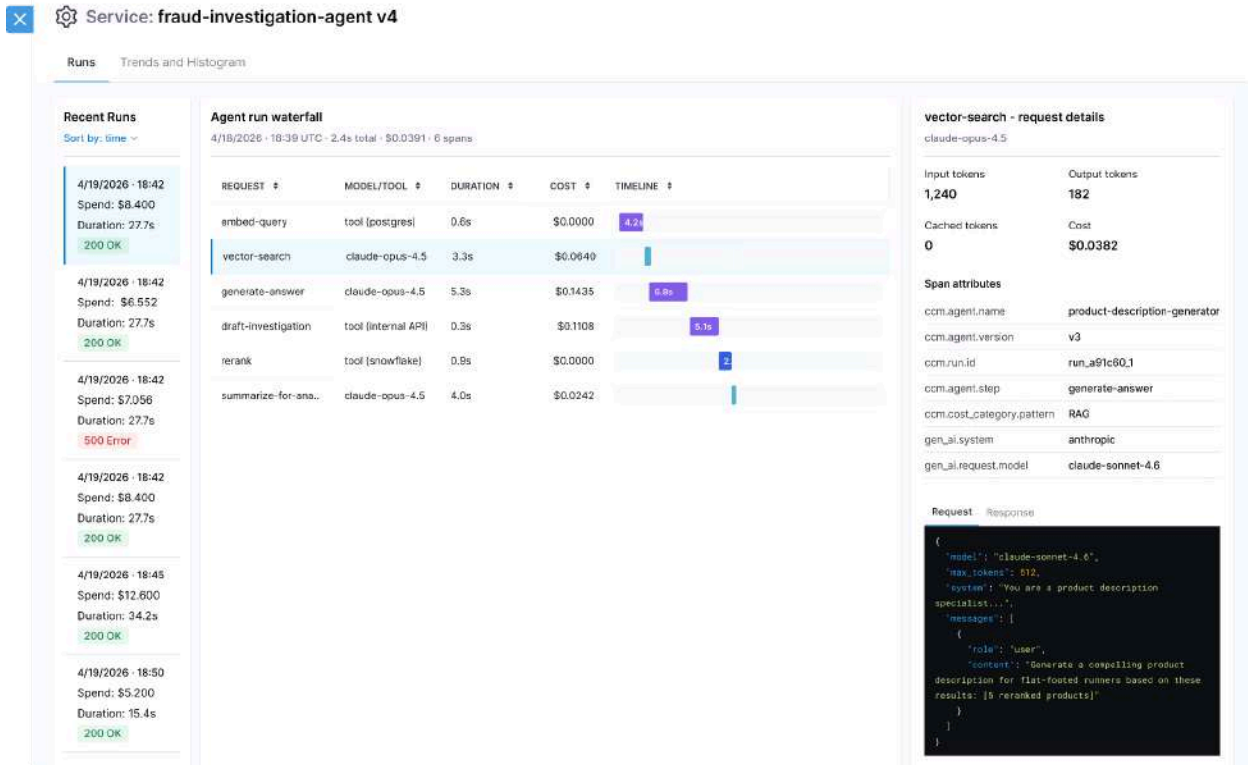
AI spend is the fastest-growing line on most cloud bills, and visibility hasn't kept up. Cost Explorer has been rebuilt to bring AI costs, and all cloud spend, into the units your business actually cares about: cost per token, per customer, per developer, per dollar of revenue. One place to go from "what did we spend?" to "what does it cost us per unit of value?" and straight into why it changed.

*A note for existing users: **Views** are the evolution of **Perspectives**. Everything you built carries over automatically. Want to see what's different, or need to switch back? See "For existing Perspectives users" at the end of this document.*

AI costs are climbing fast, and visibility hasn't kept up



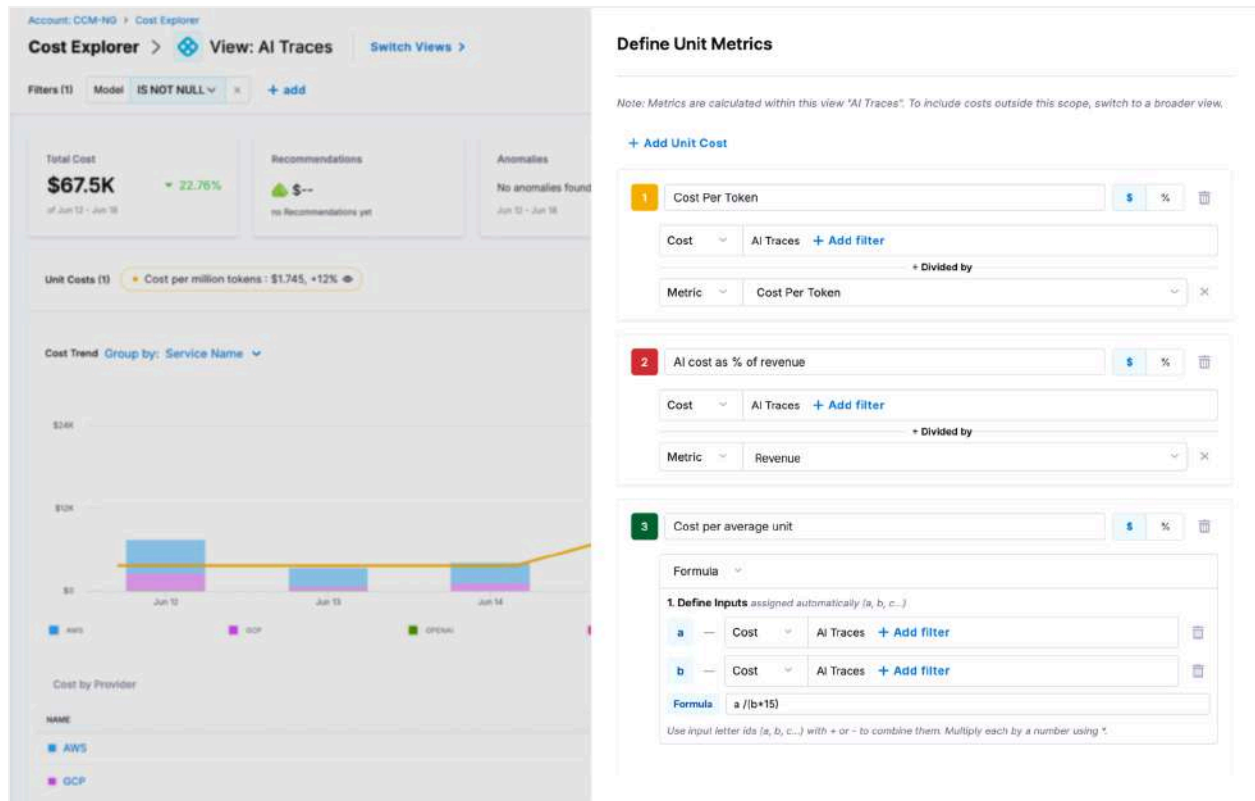
Uber burned through its entire 2026 AI coding tools budget in four months. They're not alone. AI spend is rising fast, and most teams have little visibility into where it goes.



Cost Explorer now includes AI cost management, complete with unit cost economics. Two Views ship out-of-the-box:

- GenAI Costs.** Your source of truth for Gen AI billing. Pulls spend directly from AI and Cloud providers (Anthropic, OpenAI, GCP Vertex, AWS Bedrock and others) through connectors, with budgets and anomaly alerts included. Cost per million tokens is the OOTB unit cost; add your own metrics to track cost per session, agent, run, inference or any unit that matters.
- AI Traces** (*coming soon*). For granular dimensional drilldowns into how AI is actually being used. Data flows from your own applications via OpenTelemetry instrumentation: small additions to your AI call paths that capture traces and metrics. Because it taps the application layer rather than the vendor bill, it covers every AI provider your code touches in one View. Investigate cost by model, agent, session, environment, or any custom dimension. For billing reconciliation, use GenAI Costs View.

Introducing Unit Cost Economics



The AI Views above use unit costs (cost per token, cost per session). That same capability is available for all of your cloud and AI spend.

Raw cloud spend tells you how much. Unit cost tells you how much per unit of value. This new capability brings your business metrics (developers, customers, orders, revenue, etc.) directly into Cost Explorer, as a trend line alongside the cost reports you already use. All this in just 2 steps:

- **Scope the cost.** Start from a View or filter down to the costs that matter: non-production spend, just the checkout services, whatever's relevant. You're working with costs you already understand.
- **Add a business metric.** Bring in your metric (total developers, active customers, orders, etc.) and Cost Explorer divides automatically. You can build custom formulae across multiple costs and metrics and a trend line appears over your cost report.

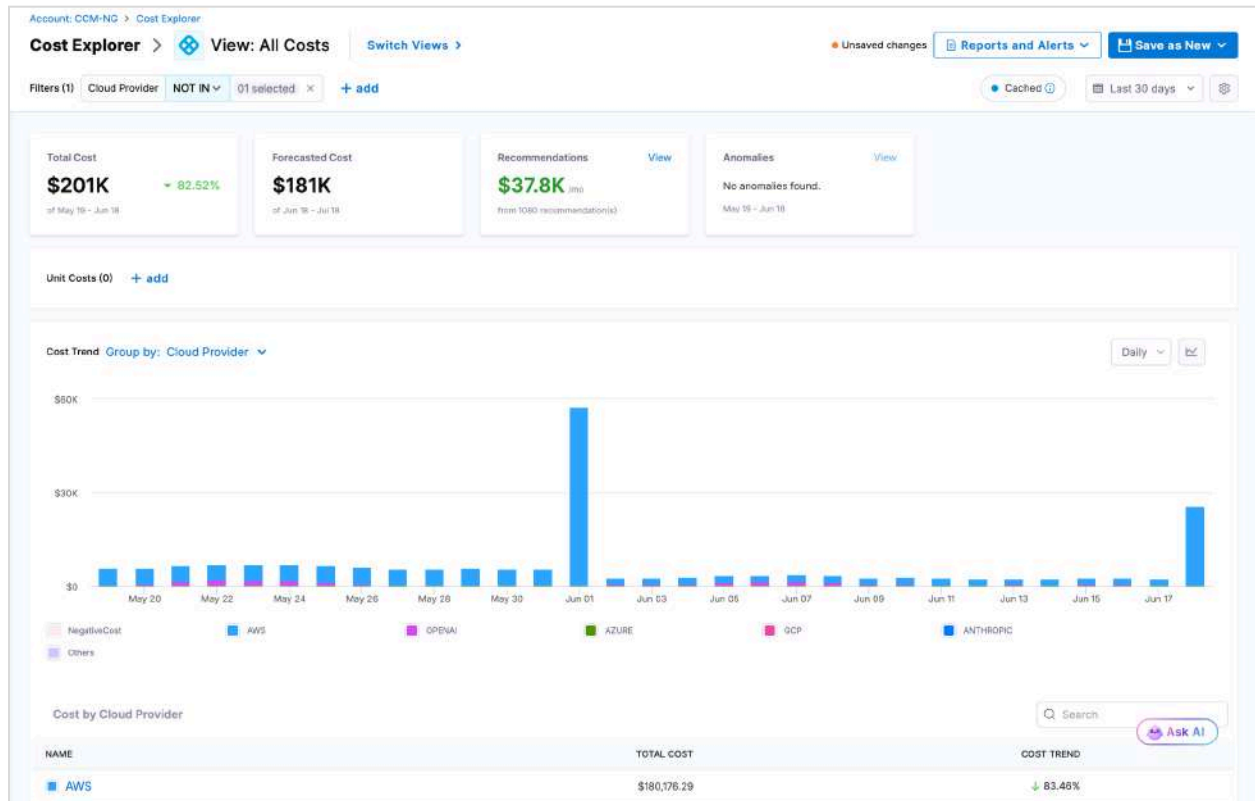
Compare across teams or products by splitting a View into buckets via Cost Categories and showing a unit cost line for each. An SVP can watch every team's efficiency side by side.

Always know what you're dividing. Each unit cost displays the exact cost scope behind it, so the number is never a mystery and is safe to put in a board deck.

To learn how to ingest your business metrics and get started with Unit Cost Economics, see [the Unit Costs documentation](#).

The foundation: What's new in Cost Explorer

Everything above is built on a rebuilt Cost Explorer.



Save what you see, no rule builder. Slice and dice your costs, then save the result as a named View. There's no rule-building step. You save exactly the analysis in front of you, and reopening it restores the full setup: filters, group-by, time period, granularity, and preferences (including cost types and adjustments). Out-of-the-box Views ship ready to use, and any of them can be the starting point for a custom View.

Click to filter, drill down to the cause. Click any bar in the chart or row in the table to filter in place. Intelligent drilldown steps to the next logical level automatically (Service to Usage Type, for example), so root-cause analysis happens in clicks, not configuration.

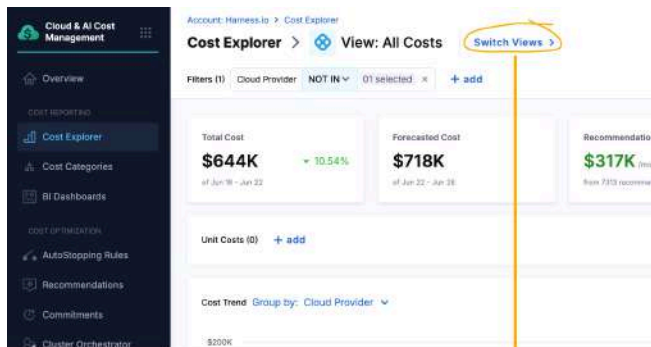
From spike to answer: built-in root-cause analysis

Because everything above lives inside Cost Explorer, investigating a cost spike is one continuous flow. When cost-per-unit jumps, keep the unit cost baseline fixed and drill into the underlying cost slices by team, service, or provider, all in the same context. From there, set an alert or save the investigation as a shareable View.

For existing Perspectives users

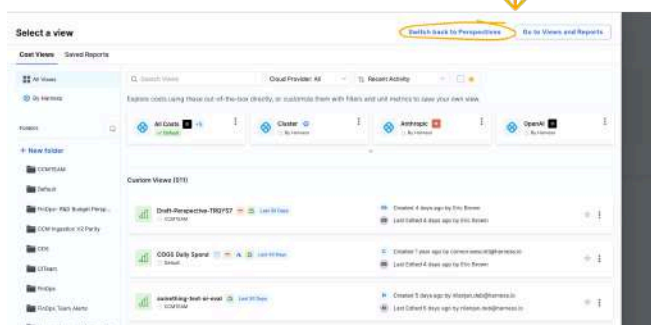
The new Cost Explorer is now the default for all accounts. All your Perspectives have been migrated to Views automatically, and nothing is lost. Here's what changes:

	Legacy (Perspectives)	New (Views)
Creation	Rule builder: construct AND/OR conditions up front	Save what you see: filter, group, save
Unit Costs	<i>Not available</i>	Cost per developer, per customer, per token, with custom formulae
AI Cost Management	<i>Not available</i>	GenAI Costs and AI Traces Views, out of the box
Drilldown	Manual re-filtering	Click to filter, intelligent hierarchical drilldown
Preferences	Set separately	Saved with the View, so everyone sees the same numbers



If you would like to switch back:

Once the new Cost Explorer is enabled for your account, it becomes the default. All your existing Perspectives are automatically migrated to Views, and nothing is lost.



If you'd like to switch back to the classic Perspectives experience, open the View selector and choose "Switch back to Perspectives" at the top right of the panel. You can toggle freely between the two experiences at any time.