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

30 Regulatory Questions Before Your Ship Your AI Agent

Built through a Singapore AI Verify & EU AI Act lens.

For Legal, Data, Compliance, Risk & AI Governance teams.

Operational readiness keeps the agent effective.

Regulatory readiness keeps the business protected.

→ SAVE & SHARE

Swipe →

LEGAL · DATA · COMPLIANCE · RISK · GOVERNANCE

Most teams test the model.

The integrations. The prompts.

Then hope nothing regulatory comes knocking.

Operational readiness ≠ regulatory readiness.

Your engineering checklist keeps the agent working.

It doesn't keep the business protected.

5 sections · 6 questions each · Risk · Boundaries · Data · Transparency · Auditability

01

Risk Classification

1	Have we clearly identified our role under the EU AI Act (provider, deployer, distributor) and mapped corresponding obligations?
2	Have we classified this use case under EU AI Act risk tiers, and documented why?
3	Have we explicitly assessed whether this use case falls under prohibited practices (EU AI Act Article 5)?
4	Is this system in a regulated sector or high-impact context requiring enhanced controls?
5	Has legal or compliance signed off on the final risk classification?
6	If this relies on a GPAI model, have we reviewed the provider's compliance documentation, assessed systemic-risk status, and confirmed its controls are sufficient for our deployment context?

02

Lawful Purpose & Boundaries

7	Is the intended purpose explicitly documented with clear out-of-scope uses?
8	What level of autonomy does the agent operate at (advisory, semi-autonomous, fully autonomous), and are controls aligned to that level?
9	Are there hard policy limits the agent cannot bypass, even if performance improves?
10	Are there defined escalation or human intervention thresholds for high-impact decisions?
11	Are these limits technically enforced, not only written in policy?
12	Do we have an emergency shutdown and rollback capability that can disable the agent within a defined time window if a serious issue is detected?

03

Data Governance & Privacy

13	Do we have traceable data provenance for training, grounding, and memory sources?
14	Have we distinguished between training, fine-tuning, and runtime/grounding data, applied appropriate controls to each, and used only what's strictly necessary?
15	Is sensitive data handling compliant with local privacy and cross-border transfer rules?
16	Is there a documented lawful basis and, where required, valid consent for each category of data processed by the agent?
17	Are retention and deletion policies defined for prompts, memory, logs, and outputs?
18	Is there a periodic review to remove stale or unjustified retained data?

04

Transparency, Rights & Fairness

19	Are users clearly informed when they are interacting with AI?
20	Where applicable, are AI-generated outputs clearly labeled (including synthetic or manipulated content)?
21	Can we explain meaningful reasons for consequential outputs?
22	Can affected users contest decisions and request human review?
23	Have we assessed potential harm to vulnerable groups and added safeguards?
24	Have we measured and documented bias across relevant demographic and contextual dimensions, with defined thresholds for acceptable disparity?

05

Safety, Docs & Auditability

25	Is technical documentation complete for purpose, architecture, limits, and controls?
26	Are decision logs tamper-evident and sufficient for incident reconstruction?
27	Is there a serious-incident process with owners, SLAs, and notification pathways?
28	Do major model, tool, or prompt changes trigger mandatory re-assessment before rollout?
29	Do we have continuous monitoring for performance degradation, drift, and emerging risks in production?
30	Have we conducted adversarial testing AND deployed runtime controls against prompt injection, jailbreaks, and tool misuse, with documented mitigations?

Operational readiness keeps the agent effective.

Regulatory readiness keeps the business protected.

Built through a Singapore AI Verify and EU AI Act lens.

5 sections. 30 questions. Send it to your legal & compliance team.

💾 **SAVE THIS POST**

If your team can't answer most of these, it isn't compliant yet.

♻️ Repost · 💬 Which of these is your org least ready to answer?

#AIGovernance #EUAIAct #ResponsibleAI #AICompliance

www.damandavidpant.com