

Answer Key — The Personal / Corporate Boundary

Admin / self-check reference

Compare each item against your own written answer. The direction label matters as much as the verdict — being able to name the direction is the sign you have internalized the two-sided threat model.

Scenario 1 — Stack trace into personal ChatGPT

Verdict

Violation? Yes.

Direction: personal→corporate (corporate data flowing into a personal AI account).

Risk: Source paths, internal hostnames, and a live token have just left the corporate perimeter and entered a personal-tier provider with default retention and possible training eligibility. The company has no contract, no audit trail, and no way to recall the data. The token is the worst part: you have no reason to believe it is "redacted-looking enough" — it is a live credential now sitting in a third-party log pipeline.

Fix: Use the corporate enterprise AI tier for this kind of debugging, or ask a colleague. Regardless, *rotate the token immediately* — you now have to assume it is compromised, because from the company's perspective it is.

Scenario 2 — Friend's side project into corporate Claude

Verdict

Violation? Yes.

Direction: corporate→personal (external, personal-origin content flowing into a corporate AI account).

Risk: Several layered problems. IP assignment gets confused because anything produced with corporate tools can be pulled into the scope of your employment contract. The friend's code is now in corporate retention and audit logs forever. Audit tooling will see corporate resources being used on code that does not match any corporate project. And it is a straightforward acceptable-use policy violation, which can carry disciplinary consequences.

Fix: Switch to a personal account for personal help. Better still, help your friend over a call or a screen share where nothing gets pasted into any tool at all — that keeps both jurisdictions clean.

Scenario 3 — AI email summarizer in a shared browser profile

Verdict

Violation? Yes.

Direction: both, simultaneously — the extension unifies the blast radius. It is the canonical linked-account trap.

Risk: Anything the extension can read in corporate Gmail flows out to the extension provider, which is a personal-tier vendor with no corporate contract. At the same time, anything in personal Gmail is now mixed into the corporate browser-profile context and may end up in corporate logs or get correlated with corporate activity. The extension does not maintain a wall between the two tabs — there is no such wall. The union of what it can see is its blast radius.

Fix: Never install unvetted AI extensions, especially in a profile that is signed into corporate services. Use separate browser profiles per jurisdiction — one for corporate, one for personal — and treat corporate profile extensions as a list the company approves. Use only enterprise-approved tools for corporate work.

Scenario 4 — Corporate enterprise tier on an open-source side contribution

Verdict

Violation? Yes, subtle.

Direction: corporate→personal (corporate AI account used on personal work).

Risk: This is the case that trips up careful people, because "zero-retention" feels like a complete answer. It isn't. Zero-retention is a promise about what the *provider* stores, not about what your *employer* sees. Corporate admin logs still show your prompts and usage patterns. IP-assignment clauses in your employment contract can cloud the ownership of your open-source contributions if they were produced with corporate tooling. And the enterprise contract almost certainly covers corporate work only, meaning the "side project" use is outside its scope and therefore a policy violation regardless of retention posture.

Fix: Use a personal-tier AI account for open-source contribution work. The zero-retention corporate tier is not a universal-purpose tool; it is a corporate tool with a corporate scope.