The Engage-Meta Reference Vocabulary

Reference Vocabulary

Version 1.0

July 2026

Auteur: Pierre Bonnet (pierre.bonnet@hlfl-consulting.com)

Available on www.engage-meta.com

ENGAGE META

Engage-Meta Helps Companies Build the Cognitive Architecture Required For Reliable Enterprise AI

Introduction

The Engage-Meta Reference Vocabulary is organized around sixteen fundamental concepts that collectively describe the architecture of the framework. These concepts cover the complete lifecycle of enterprise knowledge, from business understanding to AI execution.

The vocabulary includes:

- 2 Business Modeling concepts: *Data Narrative* and *Business Glossary*.
- 2 Data Modeling concepts: *Conceptual Data Model (CDM)* and *Logical Data Model (LDM)*.
- 2 Integration concepts: *Business Object* and *Semantic API*.
- 2 Data Execution concepts: *Operational Data Store (ODS)* and *Data Transfer Object (DTO)*.
- 2 Knowledge Management concepts: *Ontology* and *Enterprise Knowledge Graph (EKG)*.
- 3 AI concepts: *AI Harness Loading*, *Internal Knowledge Distillation (IKD)* and *Reasoning Processing Unit (RPU)*.
- 2 Structural concepts: *Semantic Processing Unit (SPU)* and *Data Domain Service*.
- 1 Organizational concept: *Root Package*.

Together, these concepts establish a common architectural language for designing scalable, governable and profitable Enterprise AI systems.

AI Harness Loading

AI Harness Loading *is the mechanism by which the Engage-Meta framework loads the resources required to execute an AI Harness from the Semantic Processing Unit (SPU). It enables AI agents to retrieve their prompts, briefs, constitutions, procedures and other execution resources independently of the underlying AI platform.*

AI Harnesses are now widely used to structure the behavior of AI agents. They typically include prompts, briefs, constitutions, procedures, security rules and other artifacts that guide the execution of a business task.

Rather than storing these resources inside a specific AI platform, Engage-Meta centralizes them within the Semantic Processing Unit (SPU). AI Harness Loading is the mechanism responsible for retrieving these resources and making them available to the execution environment.

This architecture decouples enterprise knowledge from the technical implementation of AI agents. The same AI Harness can therefore be loaded and executed by different AI platforms, including OpenAI, Anthropic, Hermes or future agentic frameworks, without modifying its business content.

By treating AI Harnesses as enterprise assets managed by the SPU, organizations preserve their knowledge independently of AI vendors, simplify governance and version management, and facilitate the long-term evolution of their AI ecosystem.

Business Glossary (BGL)

The ***Business Glossary (BGL)*** *is the semantic foundation of the Semantic Processing Unit (SPU). It defines the official business vocabulary of the enterprise from which the Logical Data Model, Semantic API, ontologies and Enterprise Knowledge Graph are derived.*

The Business Glossary is the central repository of business terminology within the Engage-Meta framework. It captures the meaning of business concepts using natural language definitions that are understandable by business experts, architects and AI agents alike.

Unlike application-specific vocabularies, the Business Glossary has an enterprise-wide scope. It is intentionally independent of operational systems such as ERP, CRM or MES platforms, although it may reuse terminology commonly adopted within the organization's industry. Its purpose is not to reproduce the vocabulary of existing applications but to establish the enterprise's own semantic reference.

Business terms are organized into Root Packages, which provide the first level of semantic organization for enterprise knowledge and serve as the foundation for defining Data Domains. This organization supports the progressive derivation of the Conceptual Data Model (CDM), the Logical Data Model (LDM), the Semantic API, Conceptual and Operational Ontologies, and the Enterprise Knowledge Graph (EKG).

Within the Semantic Processing Unit, the Business Glossary also supports AI-assisted modeling. From enterprise Data Narratives, AI agents can propose candidate business terms, semantic relationships and Root Packages to accelerate the construction and maintenance of the enterprise vocabulary.

By establishing an explicit and authoritative business language, the Business Glossary enables humans and AI systems to interpret business information consistently, reducing semantic ambiguity while preserving the enterprise's own knowledge independently of its software landscape.

Business Object

*A **Business Object** is a logical aggregation of business information derived from the Logical Data Model (LDM). It defines the business structures exchanged through the Semantic API independently of their physical implementation in operational systems.*

Business Objects represent the information required to perform a business capability through the Semantic API. They are not designed independently but are derived from the Logical Data Model, which itself originates from the Conceptual Data Model and the Business Glossary. As a result, every Business Object inherits the semantic consistency established throughout the modeling process.

At implementation level, a Business Object does not correspond to a single physical structure. Instead, it is assembled from one or more Data Transfer Objects (DTOs) managed by the Operational Data Store (ODS). This abstraction isolates AI agents and consuming applications from the internal organization of source systems while preserving a stable business contract.

Business Objects constitute the request and response structures of Semantic API operations. They expose business semantics rather than technical representations, allowing enterprise capabilities to evolve independently from changes in databases, applications or integration technologies.

Although Business Objects are primarily assembled from operational data managed by the ODS, they may also incorporate semantic information retrieved from the Enterprise Knowledge Graph (EKG). This allows operational data and enterprise knowledge to be combined transparently while presenting a single coherent business representation to AI agents and enterprise applications.

Conceptual Data Model (CDM)

*The **Conceptual Data Model (CDM)** represents the enterprise business knowledge using UML class diagrams expressed exclusively with the terminology defined in the Business Glossary. It provides the authoritative conceptual representation from which the Logical Data Model, ontologies and Enterprise Knowledge Graph are derived.*

The Conceptual Data Model organizes enterprise knowledge into UML packages derived from the Root Packages defined in the Business Glossary. Each Root Package corresponds to a first-level package that may be further decomposed into smaller sub-packages, each represented by its own UML class diagram. This organization maintains a readable and scalable model while preserving a coherent semantic structure across the enterprise.

Within each package, one class is identified as the primary business concept through a UML stereotype. This class represents the concept with the greatest semantic significance and provides the natural entry point for understanding the surrounding business domain.

The Conceptual Data Model intentionally follows a closed-world approach. It defines the business knowledge that is considered valid within the enterprise, while information not represented in the model is assumed to be either unknown or outside the enterprise reference model. This property is essential for Enterprise AI because it reduces semantic ambiguity and prevents AI systems from compensating for missing knowledge through hallucination. When AI agents access the Semantic Processing Unit, they rely on both the Business Glossary and the Conceptual Data Model as the authoritative semantic reference.

Although the Conceptual Data Model is authoritative, it is not static. It is designed to evolve iteratively as enterprise knowledge matures. Engage-Meta therefore supports conceptual modeling through AI assistants integrated into the Semantic Processing Unit, helping architects identify candidate concepts, relationships and model refinements while preserving the overall semantic consistency of the model.

Knowledge that emerges during the life of the enterprise but has not yet been formally incorporated into the Conceptual Data Model can be represented within the Enterprise Knowledge Graph (EKG). Rather than replacing the conceptual model, the EKG extends it through an ontology derived from the concepts already defined in the CDM. This architectural principle ensures that new knowledge remains semantically aligned with the enterprise model instead of creating an independent or conflicting representation.

The Conceptual Data Model therefore establishes the stable semantic foundation of the enterprise, while the Enterprise Knowledge Graph provides controlled flexibility for capturing emerging knowledge. From this conceptual foundation, the Logical Data Model (LDM) is derived and subsequently serves as the basis for constructing the Operational Data Store (ODS) and the business contracts exposed through the Semantic API.

Data Domain Service

*A **Data Domain Service** is a software component derived from a Root Package through the Conceptual and Logical Data Models. It provides data access services for a single business domain and constitutes the foundational service layer of the Operational Data Store (ODS).*

Within the Engage-Meta framework, each Root Package defined in the Business Glossary is progressively transformed into a Conceptual Data Model and subsequently into a Logical Data Model. At implementation level, each logical package becomes a Data Domain Service, a software component responsible for exposing data access services within one and only one business domain.

A Data Domain Service encapsulates the persistence logic of its domain and exposes service contracts that exchange Data Transfer Objects (DTOs) expressed using the enterprise logical naming convention. Its responsibility is limited to retrieving, creating, updating and managing business data belonging to its own domain.

Relationships involving multiple business domains are intentionally excluded from the responsibilities of a Data Domain Service. Cross-domain orchestration is performed at a higher architectural layer by the Semantic API, which combines services from multiple domains to expose complete business capabilities.

This separation of responsibilities is a fundamental architectural principle of Engage-Meta. It promotes service reuse around stable business domains, simplifies software engineering, enables independent testing and maintenance of each domain, and reduces coupling between different parts of the information system.

Data Domain Services are internal implementation components. They are not intended to expose the enterprise business semantics directly to applications or AI agents. Enterprise semantics are exposed exclusively through the Semantic API, while Data Domain Services remain dedicated to efficient and governable access to operational data.

Data Narrative

*A **Data Narrative** is a structured description of the business knowledge expressed through the data manipulated by the enterprise. Within the Engage-Meta framework, it constitutes the primary source for constructing the Business Glossary and the Conceptual Data Model.*

Rather than describing business activities through organizational structures, workflows or use cases, Engage-Meta encourages organizations to express their business knowledge through Data Narratives. A Data Narrative focuses on the meaning of business data, the relationships between business concepts and the way information is created, transformed and consumed across the enterprise.

Writing Data Narratives is therefore not simply a documentation exercise. It is a knowledge engineering activity whose objective is to progressively externalize the enterprise's business knowledge into a form that can be understood by both humans and AI systems. Within Engage-Meta, this capability is supported by the WASI principles (*Write, Analyze, Share and Innovate*), which encourage organizations to continuously improve the quality of their business knowledge.

The AI modeling agents of the Semantic Processing Unit use Data Narratives as their primary source for proposing Business Glossaries, Root Packages, Conceptual Data Models and other semantic artifacts. Because the business knowledge is explicitly expressed, AI agents can infer semantic structures with significantly greater accuracy and consistency.

When Data Narratives are not available, AI modeling agents remain capable of inferring business concepts from traditional use cases. However, this alternative presents two important limitations. First, business teams do not progressively acquire the capability to formally describe their own knowledge, reducing the organization's long-term readiness for Enterprise AI. Second, AI-generated models require more extensive validation because business semantics must be inferred rather than explicitly described, increasing the risk of interpretation errors.

For these reasons, Engage-Meta considers Data Narratives to be the preferred mechanism for capturing enterprise knowledge. By progressively building a corpus of business narratives, organizations improve the quality of the Semantic Processing Unit while establishing a durable knowledge base that supports both conceptual modeling and future AI training initiatives.

Data Transfer Object (DTO)

*A **Data Transfer Object (DTO)** is the implementation counterpart of a **Business Object** within a single Data Domain. It defines the data structures exchanged by a Data Domain Service while remaining strictly limited to the business data managed by that domain.*

The Data Transfer Object (DTO) is a well-established software engineering concept commonly used to exchange data between software components or service interfaces. While DTOs are frequently associated with implementation technologies such as REST or messaging frameworks, their semantic boundaries are rarely defined explicitly.

Within the Engage-Meta framework, the scope of a DTO is intentionally constrained. A DTO may only represent business data belonging to a single Data Domain. Cross-domain business structures are intentionally excluded and are assembled by the Semantic API through the orchestration of multiple Data Domain Services.

This architectural rule preserves the autonomy of each business domain and prevents implementation services from embedding cross-domain business logic.

The Business Object and the DTO represent two complementary views of the same enterprise information. A Business Object belongs to the semantic layer and exposes business capabilities through the Semantic API. A DTO belongs to the implementation layer and provides efficient access to operational data within a single Data Domain. The mapping between Business Objects and DTOs is therefore performed by the Semantic API during service orchestration.

Whenever information from multiple Data Domains must be combined, responsibility is transferred to the Semantic API. The Semantic API orchestrates multiple Data Domain Services and assembles their DTOs into Business Objects that expose complete business capabilities to enterprise applications and AI agents.

This separation of responsibilities is a fundamental principle of the Engage-Meta service-oriented architecture. Data Domain Services remain responsible for reliable and reusable access to operational data, while the Semantic API centralizes all cross-domain business orchestration. The resulting architecture promotes modularity, service reuse, long-term maintainability and a clear separation between implementation concerns and enterprise business semantics.

Enterprise Knowledge Graph (EKG)

*The **Enterprise Knowledge Graph (EKG)** is the graph-based knowledge repository of the Engage-Meta framework. It extends the enterprise semantic foundation established by the Conceptual Data Model while remaining aligned with the Business Glossary and the Operational Data Store.*

Within the Engage-Meta framework, ontologies are derived from the Conceptual Data Model and progressively enriched with additional knowledge. These ontologies are implemented within the Enterprise Knowledge Graph (EKG), which provides a graph-based representation of enterprise knowledge beyond the operational data managed by the Operational Data Store (ODS).

Unlike many AI architectures that position the Knowledge Graph as the central repository of enterprise information, Engage-Meta places the Operational Data Store at the core of the enterprise architecture. The ODS remains the authoritative repository for transactional, master and reference data, providing a unified operational view of the business. The Enterprise Knowledge Graph complements this foundation by capturing knowledge that cannot be efficiently represented within the relational operational model.

Because the Enterprise Knowledge Graph shares the same Business Glossary and conceptual foundation as the ODS, both repositories expose a consistent semantic view of the enterprise. This enables the Semantic API to orchestrate operational data retrieved from the ODS together with extended knowledge retrieved from the EKG within a single business request, without introducing semantic inconsistencies.

Engage-Meta recommends a two-level Enterprise Knowledge Graph architecture. A first, enterprise-wide EKG extends the ODS with unstructured and semantically enriched knowledge while preserving alignment with the enterprise conceptual model. Additional domain-specific EKGs may then be created to support specialized analytical, research or innovation activities. These operational knowledge graphs may progressively diverge from the enterprise conceptual structure or even from the Business Glossary whenever greater analytical flexibility is required.

This architectural approach enables organizations to become AI-native by integrating operational data management and knowledge representation within a coherent semantic foundation, rather than limiting knowledge graphs to analytical use cases alone.

Internal Knowledge Distillation (IKD)

Internal Knowledge Distillation (IKD) *is the process by which enterprise knowledge acquired through Frontier or Premium LLMs is captured within the Semantic Processing Unit (SPU) and subsequently reused to improve AI Harnesses and train Open-Source LLMs. Its purpose is to preserve and continuously increase the value of the enterprise's AI investments.*

Knowledge distillation between Large Language Models is a well-established technique. However, it is increasingly restricted by licensing agreements and, in some cases, by technical mechanisms designed to prevent the transfer of a model's general capabilities to competing models.

Within the Engage-Meta framework, Internal Knowledge Distillation follows a fundamentally different objective. Rather than attempting to reproduce the general intelligence of a Frontier or Premium LLM, it focuses exclusively on capturing the enterprise-specific knowledge generated during AI-assisted activities.

Whenever a Frontier or Premium LLM produces new business knowledge, this knowledge is progressively capitalized within the Semantic Processing Unit (SPU) through its semantic assets, including the Operational Data Store and the Enterprise Knowledge Graph. Once captured, it becomes a reusable enterprise asset that can enrich future AI Harnesses, improve subsequent reasoning processes and serve as training material for Open-Source LLMs operating within the organization's Reasoning Processing Unit (RPU).

Without the semantic foundation provided by the SPU, this accumulation and reuse of enterprise knowledge would require dedicated knowledge management systems built independently from the operational architecture. Engage-Meta makes Internal Knowledge Distillation a natural consequence of semantic modeling, allowing knowledge capture, organization and reuse to emerge directly from the enterprise semantic foundation.

As a result, every token invested in Frontier or Premium LLMs has the potential to generate long-term economic value. Instead of benefiting a single AI automation, newly acquired enterprise knowledge is continuously capitalized and recycled across future AI applications. This creates a cumulative learning process in which the return on AI investment increases over time through the systematic reuse of enterprise knowledge.

Logical Data Model (LDM)

*The **Logical Data Model (LDM)** is the implementation-oriented derivation of the Conceptual Data Model. It transforms enterprise concepts into a normalized logical representation that serves as the foundation for the Operational Data Store, Data Domain Services and the Semantic API.*

The Logical Data Model is derived directly from the Conceptual Data Model, preserving its semantic structure while transforming business terminology into a formal logical vocabulary suitable for software implementation. This logical vocabulary follows simple and consistent naming conventions, such as PascalCase and camelCase, ensuring readability and implementation independence.

The package structure established by the Conceptual Data Model is preserved throughout the logical model. Each Root Package becomes a logical package representing a Data Domain, providing the architectural foundation from which Data Domain Services are implemented. This continuity guarantees that the software architecture remains directly aligned with the enterprise semantic model.

The Logical Data Model is also the level at which implementation decisions are stabilized. Business identifiers, foreign key relationships and architectural patterns for source system identity resolution are formally defined before physical implementation. In particular, Engage-Meta uses the concept of a universal technical identifier, the Unified Object Identifier (uoid), which uniquely identifies every business object independently of the identifiers used by operational source systems.

The Logical Data Model is represented using an Entity-Relationship Diagram (ERD). Conceptual classes become logical tables, while conceptual relationships are transformed into normalized relational structures suitable for implementation within the Operational Data Store.

The Semantic Processing Unit provides AI modeling assistants capable of supporting the derivation of the Conceptual Data Model into the Logical Data Model while preserving semantic consistency across the enterprise architecture.

The Logical Data Model is subsequently implemented as the physical schema of the Operational Data Store (ODS). It also constitutes the structural backbone of the Semantic API. Business Objects exposed through the Semantic API are expressed using the logical structures defined by the LDM, while the Semantic API

orchestrates the mapping between these Business Objects and the DTOs exposed by individual Data Domain Services.

Ontology

*An **Ontology** provides an open semantic representation of enterprise knowledge. Within the Engage-Meta framework, two complementary types of ontologies are distinguished: Conceptual Ontologies, which extend the enterprise conceptual model, and Operational Ontologies, which support specific analytical or research objectives.*

Unlike the Conceptual Data Model, which follows a closed-world approach, an ontology adopts an open-world representation of knowledge. New concepts, relationships and information can be incorporated without requiring the prior extension of the formal enterprise model. This flexibility makes ontologies particularly well suited for integrating knowledge extracted from unstructured or semi-structured sources such as technical documentation, reports or enterprise content repositories.

Within Engage-Meta, a Conceptual Ontology is constructed directly from the concepts and relationships defined in the Conceptual Data Model. The CDM therefore provides the semantic backbone of the ontology, while the ontology extends this foundation with additional knowledge that has not yet been formally modeled. For example, documents, specifications or other unstructured information may be associated with a Customer, Product or Equipment concept already defined in the enterprise conceptual model.

Engage-Meta also distinguishes Operational Ontologies, whose purpose is different. Rather than extending the enterprise semantic reference, an Operational Ontology is created to support a specific analytical, operational or research context. It may introduce specialized concepts, relationships or viewpoints that are meaningful for a particular activity without necessarily becoming part of the enterprise conceptual model.

The boundary between Conceptual and Operational Ontologies is not always absolute and should be evaluated according to the objectives of each project. For example, a knowledge graph developed for a Research & Development department may either reuse the enterprise concepts already defined in the Conceptual Data Model or introduce new domain-specific concepts to encourage exploratory research. The appropriate balance depends on the degree of semantic consistency or analytical freedom required by the organization.

By distinguishing these two complementary ontology patterns, Engage-Meta preserves the stability of the enterprise semantic foundation while allowing knowledge graphs to evolve according to the needs of individual business domains and innovation initiatives.

Operational Data Store (ODS)

*The **Operational Data Store (ODS)** is the transactional data platform of the Engage-Meta framework. It implements the Logical Data Model by unifying transactional, master and reference data into a single operational repository that provides AI systems with a consistent and real-time business view of the enterprise.*

The Operational Data Store is an OLTP relational database that embodies the Logical Data Model. Positioned downstream from operational source systems, it progressively consolidates enterprise data into a unified operational repository while preserving the ability to coexist with legacy applications during the transition toward a unified architecture.

Unlike data warehouses or data lakes, the Operational Data Store is not designed for analytical reporting. Its purpose is to execute the operational activities of the enterprise by providing applications and AI systems with direct access to live business data through a unified semantic model. Engage-Meta therefore places the ODS at the center of the enterprise architecture, enabling organizations to become Enterprise AI-native rather than Analytics AI-native

The construction of the Operational Data Store is one of the primary objectives of the Engage-Meta framework. It results from the progressive derivation of the Business Glossary into the Conceptual Data Model, the Logical Data Model and finally the physical database schema. Data Domain Services expose operational DTOs from the ODS, while the Semantic API orchestrates these services to provide business-oriented access to enterprise capabilities.

The ODS also incorporates an internal identity management architecture capable of resolving the identifiers used by heterogeneous source systems against the unified enterprise concepts defined by the Logical Data Model. Data integration is supported through established patterns such as Change Data Capture (CDC), Data Landing and Data Staging, enabling continuous synchronization with operational systems using mature and widely available technologies.

Beyond data persistence, the Operational Data Store provides the enterprise with a common operational foundation for data governance. Historical management, lifecycle management, business state machines, workflow integration, security and auditing become centralized capabilities shared across the information system rather than being repeatedly implemented within individual applications. Modern relational platforms such as PostgreSQL provide an appropriate technological foundation for implementing these capabilities at enterprise scale.

The Operational Data Store is therefore much more than a database. It is the operational foundation upon which Enterprise AI systems can safely execute business processes, access trusted operational data and continuously interact with the Semantic Processing Unit.

Reasoning Processing Unit (RPU)

*The **Reasoning Processing Unit (RPU)** is the AI execution environment of the Engage-Meta framework. It orchestrates AI Harnesses, LLM routing and Internal Knowledge Distillation while interacting with the Semantic Processing Unit exclusively through the Semantic API.*

The Reasoning Processing Unit constitutes the execution layer of Enterprise AI within the Engage-Meta framework. It is responsible for coordinating AI reasoning through three primary capabilities: AI Harness Loading, LLM Routing and Internal Knowledge Distillation. Together, these capabilities enable AI agents to execute business tasks independently of the underlying AI technologies.

The RPU interacts with the Semantic Processing Unit (SPU) exclusively through the Semantic API, ensuring a strict separation between AI execution and enterprise knowledge management.

Its first interaction consists of retrieving enterprise information from the SPU. Through the Semantic API, AI agents access both operational data managed by the Operational Data Store (ODS) and extended knowledge managed by the Enterprise Knowledge Graph (EKG). Because these repositories share the same semantic foundation, every information flow remains interpretable using the Business Glossary and the enterprise Data Narratives.

The second interaction concerns AI Harness Loading. Prompts, briefs, constitutions, procedures and other reusable AI resources are stored within the SPU and dynamically loaded into the RPU when required. This architecture makes AI Harnesses independent of specific agentic frameworks and allows the same enterprise knowledge to be reused across different AI platforms.

The third interaction supports Internal Knowledge Distillation. Enterprise knowledge acquired through Frontier or Premium LLMs is continuously capitalized within the SPU before being reused to enrich AI Harnesses or train Open-Source LLMs. The RPU therefore participates in a continuous enterprise learning cycle while the SPU remains the persistent repository of accumulated knowledge.

The Reasoning Processing Unit does not maintain an enterprise knowledge repository of its own. It manages only the execution context required by AI systems, whether transient or persistent according to implementation requirements. The Operational Data Store and the Enterprise Knowledge Graph remain the exclusive responsibility of the Semantic Processing Unit. By enforcing all interactions through the Semantic API, Engage-Meta establishes a clear architectural separation between enterprise knowledge management and AI reasoning, allowing both layers to evolve independently.

Root Package

*A **Root Package** defines the highest level of semantic organization within the Engage-Meta framework. It groups business concepts according to a common business concern and provides the structural foundation for the Business Glossary, data models and service architecture.*

Root Packages are established during the construction of the Business Glossary, where business terms are organized into coherent business domains. This organization is neither functional nor organizational. Instead, it reflects the major business concerns of the enterprise and provides a stable semantic structure from which the entire enterprise architecture is derived.

Typical Root Packages may include domains such as Sales, Marketing, Finance, Human Resources, Planning, Supply, Administration or Research & Development. Their purpose is not to mirror the organizational chart, but to organize enterprise knowledge according to meaningful business concepts that remain relatively stable over time.

During conceptual and logical modeling, each Root Package becomes a corresponding modeling package that may be further decomposed into smaller sub-packages in order to maintain readability and semantic cohesion. This hierarchical organization supports the progressive construction of the Conceptual Data

Model and the Logical Data Model while preserving the original business structure established by the Business Glossary.

At implementation level, every Root Package becomes a Data Domain represented by a dedicated Data Domain Service. Each service exposes operational data exclusively within the boundaries of its own business domain through Data Transfer Objects (DTOs). Business capabilities requiring information from multiple Data Domains are intentionally excluded from these services and are instead orchestrated by the Semantic API, preserving a clear separation between domain-level data access and enterprise-level business semantics.

The Root Package therefore provides architectural continuity across the entire Engage-Meta framework. From business terminology to software implementation, it remains the fundamental organizational unit that ensures semantic consistency throughout the enterprise architecture.

Semantic API

*The **Semantic API** is the business interface of the Semantic Processing Unit (SPU). It exposes enterprise capabilities through semantically defined service contracts while shielding AI systems and applications from the implementation details of the underlying operational architecture.*

The Semantic API is implemented within the Semantic Processing Unit (SPU) and constitutes the exclusive access point to the enterprise semantic foundation. Rather than exposing technical services or database structures, it provides business-oriented service contracts expressed entirely in the logical vocabulary derived from the Business Glossary and the Conceptual Data Model.

Because every logical term is formally derived from the Business Glossary, each element exposed by the Semantic API can be traced back to its business definition and corresponding conceptual and logical models. This self-describing semantic foundation significantly improves the understanding of service contracts by both software engineers and AI systems, allowing them to discover, interpret and invoke business capabilities with minimal semantic ambiguity.

The Semantic API also acts as a protective façade for all resources managed by the Semantic Processing Unit, including the Operational Data Store (ODS) and the Enterprise Knowledge Graph (EKG). AI systems operating within the Reasoning Processing Unit (RPU) never interact directly with these repositories. Instead, they consume business capabilities exclusively through the Semantic API, preserving a complete separation between enterprise knowledge management and AI execution.

Each Semantic API operation exposes a complete business service contract composed of business capability descriptions, execution context, preconditions, postconditions, request Business Objects, response Business Objects and execution results. All exchanged structures are expressed as Business Objects, which aggregate information retrieved from one or several Data Domain Services.

The Semantic API is responsible for orchestrating Data Domain Services in order to assemble Business Objects from the underlying Data Transfer Objects (DTOs) exposed by each domain. This orchestration

layer centralizes all cross-domain business logic while allowing Data Domain Services to remain focused exclusively on operational data management.

The Semantic API itself is stateless. Business execution context is carried within service requests and responses, allowing AI systems and applications to manage their own conversational or operational context independently of the API. Temporary execution context may nevertheless be maintained internally while orchestrating multiple Data Domain Services during the execution of a single business operation.

Semantic Processing Unit (SPU)

*The **Semantic Processing Unit (SPU)** is the enterprise semantic foundation of the Engage-Meta framework. It is responsible for constructing, governing and exposing enterprise knowledge through an unambiguous semantic representation that can be consistently understood by both humans and AI systems.*

The Semantic Processing Unit constitutes the persistent semantic memory of the enterprise. It captures, organizes and governs business knowledge independently of any AI platform or software implementation. Rather than storing isolated technical artifacts, the SPU maintains a coherent semantic foundation from which enterprise data, business capabilities and AI resources are progressively derived.

The Semantic Processing Unit integrates the complete semantic engineering process defined by Engage-Meta, including Data Narratives, the Business Glossary, the Conceptual Data Model, the Logical Data Model, Business Objects, Data Transfer Objects, the Operational Data Store, the Enterprise Knowledge Graph, Data Domain Services and the Semantic API. Together, these components establish a unified semantic representation of the enterprise that remains consistent from business knowledge to software implementation.

The construction and maintenance of this semantic foundation are supported by a family of AI modeling agents integrated into the SPU. These agents assist architects and business experts in capturing enterprise knowledge, deriving semantic models and maintaining consistency across the entire semantic engineering lifecycle. The SPU also represents the strategic knowledge asset of the enterprise. AI systems, including those executing within the Reasoning Processing Unit (RPU), never access its internal repositories directly. All interactions occur exclusively through the Semantic API, which preserves a strict separation between enterprise knowledge management and AI reasoning.

This architecture ensures that the enterprise knowledge remains independent of any particular Large Language Model or AI platform. Even in the unlikely event that Frontier or Premium LLMs become temporarily or permanently unavailable, the semantic foundation of the enterprise remains fully preserved within the SPU. New AI technologies, software systems or even manual operational procedures can therefore be connected to the same enterprise knowledge without requiring its reconstruction.

The Semantic Processing Unit is therefore not simply a repository of enterprise information. It is the strategic semantic asset that preserves the enterprise's knowledge, guarantees its long-term continuity and enables its progressive enrichment independently of the evolution of AI technologies.

--- END ---