MAGELLAN AI · SEMANTIC MATCH ENGINE

Algorithm Classes

Five whitelisted decision bases for every semantic match

Overview

The Magellan Semantic Match Engine (**semantic-match** edge function, v1.0.0) calls **google/gemini-2.5-flash** via the Lovable AI Gateway at temperature 0.1 and returns a strict JSON response per attribute pair. Every response carries an integer confidence (0–100), a short rationale, and exactly one of **five whitelisted algorithm classes**. Anything outside the whitelist normalizes to `unknown` and is excluded from policy. Each accepted decision is appended to a ULID-keyed, SHA-256–chained audit ledger.

The Five Classes

Class	When it fires	Worked example
`exact_match`	Normalized attribute names are identical.	patient_id ↔ patient_id
`synonym_match`	Different names, same concept.	subscriber_id ↔ member_number
`abbreviation_match`	One name is a known abbreviation/expansion of the other.	dx_code ↔ diagnosis_code
`context_hierarchy`	Names align by structural context within the schema path.	patient.address.zip ↔ member.mailing.postal_code
`domain_equivalence`	Domain-specific equivalence (HL7/FHIR, claims, clinical).	ICD10_cd ↔ diagnosis_code

Why a whitelist?

A reviewer who sees **confidence = 87, class = domain_equivalence** knows exactly the kind of evidence the engine is asserting and can accept, reject, or override it. Free-text justifications are not auditable; an enum is. The class set is governed and versioned alongside the API.

Response shape

```
{ "confidence": 0–100, "algorithm": "exact_match" | "synonym_match" | "abbreviation_match" |
"context_hierarchy" | "domain_equivalence", "rationale": "short natural-language justification" }
```