Path to Production for AI Agents

A practical blueprint to move from prototype to production-ready multi-agent systems.

1) Mobilize

- Define the business outcome, decision boundaries, and success metrics.
- Choose an agent type: assistant, workflow agent, or autonomous agent with approvals.
- Inventory data sources + tools the agent will be allowed to use.

2) Govern

- Set a governance model (AI CoE or designated owners): policy, risk, and approvals.
- Define safety requirements, logging/audit, and access controls from day one.
- Create evaluation criteria (quality, safety, latency, cost) before scaling.

3) Prototype

- Start small with one workflow + one tool; add retrieval only if needed.
- Write test prompts and golden questions; measure failure modes early.
- Add guardrails: tool allowlists, rate limits, and rollback switches.

4) Deploy

- Use repeatable pipelines: build/test/deploy; environment separation (dev/stage/prod).
- Add policy gates: approvals for high-impact actions and sensitive data access.
- Launch canaries and capture telemetry from day one.

5) Operate & Optimize

- Monitor drift, quality regressions, and cost spikes; alert on thresholds.
- Tune prompts/tools; improve retrieval; add caching where appropriate.
- Run post-incident reviews and continuously update your evaluations.