

Sticker price is \$60/seat. Real cost runs \$1.1M–\$2.5M annually for a 1,000-person deployment once API, custom GPTs, shadow accounts, and governance overhead are counted. These controls keep the gap closed.

▲ THE TRUE COST STACK — WHAT "\$60/SEAT" ACTUALLY MEANS

Seats	API Usage	Custom GPTs	Shadow Accts	Gov./Integ.
<p>Seats are the contract floor. API usage on top of seats often matches the seat spend within 12 months. Custom GPTs and connectors drive token consumption invisibly. Shadow accounts (Plus, Pro, Team) layer on additional spend. Governance sits outside the OpenAI invoice but is real money.</p>				

150

seat minimum for Enterprise

\$30

per M output tokens, GPT-5.5

6×

output vs. input token cost

\$670K

avg. shadow-AI breach cost

1 CENTRALIZE PROCUREMENT

Stop paying for four parallel tiers

Marketing has Plus. Engineering has Pro. Finance has Team. The interns are on Go. Pull every ChatGPT subscription into one inventory — the sum is usually more than one Enterprise contract.

- ▶ Audit corporate cards, expense reports, SSO sign-ins
- ▶ Consolidate under one Enterprise contract for volume pricing
- ▶ Block personal-account access from corporate devices

2 CUSTOM GPT INVENTORY

The #1 governance gap most enterprises miss

Every custom GPT can connect to internal data and external services. Most enterprises don't know how many exist, who built them, or what they touch.

- ▶ List every GPT: owner, data sources, external actions
- ▶ Enable approved-domain lists and group permissions
- ▶ Require security review before new GPTs touch data

3 PIN YOUR API MODEL VERSIONS

Stop the silent upgrade tax

OpenAI changes default models. When GPT-5.5 launched, code using "default" doubled in cost overnight. Pin versions explicitly in every internal tool.

- ▶ Explicit model IDs in all API calls, no defaults
- ▶ Test new models in staging before promoting
- ▶ Use Batch API or Flex for non-urgent work (50% discount)

4 SET WORKSPACE ACTIVITY THRESHOLDS

"Some use last month" is not active

Define what active means: **one meaningful use per workday**. Workspace analytics is in the admin console — most companies have never opened it.

- ▶ Monthly review of per-user message and project counts
- ▶ Re-engage below-threshold users; reassign if no fit
- ▶ Treat power users as champions — they drive ROI

5 AUDIT YOUR CONNECTOR FOOTPRINT

Every connector is a data flow that needs governance

Google Drive, Slack, Teams, SharePoint, Gmail — each connector creates a token stream and an exposure surface. Many were enabled "to test" and never disabled.

- ▶ Inventory active connectors; disable unused ones
- ▶ Apply group-level permissions, not workspace-wide

6 DEPLOY PROMPT-LEVEL DLP

Scan what employees type, before it reaches OpenAI

PII, financial records, source code pasted into prompts is the most common security incident. SOC2 infrastructure doesn't cover this.

- ▶ Prompt-level DLP scanning for PII, PHI, source code
- ▶ Real-time block or alert for high-risk patterns

7 BASELINE OUTCOMES BEFORE THE BOARD ASKS

"People like it" is not a renewal argument

Capture workflow metrics **before** rollout. Re-measure at 60 days. Proposal cycle time, contract review duration, customer response time — outcomes that survive a CFO conversation. Self-reported productivity from Impact Surveys is directional only.

- ▶ Pick 3 measurable workflow outcomes per major user group
- ▶ Track total true cost (seats + API + governance) monthly
- ▶ Build the renewal slide before the renewal meeting

⚡ MONTHLY CHECKLIST FOR CHATGPT ENTERPRISE ADMINS

- ✓ Pull workspace analytics: messages, GPTs, projects per user
- ✓ Track API spend vs. seat spend trajectory month-over-month
- ✓ Reconcile assigned seats against active threshold (1/workday)
- ✓ Review custom GPT inventory — any new ones since last audit?
- ✓ Audit prompt-DLP alerts and incident response actions
- ✓ Update the CFO/CIO cost-and-outcome dashboard slide

"The sticker price is 30–50% of your true cost. The other half is API, shadow accounts, custom GPTs, and governance overhead. You don't need to be smarter than OpenAI to get this right — you just need to be more disciplined than the average buyer.