SecuriX

AI GATEWAY · COST CONTROL

THE 151ST SEAT

Seat 151 shouldn't rewrite your AI budget.

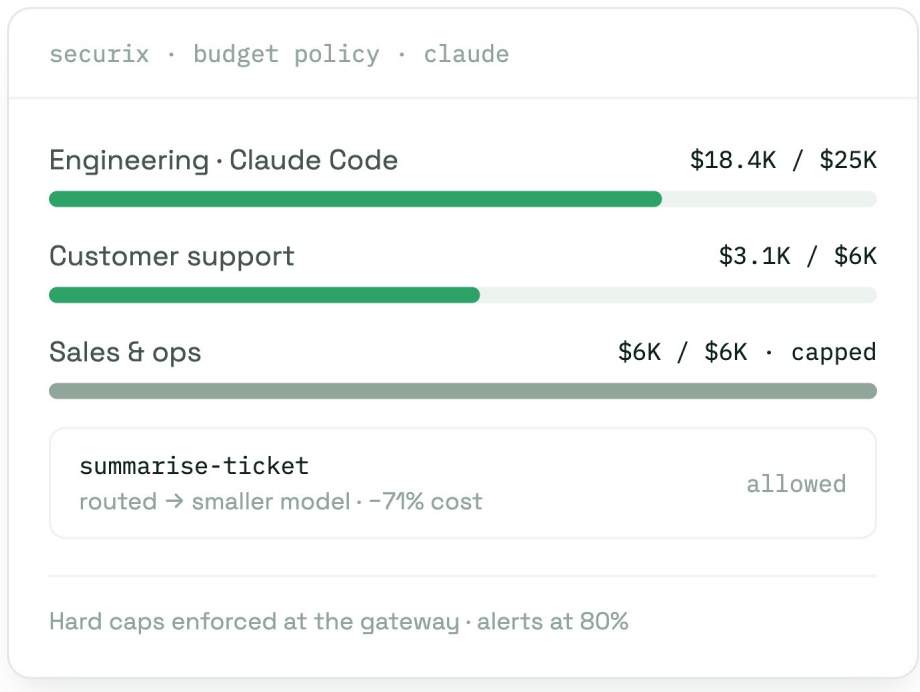
Claude's Team plan tops out at 150 seats. On Enterprise, every token is billed at API rates on top of the seat. Put a gateway in front **before** you switch.

WITHOUT A GATEWAY

$36K → $1.36M / yr

One public estimate for 151 heavy users once usage is billed per token.

- Seat covers access; usage is billed on top
- Bill tracks each person — hard to forecast
- Heavy users can far outrun the seat fee



How SecuriX makes API billing predictable.

01

Hard spend caps

Budgets per team, department or key. The cap stops requests — not your quarter.

02

Right model, every prompt

Policy-aware routing and allowlists keep routine work off premium models.

03

Token-level attribution

Every request tied to a user, agent, model and tool — Claude Code included.

04

Pay per token, not per seat

One view of spend across every provider. Runs on-prem or in your VPC.

60–80% lower AI cost

vs. over-provisioned seats. Weighing the Enterprise jump? Talk to us first.

$495/month per deployment. No fee per employee, no markup on tokens.

Book a demo →

getsecurix.com