StackGen

AI ENGINEER WORLD'S FAIR 2026 EDITION

Aiden Agentic OS Cheat Sheet

What every AI engineer needs to know about running agents on real infrastructure.

BRING YOUR OWN LLM & AGENTS | Claude · GPT-5 · Gemini · Claude Code · Codex · Gemini CLI · CUSTOM

ALL ARE WELCOME.

1. The Problem

The model is not the bottleneck.
The infrastructure underneath it is.

Demo Shows:

- ✓ Agent completes task successfully
- ✗ One LLM, one harness
- ✗ Governance in the system prompt
- ✗ Agent provisions environment in seconds
- ✗ War room done in 90 minutes

2. Mental Model

An OS doesn't care which app runs on it.
An Agentic OS doesn't care which LLM or agent runs on it.

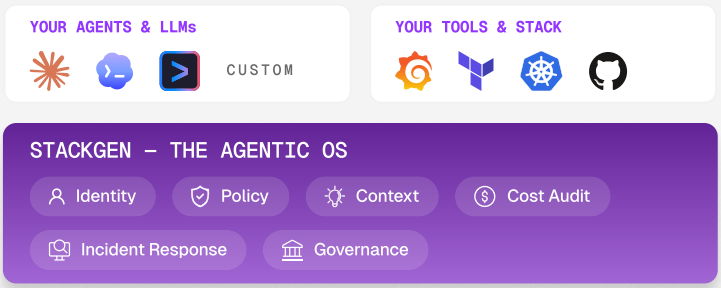

StackGen is not in the model business.
StackGen is not in the agent framework business.
Aiden from StackGen is the OS. **Everything else runs on it.**

3. The Three Failure Modes — and the fix

FAILURE MODE 1

The 90-min war room

- ✗ Alert fires, 4 engineers join with.
- ✗ Different context.
- ✗ 45 mins of log spelunking.
- ✗ Root cause maybe found. Recurs.

FIX → **Aiden for SRE**

Alert fires.

- ✓ Aiden pulls metrics, logs, topology
- ✓ RCA in Slack before first engineer
- ✓ Optional: remediation PR opened
- ✓ Optional: deployed after approval

Works: Grafana, Datadog, PagerDuty
LLM: yours. Free for 1+1 users → ai.stackgen.com/app/sre

FAILURE MODE 2

Agents + Infra = No Guardrails

- ✗ Claude Code spins up an environment.
- ✗ Wrong billing mode. No policy ran.
- ✗ You find out from the cloud bill.
- ✗ Or worse: from the security audit.

FIX → **Aiden for Infrastructure**

Agent requests infrastructure.

- ✓ AppStack golden-path template
- ✓ Governance at PLAN TIME ← key
- ✓ Non-compliant: BLOCKED
- ✓ Compliant: deployed

Rule: governance after apply = postmortem, not policy.

FAILURE MODE 3

Agent sprawl across SDLC

- ✗ Claude Code + Codex + Gemini + ops agent.
- ✗ 4 identity models. 4 token bills.
- ✗ 0 shared context.
- ✗ 0 audit trail. Compliance team calling.

FIX → **Aiden for DevOps**

One plane for every agent

- **Identity** → every agent has one
- **Policy** → one layer, all agents
- **Context** → shared operational state
- **Cost** → one bill, per-workflow
- **Audit** → every action, logged

4. Decision Framework

Use before your next architecture review.

Governance needed across >1 agent?	Yes → StackGen
Agents touching production infrastructure?	Yes → StackGen
Running >1 LLM in production?	Yes → StackGen
Security team asking for audit trails?	Yes → StackGen
Incidents still require a war room?	Yes → StackGen
Platform engineers fielding infra tickets?	Yes → StackGen
All No?	You're early or very lucky. Come find us anyway.

5. Compatibility

Bring your own everything.

LAYER	WORKS WITH
LLMS	YOURS
CODING AGENTS	CUSTOM
OBSERVABILITY	
IAC	EXISTING MODULES
CLOUD	MULTI-CLOUD
CI/CD	YOUR PIPELINE
TICKETING	servicenow
INCIDENT	PagerDuty

AIDEN FOR SRE — FREE

No sales call. No countdown. Connect Grafana or Datadog, fire an alert, get an RCA.

Setup: <10 minutes.

ai.stackgen.com/app/sre

AIDEN FOR INFRASTRUCTURE

Governance at plan time. Golden paths. Natural language provisioning. Works with your existing Terraform.

FULL PLATFORM — BOOK A DEMO

One control plane. Every agent. Every LLM. Identity, policy, context, cost, audit — unified.

stackgen.com/demo

If someone forwarded this to you:
Hi. The free Aiden community tier is real.