

DRAFT — UNOFFICIAL — NOT FOR OPERATIONAL USE

ARCHITECTURE REFERENCE

ODT-GDAP



ADR-0001: LlamaIndex vs DuckDB Responsibility Boundaries

Context

HEADQUARTERS
UNITED STATES ARMY EUROPE AND AFRICA
(USAREUR-AF)
Wiesbaden, Germany

DRAFT — NOT FOR OFFICIAL USE. FOR TRAINING PLANNING PURPOSES ONLY.

20 MARCH 2026

DRAFT — UNOFFICIAL — NOT FOR OPERATIONAL USE

ADR-0001: LLAMAINDEX VS DUCKDB RESPONSIBILITY BOUNDARIES

Date: 2026-03-07 Status: Accepted

CONTEXT

GDAP already uses DuckDB ([GraphStore](#)) as a durable, queryable system of record for doctrine atoms, alignment outputs, temporal history, and DVEE entities.

The LlamaIndex roadmap introduces a retrieval and orchestration layer (ingestion pipeline, indexes, query engines, tools, and evaluation).

Without explicit boundaries, we risk duplicate truth stores, data drift, and unclear recovery procedures.

DECISION

DuckDB remains the source of truth. LlamaIndex indexes are derived, rebuildable serving artifacts.

DuckDB Owns

- Canonical doctrine records: [documents](#), [sections](#), [elements](#)
- Temporal lineage and auditability: [version_diffs](#), [concept_drift](#), [superseded_at](#)
- Coalition/domain analytics: [alignment_pairs](#), [authority_*](#), [exploitability_scores](#), [coalition_seams](#)
- Read-only SQL query surface and deterministic filtering

LlamaIndex Owns

- Document-to-node transformations for retrieval
- Embedding generation and vector/sparse/graph index serving
- Query orchestration: routing, post-processing, synthesis, tool composition
- Evaluation and observability hooks for retrieval/response quality

CONSEQUENCES

Positive: - Clear failure domain separation - Deterministic rebuild path from DuckDB to indexes - Simplified compliance/audit story

Tradeoffs: - Requires robust index lifecycle management - Requires explicit sync/rebuild triggers for version changes

IMPLEMENTATION NOTES

- All LlamaIndex nodes must carry GDAP primary identity metadata (`unique_id` , nation/doc/version/temporal fields).
- Index writes are idempotent and derived from DuckDB snapshots/diffs.
- No workflow may mutate DuckDB through LlamaIndex agent tools.

SEE ALSO

- [GDAP_PERSISTENCE_STRATEGY.md](#) — concrete persistence layers and rebuild policy
- [GDAP_ACCEPTANCE_TESTS.md](#) — Section H: Operations and Lifecycle gates