



BareBear
OPEN SOURCE AGENTIC FRAMEWORK

Every step leaves a print.

A minimal, policy-first Python agent framework.

Most agents fail the same way in **production**



Hidden state no one can inspect



Tool calls that fire without gates



Unbounded loops with no kill switch



No answer to "why did it do that?"

Make agent behaviour governable



01 Explicit state

02 Policy before action

03 Bounded autonomy

04 Run receipts



7 primitives. Nothing extra.

Bear	The agent runtime
Task	A unit of work
State	Explicit, diffable, snapshotable
Tool	Gated callable with risk metadata
Policy	Pre-execution constraint set
Checkpoint	Human approval gate
Report	Full run receipt

Policy checks **before** execution

Because you can't un-send an email.

 **Tool call** requested



 Blocked? Side effects? Approval?



✓ Allow
Execute

✕ Block
Skip


⏸ Pause
Checkpoint

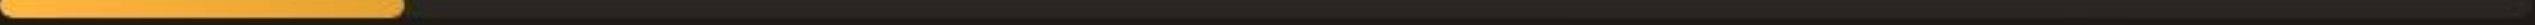



An agent without limits is a while- True loop with a credit card


Steps 7 / 10




Tokens 8,200 / 50,000



Cost \$0.04 / \$0.10



send_email → approval required



Every run is a receipt



```
task_id a1b2c3d4 status completed steps 5 tokens 8,200 cost $0.0340 duration 1.34s ----- assumptions "Tier
assumed from email domain" uncertainties "Could not verify account status"
```



Different tools, different jobs

No orchestration theatre here.

Category	Optimised for
Orchestration-first e.g. LangChain	Breadth & composability
Vendor-native stacks e.g. OpenAI Agents	Fast onboarding, ecosystem lock-in
Autonomous systems e.g. OpenClaw	High autonomy, broad tasking
BareBear	Explicit, policy-first, auditable



Choose BareBear if...

- ✓ You need policy controls before side effects
- ✓ You need explicit state you can inspect and diff
- ✓ You need hard runtime budgets
- ✓ You need approval checkpoints for high-risk actions
- ✓ You need report-grade traces for audit



BareBear
OPEN SOURCE AGENTIC FRAMEWORK

Built for agents you can **explain after they run.**

MIT licensed · Python 3.9+ · 83 tests
github.com/richey-malhotra/barebear

What would you need before trusting an agent in production?

