MAGELLAN AI · SEMANTIC MATCH ENGINE

Decision Policy

Versioned threshold profiles convert confidence into action

Policy rule

```
decision = bucket(response.confidence, threshold[profile_ver])
```

For each pair, the **semantic-match** edge function (v1.0.0) returns an integer confidence in **0–100** alongside an algorithm class and rationale. A versioned **threshold profile** maps that integer into one of three buckets. The profile, the model identifier, and the resulting decision are written into the ULID-keyed, SHA-256–chained audit ledger so any decision is reproducible under HIPAA §164.312(b), 21 CFR Part 11 §11.10(e), and HHS Information Blocking review.

Default healthcare profile

Bucket	Confidence band	Action
AUTO-ACCEPT	≥ 85	Mapping is committed; written to ledger and knowledge graph.
REVIEW	65 – 84	Queued for steward; rationale + class shown in UI.
REJECT	< 65	Discarded; pair recorded for drift/coverage analytics.

Why threshold profiles are versioned

- A profile change is itself a governed event — bumping AUTO-ACCEPT from 85 to 88 creates a new profile version that is referenced by every subsequent decision.
- Reproducibility: any historical decision can be re-evaluated against the exact threshold profile that was active when it was made.
- Domain tuning: healthcare, finance, and government deployments can carry different default profiles without forking the engine.

Ledger record (excerpt)


```
{ "id": "01HM…ULID…", "prev_hash": "sha256:…", "model": "google/gemini-2.5-flash", "api_version":
"1.0.0", "profile_ver": "hc-2026.01", "confidence": 87, "algorithm": "domain_equivalence",
"decision": "AUTO-ACCEPT", "actor": "user@org" }
```