

Agent Spend Limit Policy

Sample governance artifact · ghostcorpnet Pack Standard v1 starter kit

1. Purpose

This policy caps what AI agents may spend — per agent, per run, and per billing period — so no run can burn budget unattended.

2. Scope

Applies to all agents operated by [ORG NAME] that call paid model APIs, tools, or third-party services, including agents run by contractors.

3. Limits

Per-agent lifetime cap	[AMOUNT]	billing system / API gateway
Per-run cap	[AMOUNT]	orchestrator / runner
Per-day org cap	[AMOUNT]	finops dashboard
Single tool-call cap	[AMOUNT]	tool gateway

4. Approval gates

- Any run projected to exceed [X]% of its per-run cap requires human approval before starting.
- Raising any cap requires sign-off from [ROLE, e.g. Eng Lead] and is logged in the audit trail.
- Emergency override: [ROLE] may lift caps for [N] hours during an incident; the override is recorded and reviewed within [N] business days.

5. Alerts

- [X]% of any cap → warning to [CHANNEL].
- 100% of any cap → agent run pauses; page [ON-CALL].
- See the cost-alert runbook for the response procedure.

6. Audit

Every spend event is logged with: timestamp, agent identity, run id, tool called, cost, and remaining

budget. Logs are retained [N] days and reviewed [weekly/monthly] by [ROLE].

7. Exceptions

Exceptions are documented, time-boxed, and expire automatically. No standing exceptions.

How to use this template

Fill in the [BRACKETED] values for your organization. This PDF is the governance artifact of the starter pack — the Markdown source ships alongside it so you can adapt it freely under the MIT license.