GRIDBASE LEGAL TECH AUDIT
DOC ID: GB-LEGAL-2026-08
SUBJECT: LEGAL VECTOR DB SECURITY & PRIVILEGE DRIFT
PHASE: ARCHITECTURAL FORTIFICATION

[SECTION 1: SEMANTIC PERIMETER ASSESSMENT]
[ ] Map all vector databases containing attorney-client privileged information.
[ ] Assess the retrieval layer for semantic probing vulnerabilities (latent neighborhood reconstruction).
[ ] Isolate the core inference model from direct read-access to the raw vector indices.

[SECTION 2: PROMPT & INSTRUCTION SANITIZATION]
[ ] Decouple proprietary legal logic (System Prompts) from user-facing query interfaces.
[ ] Implement defensive delimiters to mitigate instruction extraction via direct or indirect injection.
[ ] Design fallback routing to reject queries exhibiting high semantic similarity to core system instructions.

[SECTION 3: PRIVILEGE DRIFT MITIGATION]
[ ] Enforce role-based access control (RBAC) at the embedding layer, not just the application UI.
[ ] Cryptographically hash and log all vector similarity searches to establish an auditable chain of custody.
[ ] Validate that cross-tenant data contamination is structurally mitigated within the Private VPC infrastructure.

[SECTION 4: THE SNAPSHOT PROTOCOL]
[ ] Generate a timestamped baseline of the vector DB access controls and active system prompts.
[ ] Document the current architectural state to satisfy the professional competence mandates of regulatory bodies.

ADVISORY: IN THE AI ERA, THE SYSTEM PROMPT IS PROPRIETARY LEGAL STRATEGY. EXPOSURE OF THE KNOWLEDGE LAYER CONSTITUTES A TERMINAL BREACH OF PRIVILEGE.