

FREE PRACTITIONER PACK

Agentic AI Deployment Pack

Score up to four candidate use cases against eight candidacy criteria. Kill weak candidates in writing — before the budget does it for you.

Pack	Agentic AI Deployment Pack
Document	Use-case scoring matrix
Use	Prioritization · pilot selection
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. © 2026 AnovaCloud.

AnovaCloud — anovacloud.com

SCORE EACH USE CASE 1-5 (5 = strongest candidacy)

For criteria marked (inv), score 5 when the risk is LOW (e.g. error tolerance: 5 = mistakes are cheap).

Criterion	Use case 1	Use case 2	Use case 3	Use case 4
Value at stake Annual cost, revenue, or risk tied to this workflow. Bigger = better candidate.				
Data & tool readiness Systems of record reachable via API; data quality adequate. Agents cannot outperform their tools.				
Error tolerance (inv) How forgiving the workflow is of mistakes. Low tolerance demands heavy human oversight.				
Human oversight available A named person can review, approve, or intervene without becoming the bottleneck.				
Process stability The workflow is documented and stable — not changing every quarter.				
Regulatory exposure (inv) Regulated decisions (credit, hiring, medical, legal) raise the bar sharply.				
Reversibility Agent mistakes can be detected and undone cheaply.				
Time to value A meaningful pilot is feasible within one quarter.				
TOTAL (of 40)				

INTERPRETATION

32–40 · PILOT-READY — Strong candidacy. Proceed to the discovery questionnaire and pick the autonomy tier deliberately.

24–31 · COPILOT-FIRST — Viable with guardrails. Start as a human-in-the-loop copilot; earn autonomy with measured performance.

Below 24 · DEFER — Do not pilot yet. Strengthen data/tool foundations or pick a sharper use case — and record the decision in writing.

NEXT STEPS

- Run the top-scoring use case through the Discovery Canvas workshop before any architecture discussion.
- Any criterion scoring 1–2 is a named risk: assign an owner and a mitigation, or defer the use case.
- Re-score quarterly — candidacy changes as tools, data, and oversight capacity change.