FREE PRACTITIONER PACK

Agentic AI Deployment Pack

20 checks across 7 governance areas. Lightweight enough to actually run; serious enough to survive an audit or a board question.

Pack	Agentic AI Deployment Pack
Document	Governance checklist (20 items)
Use	Oversight · audit readiness
Date	2026-09

AnovaCloud — Engineering the AI-Native Enterprise. Frisco, Texas. This pack is a working instrument, not a brochure: complete it with your team, score honestly, and bring it to your discovery call.

AnovaCloud — anovacloud.com

OWNERSHIP & ACCOUNTABILITY

☐ **Every agent has a named business owner**
A human accountable for what the agent does — not a team alias, a name.

☐ **A technical owner is accountable for behavior**
Someone owns the evals, the guardrails, and the incident response for each agent.

☐ **Ownership survives reorgs**
Owners are recorded in a registry and re-confirmed when teams change.

RISK CLASSIFICATION

☐ **Each agent is risk-tiered**
High / medium / low based on autonomy, data sensitivity, and blast radius — using the HITL tiers as input.

☐ **High-risk agents get enhanced controls**
Additional review, tighter permissions, and human approval where the tier alone is not enough.

INVENTORY & TRANSPARENCY

☐ **A registry of agents exists**
Name, owner, purpose, tier, tools, data sources, model version — one page per agent, kept current.

☐ **Users know when they interact with an agent**
Disclosure where agents act on behalf of the business — in product, in messages, in decisions.

CHANGE CONTROL

☐ **Tier promotions are change requests**
Moving up an autonomy tier requires review, updated guardrails, and rollback criteria — per the HITL framework.

☐ **Model and prompt changes are versioned**
What changed, why, eval results before/after — no silent upgrades to production agents.

☐ **New tools follow an approval path**
Connecting a new API or data source is a reviewed decision, not a config edit.

EVALUATION & MONITORING

☐ **Eval cadence is defined and followed**
Golden-set runs on a schedule and on every material change — per the Evaluation Framework.

☐ **Production behavior is reviewed**
Sampling of real trajectories reviewed by humans; drift triggers investigation.

DOCUMENTATION

☐ **Each agent has a system card**
Purpose, tier, permissions, data, evals, known limitations, incident history — one living document.

☐ **Decisions are recorded**
Why this tier, why these tools, why this scope — the reasoning an auditor or successor will need.

REVIEW & LIFECYCLE

☐ **A review board exists with a real cadence**
Cross-functional review (engineering, security, legal, business) that actually meets and decides.

☐ **Incident and rollback procedures are tested**
Kill switch, runbook, and communication plan exercised — not just written.

☐ **Decommissioning is planned**
How an agent is retired: data cleanup, credential revocation, user communication.

☐ **Regulatory mapping is current**
Applicable obligations mapped to controls, with evidence the controls operate — reviewed when regulations change.

SIGN-OFF

Governance lead: ____________________ Date: _________ Business owner: ____________________ Date: _________