

FREE PRACTITIONER PACK

# Agentic AI Deployment Pack

20 checks across 5 categories. Agents cannot outperform their data — run this before you promise any timeline.

|          |                                  |
|----------|----------------------------------|
| Pack     | Agentic AI Deployment Pack       |
| Document | Readiness checklist (20 items)   |
| Use      | Discovery · phase entry criteria |
| 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](https://anovacloud.com)

## ACCESS & COVERAGE

☐ **Every required system of record is reachable**

Named systems, owners, and access method (API, warehouse, file). No 'we'll figure out access later'.

☐ **Authentication path is proven, not assumed**

Service accounts, OAuth flows, or keys tested end-to-end from the agent's runtime — not from a laptop.

☐ **Coverage matches the use case**

The data the agent needs exists for the full scope (all regions, products, or time ranges in scope).

☐ **Unstructured content is inventoried**

Docs, tickets, transcripts the agent must ground on are listed with owners and update cadence.

## QUALITY & LINEAGE

☐ **Known quality issues are documented**

Duplicates, stale records, conflicting sources — listed with impact on agent behavior, not hidden.

☐ **Lineage is traceable to source**

For any answer the agent gives, you can trace the underlying records back to their system of origin.

☐ **PII and sensitive fields are classified**

Sensitive columns/fields tagged; handling rules defined before the agent ever reads them.

☐ **A data owner signs off per source**

Named humans accountable for each source's fitness — data without an owner rots.

## FRESHNESS & LATENCY

☐ **Refresh cadence meets the use case**

If the workflow needs today's data, yesterday's pipeline is a defect, not a detail.

☐ **Latency budget is written down**

Max acceptable staleness per source, agreed with the business owner.

☐ **Backfill and catch-up are handled**

The agent behaves sanely after outages: replays, backfills, or clearly flags gaps.

☐ **Change detection exists**

Schema or content changes in sources are detected and surfaced — not silently absorbed.

## PRIVACY & GOVERNANCE

☐ **Retention and deletion rules are defined**

How long agent-accessible data lives, and how deletion requests propagate.

☐ **Access follows least privilege**

The agent's credentials see only what the use case requires — scoped per source.

☐ **Residency and contractual limits respected**

Data-residency, vendor, and customer-contract constraints mapped to where the agent runs.

☐ **Audit trail covers data access**

Which records the agent read, when, and why — loggable on demand for review or dispute.

## RETRIEVAL READINESS (RAG workloads)

☐ **Chunking strategy is chosen deliberately**

Chunk size, overlap, and metadata tested against real queries — not left at library defaults.

☐ **Retrieval quality is measured**

A labeled query set with relevance judgments; precision/recall tracked before tuning.

☐ **Stale and conflicting content is handled**

Versioning, supersession, and conflict-resolution rules for the knowledge base.

☐ **Evaluation set reflects real questions**

Golden queries sampled from actual users, including the ambiguous and adversarial ones.

## SIGN-OFF

Data owner: \_\_\_\_\_ Date: \_\_\_\_\_ Architect: \_\_\_\_\_ Date: \_\_\_\_\_

Any unchecked box is a named risk with an owner and a date — or a reason to defer the pilot.