# 2.4.4 Semantic Representations

Return to 2.4 Existing Functional Capabilities in Reporting Systems

Reporting environments commonly use semantic representations to express the meaning of reported information. These representations may include:

- Controlled vocabularies
- Taxonomies
- Classification schemes
- Conceptual models
- Ontologies
- Tagged data structures associated with reporting formats

Their purpose is to provide shared designations for the concepts, relationships, attributes, and distinctions used in reporting.

In current practice, organisations often develop and apply semantic representations in isolation or couple them tightly to specific tools, data formats, analytical workflows, or implementation environments. Their role within the broader reporting lifecycle, particularly their relationship to source artefacts, extracted observations, interpretive rules, and validation outcomes, frequently remains implicit rather than explicitly governed.

As a result, different systems, agencies, or reporting periods may interpret or operationalise the same semantic construct differently without adequate visibility, traceability, or governance. Such variation weakens comparability and makes it difficult to determine which definitions, classifications, relationships, or versions governed a particular interpretation.

From a Reference Architecture perspective, semantic representation constitutes a distinct functional capability: the ability to define, manage, version, and evolve the meanings, concepts, relationships, and classifications used by reporting systems independently of how those systems extract, interpret, calculate, or validate information.

Semantic representations provide the governed vocabulary and conceptual structure upon which interpretation and validation operate. They do not, by themselves:

- Perform extraction
- Execute interpretation
- Conduct inference
- Apply calculations
- Produce validation results
- Determine compliance
- Exercise regulatory or legal authority

The Reference Architecture therefore treats semantic representations as reusable, versioned architectural assets positioned explicitly within the reporting lifecycle and traceable to the interpretive and validation outcomes they govern.

Their architectural significance lies not in the selection of a particular representation technology, modelling language, or formalism. It lies in their function as stable semantic references governed through explicit authority, identification, versioning, provenance, and lifecycle controls independently

of any specific implementation technology.

This treatment supports:

- Consistent interpretation
- Controlled evolution of meaning
- Comparison across systems and reporting periods
- Reconstruction of historical interpretations
- Reuse across independently developed implementations
- Separation of semantic authority from tools and code

By separating semantic representation from extraction, interpretation, and validation, FDIS-RA preserves clear architectural responsibilities and prevents semantic intent from becoming entangled with execution logic.

This separation is essential for long-term semantic governance, particularly in federated environments where multiple systems, standards, jurisdictions, and implementations interoperate while retaining independent authority and lifecycle control.

---