⚡ OPERATIONAL READINESS | LEGAL | DATA | COMPLIANCE | RISK | AI

20 Questions Before Your AI Agent *Goes Live.*

Not about model performance. About operational governance.

Who owns it, what it's allowed to do, and what happens when it's wrong.

For anyone building, deploying, or managing AI agents in production.

→ SAVE & SHARE

Swipe →

PRODUCT ENGINEERING OPS AI TEAMS

You're the one who gets the call when an AI agent goes wrong.

Not the team that shipped it.

Most AI incidents aren't model failures.

They're governance gaps — unclear ownership, missing guardrails, undocumented scope creep, and no one watching the drift.

01

Objective & Authority

1. What is the agent's exact objective - and what is it NOT allowed to sacrifice to achieve it?

2. What decisions can it make independently vs. what needs human approval?

3. Is there a spending or resource threshold above which it must stop and escalate?

4. Who defined "success" - and does the metric have boundaries around it?

02

Accountability & Ownership

5	Who is the named human owner accountable for this agent's decisions?
6	If the agent makes a mistake at 2am, who gets the alert?
7	Is there a documented escalation path - or does everyone assume someone else owns it?
8	Who reviews the agent's decision quality - not just accuracy, but judgment?

03

Memory & Learning

9	What is the agent allowed to remember about users or customers?
10	Who approved what it retains - and is anyone checking if those memories are still true?
11	Does stored information expire, or does a six-month-old inference drive today's decisions?
12	Can the agent update its own behavior based on what it learns - without human review?

04

Output & Exposure

13	What customer-facing outputs can the agent produce without a human reviewing them?
14	Could any of its outputs create legal, regulatory, or brand risk?
15	If a regulator asked why the agent said what it said - could you answer?
16	Does the agent identify itself as AI in customer interactions - or does it let people assume it's human?

05

Scope & Drift

17	What systems does the agent have access to - and who approved that access?
18	Has the agent's scope expanded since launch without a formal review?
19	Would the agent pass the same compliance review today that it passed at launch?
20	When was the last time anyone checked whether the agent is still doing what it was originally deployed to do?

Model performance isn't the problem.

Governance is.

Who owns it. What it's allowed to do.

What happens when it gets it wrong.

5 sections. 20 questions. One conversation your team needs to have.

💾 **SAVE THIS POST**

Send it to your team before the next agent goes into production.

♻️ Repost · 💬 Which question does your team quietly skip?

#AIGovernance #AgenticAI #ResponsibleAI #ArtificialIntelligence