

What you are billed for

A common and reasonable question: does routing through Sentilai make AI usage cost more?

Sentilai charges per seat band

Your Sentilai subscription is a flat price for a **band** of seats — the shared pool of admins and developers. It does not vary with how many requests your team makes, and it does not vary with how much of the band you use: eleven people and twenty-five people on the same band pay the same.

AI usage is billed separately, and by whom depends on the mode

Subscription mode — the developer's own AI plan pays, exactly as it did before Sentilai existed. Nothing changes about that bill.

Managed API key mode — the request goes out on your organization's provider key and the provider bills you directly, at the provider's own rates. Sentilai does not mark this up and does not resell AI capacity.

The **Mode** column in Activity tells you which applies to any given request, and **Compliance** reports the overall split.

Controlling managed spend

- **Enabled models** on each provider key: the simplest lever. A key that does not enable the expensive model cannot be used to buy it.
- The **Tokens** column in Activity finds the agent loop that is burning far more than anyone expected.
- Your provider's own dashboard remains the authority on what you owe them.

What the classifier costs

The prompt-injection classifier we host ourselves is included. The larger EU-hosted classifier tier is part of your subscription rather than a separate meter — if you are sizing an evaluation and want the details, ask.