

FREE PRACTITIONER PACK

Agentic AI Deployment Pack

One page per candidate use case. Fill it in a 90-minute workshop before you score anything — agents fail on vague scope, not on model quality.

Pack	Agentic AI Deployment Pack
Document	Discovery canvas (1 page per use case)
Use	Workshops · scoping · go/no-go input
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

THE CANVAS — one page per candidate use case

1 • JOB TO BE DONE What decision or outcome does the requester need? Who owns that outcome today? • Trigger: what starts this work? • Who is accountable for the result? • What does 'done' look like in one sentence?	2 • CURRENT PROCESS How is it done today — steps, tools, handoffs, and where the pain concentrates. • Steps from trigger to outcome • Handoffs and waiting time • Where do errors or rework happen?	3 • SUCCESS METRIC The one objective measure that decides whether the pilot worked. • Metric + current baseline • Target and measurement method • Who signs off the result?	4 • DATA & TOOLS Everything the agent may read, call, or change — and what is missing. • Systems of record it must read • APIs/tools it may call • Gaps: what does not exist yet?
5 • AUTONOMY TIER How much freedom the agent gets on day one. Start conservative; earn autonomy. • L0 advise · L1 copilot · L2 supervised · L3 semi-auto · L4 autonomous • Why this tier — and what unlocks the next one?	6 • HUMAN CHECKPOINTS Where a person must approve, review, or can stop the agent. • Approve-before-execute points • Review-after-execute sampling • Who is on call when it misbehaves?	7 • RISKS & BLAST RADIUS Worst-case failures, reversibility, and the guardrails that contain them. • Worst realistic failure • Reversible or not? Customer impact? • Guardrails: policy checks, limits, kill switch	8 • CONSTRAINTS Budget, timeline, and compliance boundaries the design must respect. • Budget ceiling for pilot + run • Timeline and hard deadlines • Regulatory / data-residency constraints

HOW TO RUN IT — 90-minute workshop agenda

0:00-0:15	Frame	Pick one use case. Write box 1 and box 3 first — if the outcome and metric are vague, stop and pick a sharper use case.
0:15-0:35	Map	Boxes 2 and 4 with the people who do the work today. Capture tool gaps as action items with owners.
0:35-0:55	Bound	Boxes 5, 6, 7. The autonomy tier is a negotiation between value and blast radius — write the tier down, not the aspiration.
0:55-1:15	Constrain	Box 8. A pilot without a budget ceiling and deadline is a research project.

1:15-1:30	Decide	Score the canvas against the Agent Use-Case Matrix. Defer in writing anything scoring below the pilot bar.
-----------	--------	--

Rule of thumb: a canvas that cannot name the success metric, the data sources, and the human approver is not ready to score — it is ready for another workshop.