MOF-based Knowledge Management for a Digital Business Ecosystem

Fotis Kazasis, Nektarios Gioldasis, Nikos Pappas, George Anestis, and Stavros Christodoulakis

[fotis,nektarios,nikos,ganest,stavros]@ced.tuc.gr

Laboratory of Distributed Multimedia Systems and Applications
Technical University of Crete

Abstract. We describe a knowledge management model and middleware for managing the knowledge in a digital business ecosystem environment where SMEs exist cooperate, and compete. The knowledge management model follows the OMG's MOF framework and allows the semantic description of both businesses and services following a multi-layered knowledge organization architecture in a way that share specifications but also differentiate from each other. Several kinds of metamodels are developed for describing various aspects of SMEs and their services. The middleware implementing the knowledge management model is able to handle MOF models (metamodels) and uses XMI documents for metadata interchange. An ontology definition metamodel has been created to allow the import of OWL ontologies in the MOF-based knowledge management framework. This work is part of the EU-funded Digital Business Ecosystem (DBE) Integrated Project which aims at the development of a P2P infrastructure on top of which SME's collaborate, compete and co-evolve.

Introduction

J.F. Moore describes a Business Ecosystem as "An economic community supported by a foundation of interacting organizations and individuals - the organisms of the business world. This economic community produces goods and services of value to customers, who are themselves members of the ecosystem. Over time, they co-evolve their capabilities and roles, and tend to align themselves with the future directions..."[1]. The Digital Business Ecosystem EU-funded Project[2] is the enabling technology for the Business Ecosystem. A Digital Business Ecosystem is defined as "evolutionary self-organizing system aimed at creating a digital software environment for small organizations" that support regional and local development by empowering

---

[1] J.F Moore, The Death of Competition, 1996, pag.6-7

[2] The Digital Business Ecosystem Integrated Project (IP:507953): http://www.digital-ecosystem.org/html/

open, distributed and adaptive technologies and evolutionary business models for the growth of small organizations.

The Digital Business Ecosystem vision is in fact based on two fundamental ideas: those of self-organization and biological evolution. The concept of self-organization implies intelligent behavior and the ability to learn on a short time scale, whereas evolution implies an ability of the system to optimize itself through differentiation and selection of its components on a long time scale. These ideas can be only realized if the (eco) system as a whole is able to learn over time and adapt to this knowledge that itself produces as well as to knowledge that it is derived from its biosphere. So, the main issue is learning. The question however is: what do the SMEs learn in DBE for the short term and what for the long term, and what the knowledge base keeps to facilitate this learning?

Thus, a fundamental objective of DBE is that it aims to build an infrastructure for supporting Knowledge-Based Business Communities. The Knowledge Base (KB) of such a system is one of its core infrastructural components and it will be the enabling factor for the advanced features that were previously described. Knowledge sharing (a main objective of the DBE) requires consistent knowledge representation. On the other hand, DBE should give to SMEs the flexibility to describe their services (as well as their Business Models) in their desired way and differentiate itself from its competitors. Differentiation is a key capability that allows companies to avoid market commoditization which leads to an environment where companies compete only based on the price of the service or product. It is important that each SME will be able to define its own model for describing itself and at the same time this description will be understood by other SMEs (candidate partners). Fundamental to this approach is the Business Modeling Language that is under development in order to provide description of SMEs and their offerings as well as business contracts and agreements. BML is a front-end to the Knowledge Base of the DBE which manages (in a decentralized peer-to-peer manner) all the knowledge about the digital business ecosystem entities and their interaction.

This knowledge will be exploited by other core components for many purposes. One of the main usages of the Knowledge Base is for enabling valuable recommendations (in terms of partnerships and evolutionary actions) by the recommendation infrastructural service to the SMEs which "live" in the ecosystem. Personalized recommendations are provided by a special system component, the recommender, which takes into account business and service models, usage histories, and fitness parameters, in order to recommend best-suited partners and services for the user SMEs.

This paper presents the metadata modeling and management approach taken in DBE, as well as the implementation of the corresponding metadata middleware as part of the DBE Knowledge Base. In what follows we will discuss what are the different kinds of metadata that are stored in the DBE Knowledge Base, we will shortly describe two of the metamodels that are used to specify the structure of the metadata and we will discuss how this metadata are to be exploited in service specification and implementation. Finally we will present the technical approach followed for the implementation of the DBE Knowledge Base.

The DBE Metadata Spectrum

The DBE Knowledge Base (KB) provides a common and consistent description of the DBE world and its dynamics, as well as the external factors of the biosphere affecting it. Its content includes:

1. **Representations of domain specific ontologies** (common conceptualization in a particular domain);
2. **Descriptions of the SMEs** themselves in terms of business vocabularies, business models, policies, strategies, views etc.;
3. **Description of the SME value offerings** (services) and the achieved solutions (service chains/compositions) to particular SME needs.
4. **Models for gathering usage data and statistics** about the DBE environment in order to allow learning and adaptation.

The KB provides a consistent knowledge model and input for the Service Description / Business Modeling Language, the recommendation process and the Service Composition process.

The business metadata that reside in the KB are concerned, they can be classified into three different categories (i.e. the available KB information could be seen from three different perspectives): **Business Domain (Sector) Ontologies**, **SME Business Modeling**, and **SME Service Offerings** (see fig. 1). All these metadata (all the three kinds together) gradually describe business, starting from generally accepted concepts regarding a business domain (e.g. Tourism, Telecommunications, etc.) proceeding to the conceptualization and the description of SMEs (companies) which are the basic business agents in an economy, and finally describing the offerings (services) of each SME in a way that facilitates their matching to the needs of other SME's (B2B scenario) as well as end user demands (B2C scenario).

The DBE follows the OMG's[3] Model Driven Architecture (MDA) approach for structuring specifications expressed as models and support mappings between those models. MDA is considered as a big shift in software development. With a set of software standards that impose discipline in the development process, MDA aims to increase the control over the applications and also over the changes that affect businesses. The core of MDA is UML, as the medium to design and document software applications. Using models as an essential part of the development process is now recognized as a best practice, known generally as Model Driven Development (MDD). MDA, however, attempts to broaden the benefits of using models with model-to-model transformation by allowing (semi) automatic code generation from high level (even informal) business models. To ensure that applications do not become obsolete as platforms and infrastructures evolve, MDA allows to define software independently of the infrastructure or platform on which it is going to be executed. This goal is achieved by creating Platform-Independent Models (PIMs). However, in order for the software to take advantage of a specific platform or infrastructure, PIMs must be transformed to Platform-Specific Models (PSMs), which describe the application within the native constructs of the target platform.

[3] http://www.omg.org

MDA seems to be the most appropriate approach for a digital business ecosystem where flexible mechanisms are needed in order to allow SMEs to adapt to market changes and join to established collaborations. In particular the adoption of the MDA approach provides lots of benefits to SMEs like the reduced risk of lock-in to specific technologies, increased business agility, reduced development and maintenance costs and reduced time-to-market.

Independently of any information modeling approach, an important requirement that it should be stressed out is the need for meta-modeling. SMEs should have the flexibility to describe their services (as well as their Business Models) in their desired way. This requirement refers to free market competition where each enterprise is trying to differentiate itself from its competitors. Having said this, it is obvious that we should provide mechanisms that will allow SMEs to describe their service offerings in their own way. On the other hand, SMEs must be able to share knowledge and this knowledge to be understood by other SMEs. Thus, the modeling mechanisms should enable not standardized (i.e. not pre-defined) but interoperable (i.e. understandable by communicating parties) metadata modeling. The metadata modeling approach in DBE is based on the OMG's MOF (Meta-Object-Facility) metadata architecture which is which comprises the following layers:

1. **The information layer.** It consists of the data that we wish to describe.
2. **The model layer**. It comprises the metadata that describe data in the information layer. Metadata is informally aggregated into models.
3. **The metamodel layer.** It consists of the descriptions (i.e., meta-metadata) that define the structure and semantics of metadata. Meta-metadata constructs are informally aggregated into metamodels. A metamodel is essentially an "abstract language" for describing different kinds of data.
4. **The meta-metamodel layer.** It consists of the description of the structure and semantics of meta-metadata. In other words, it is the "abstract language" for defining different kinds of metadata.

According to this approach each metadata element that will be stored in the KB should be placed as instance of a model at a higher layer of the MOF meta-data architecture. That is, MOF based (abstract) languages or mechanisms will be defined for defining each segment of information. For the rest of this paper the terms language and metamodel are interchangeably used. The DBE framework has adopted the OMG's approach of business modeling[4] (i.e. Business Rules, Business Processes, and Business Organization) and the OMG's approach to ontology modeling[5] (i.e. business domain ontologies).

In the following sub-sections, we present these kinds of metadata discussing also what they describe relationships among these kinds of information (see Fig. 1).

[4] http://bei.omg.org/
[5] http://adtf.omg.org/

Business Domain-Specific Metadata

Ontologies are well known mechanisms for knowledge representation and are frequently used to mean a specification of a conceptualization. Ontologies have three main characteristics[6]:

1. **They are formal**, which means that they are machine-readable.
2. **They are explicit**, the type of concepts used and the constraints on their use are explicitly defined.
3. **They specify a shared conceptualization**, which means that an ontology captures a consensual knowledge, that is, it is not private to some individual, but accepted by a community.

Ontologies are stored into the KB in order to specify the *shared conceptualization of particular business domains*. These ontologies are called **domain specific ontologies**, and as an example, one can think of ontologies for the tourism domain, the telecommunications domain, etc.

---

[6] Studer R, Benjamins VR, Fensel D (1998) Knowledge Engineering: Principles and Methods. IEEE Transactions on Data and Knowledge Engineering 25(1-2):161-197

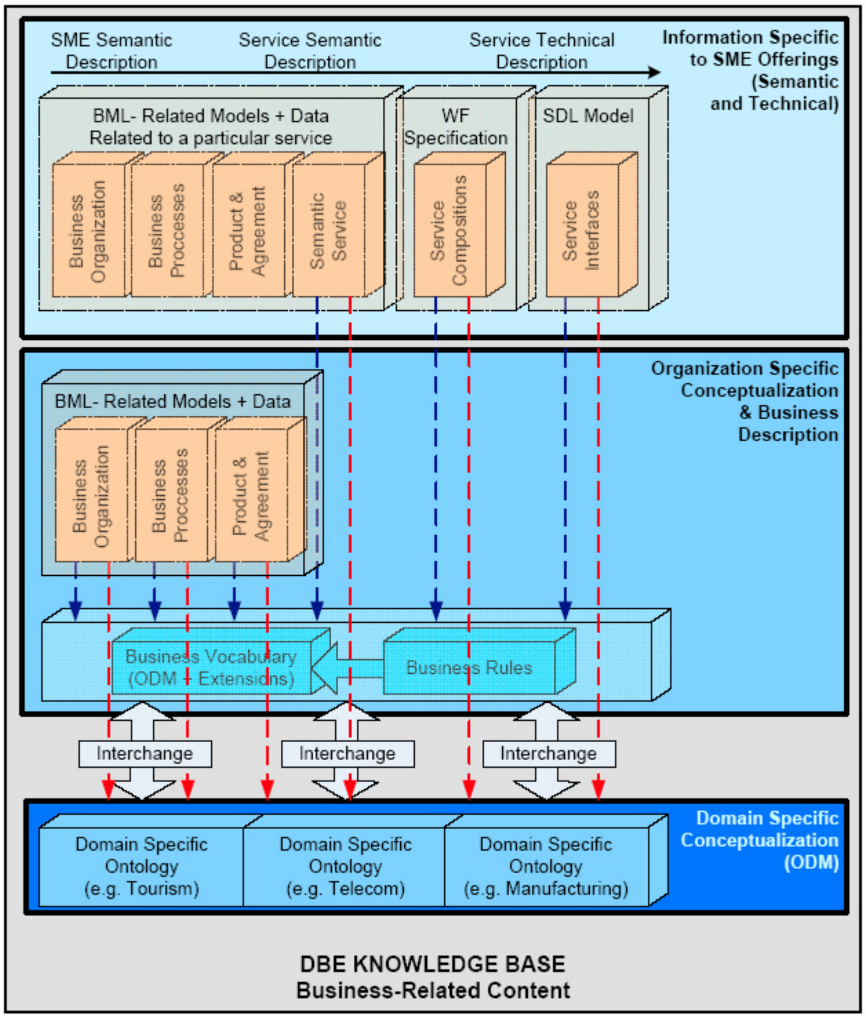

Fig. 1. Business Related Knowledge Stored in the DBE Knowledge Base

As pointed out, the DBE KB follows a MOF metadata approach. As a consequence, the MOF-based approach should be also used for representing and storing Ontologies. However, the most widely accepted language for storing ontologies nowadays is the Ontology Web Language (OWL)[7], which has been produced by the Semantic Web Community and is maintained by W3C. Existing ontologies have been described and published using OWL by the semantic web community.

[7] http://www.w3.org/TR/owl-ref/

Thus, an important feature for the DBE Knowledge Base is to be able to import ontologies described in OWL ensuring interoperability with the semantic web community. This goal has been recognized by OMG, which has issued a Request for Proposal asking for metamodels that will be used for ontology definition. OMG refers explicitly to the need for interoperability with OWL ontologies. A first version of such a metamodel has been developed for the DBE Knowledge Base. This metamodel (it will be discussed below) will be used for representing Domain Specific Ontologies that are created by domain experts and are a reference specification of the conceptualization provided for each domain. These ontologies are expected to be used in both business modeling and information system modeling (and interoperation) as explained below.

Organization-Specific Metadata

The next kind of business related metadata that resides in the DBE Knowledge Base refers to the specification of the conceptualization internally to SMEs as well as to several aspects of its operation (business processes, rules, etc.). This information is specified in the DBE Business Modeling Language (BML). Briefly, the following languages are used for Business Modeling:

- **Business Semantics for Business Rules** (BSBR)**.** This component comprises two sub-components. The first is used to define the vocabulary of terms that is used for the specification of the conceptualization used internally to an organization while the second is used for defining Business Rules and Policies that govern activities, roles, and processes inside the organization.
- **Business Processes** (BPDM). Describes the processes that have been set in an organization and used for specifying operational aspects of the organization.
- **Business Organization** (BOM). This component describes the way that a company is internally organized by specifying its structure, departments, roles, etc.

Moreover, DBE has identified the need to extend this vision by providing a mechanism for describing Product and Agreement aspects of business.

Service-Specific Metadata

The DBE Knowledge Base also stores (and makes public through the Semantic Registry Service) metadata regarding specific offerings of SMEs. The purpose of this kind of metadata is dual: a) to advertise services and b) to provide precise technical information and semantics to the potential service users.

The advertisement of a service requires semantic information not only about the service itself, but also about the provider: the company behind the service. This kind of information is already stored in the KB as described in the previous section and it will be re-used for service advertisement purposes.

What is missing, is the advertisement of the service itself; the semantic description of a service. Thus, a Semantic Service Language (SSL) has been developed as an additional component of BML and is used to provide modeling for the semantics of the services offered by SMEs. The metadata defined by this language provide a semantic description of services regardless whether these services are implemented by software components and are accessible via the network. These metadata may also reference terms in Business Vocabularies and Business Rules defining the semantics of the provider SME as well as to Domain Specific Ontologies for effective support of query interoperation and information retrieval.

In case that a service is implemented by some software component accessible through the network, the semantic description of a service is not adequate, and the technical information is also needed. There are two kinds of technical information available for a service: a) the service programmatic interface and b) the service composition specification in case that a service is a composite service.

A Service Description Language (SDL) will be used for providing the description of the service programmatic interface, while the Business Process Execution Language for Web Services (BPEL4WS)[8] will be used for providing the Workflow specification of a composite service.

From another point of view, there is a horizontal dimension related to this information (Service Specific Information) as shown in Fig. 1. In this horizontal dimension, a service description provides information about its technical front-end (interface), its technical internal logic (Workflow specification), its own semantics, and its provider's semantics. It should be also stressed that not all these kinds of metadata is required for providing a service in DBE. Dashed lines in Fig. 1 show optional metadata segments.

The Ontology Definition Metamodel

This section briefly describes the Ontology Definition Metamodel (ODM) currently developed in DBE for handling domain specific ontologies. The design goal of the ODM is the compatibility with the Web Ontology Language defined by W3 consortium. As described before, DBE uses a fixed, four-layer metamodeling architecture for metadata management. On the other hand, the number of the layers in the metamodeling architecture of the semantic web is non-fixed, and the layers are defined with different semantics than those of MOF. Fig. 2 shows the ODM in the context of MOF metadata architecture and where it stands with respect to the Semantic Web's metamodeling architecture.

The Ontology Definition Metamodel that is being developed is an M2 layer Model in the MOF metamodeling architecture. Instances of this metamodel (M1 Layer Models) are the ontologies that the experts or the users define. It should be noted that according to ODM defined in DBE, it is possible to describe both classes and their instances (individuals) in M1 layer. This has been done in order to allow DBE to be able to handle existing ontologies described with OWL (since OWL describes in the same ontology both class definitions and individuals). Ontologies that will be described with ODM and stored in the DBE KB will be encoded in XMI[9] format.

---

[8] http://www-128.ibm.com/developerworks/library/ws-bpel/

[9] http://www.omg.org/technology/documents/formal/xmi.htm

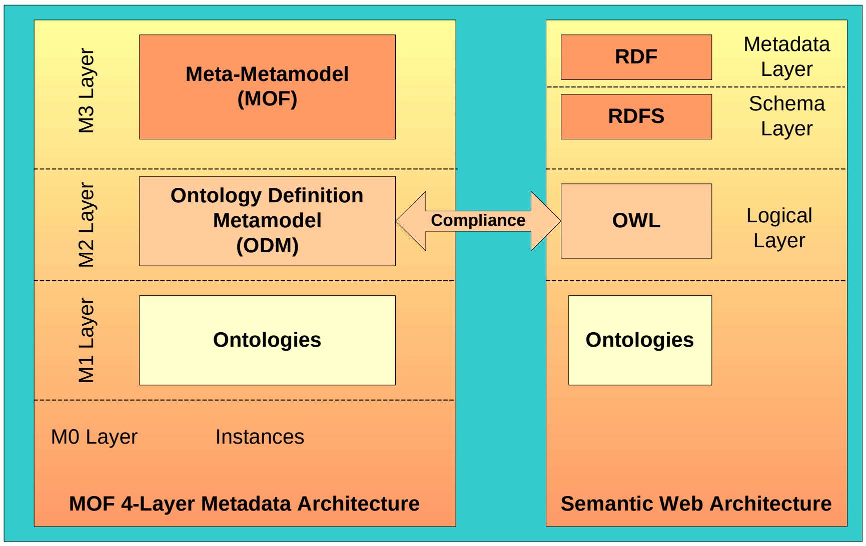

Fig. 2. The Architectural Approach of ODM

The Semantic Service Language

Intelligent service discovery, as envisioned in DBE, requires semantic description of services and service providers, description of technical details of services, knowledge about the user needs and preferences, usage history data, and so on. This section briefly describes the approach followed for the Semantic Service Language.

DBE provides a meta-model (M2) according to which SMEs will be able to develop their own model (M1) for describing their services. This meta-model will provide the required modeling primitives that the user (modeler) will have in his hands to build custom models for semantically describing his services. These models can then be instantiated and populated with the real data (M0) that refer to specific services and be published into the DBE Knowledge Base.

This metamodel is the abstract syntax of the Semantic Service Language, which has been developed. In the DBE project we envision this language to be part of an entire Meta-Modeling Framework, which once completed, it will describe (among others) the following three dimensions of DBE services:

1. ***What a service does:*** This is mainly the semantic description of the service and provides semantic information about the special service properties like the business sector that belongs to, what are its functionalities, geographic area that it applies, charging schemes that are used, etc. It may also provide description of what does service consumption require from the user and what are the outputs or the effects that it produces. This part will be covered by the Semantic Service Language and the corresponding metamodel will be described in this document.

2. ***How it works:*** This kind of description refers to composite services. It will describe the composite services as processes defining the details of the service execution pattern as well as the formal inputs and outputs that a service execution requires and produces (probably under conditions). This kind of description will be produced by the service composer and at the moment it is based on the Business Process Execution Language for Web Services specification.
3. ***How it is consumed:*** This part describes the technical details (i.e. binding information, message formats, etc.) of how the service is consumed. It is actually the technical specification of the service interface, and it will be defined using the Service Description Language.

The MDA Approach to Metadata Exploitation

MDA is concerned with organizing models used in the s/w development process so that developers can move from abstract models to more concrete models. This focus emphasizes the use of Computation Independent Models (CIM models), Platform Independent Models (PIM models), Platform Specific Models (PSM models) and mappings that allow transforming one model into another.

Fig. 3 provides an overview of how the various DBE BML and Ontology related languages and models are organized and aligned following the adopted MDA approach.

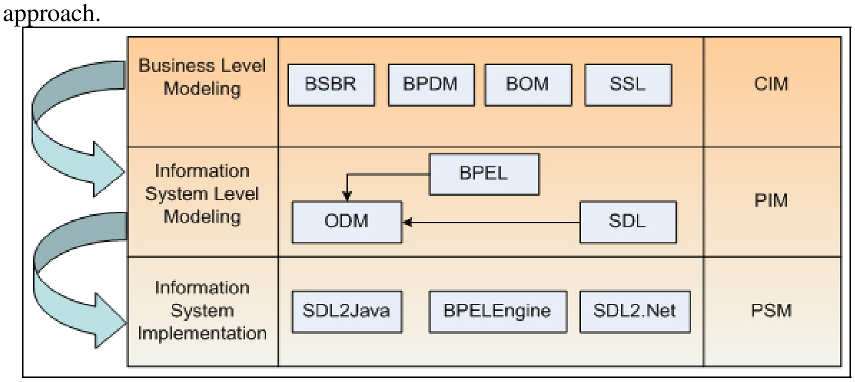

Fig. 3. The MDA Approach of DBE

At the business level (mainly used by business analysts and experts) one can see CIM models that include the models used to describe a business and its offerings. The metamodels for Business Rules, Business Processes, Organization and Semantic services are all examples of the DBE CIM Models. At the information system level (mainly used by s/w architects and designers) one can see PIM models that describe a s/w design in a way that does not assume any specific implementation. The workflow specifications of composite services and the technical description (with possible use of domain ontologies) of service interfaces are examples of the DBE PIM Models.

Last at the information system implementation level (mainly used by the s/w developers and programmers) one can see PSM models that assure a target platform. SDL2Java, SDL2.Net and BPEL Engine are examples of the DBE PSM Models.

The Implementation Approach of the Metadata Management Middleware

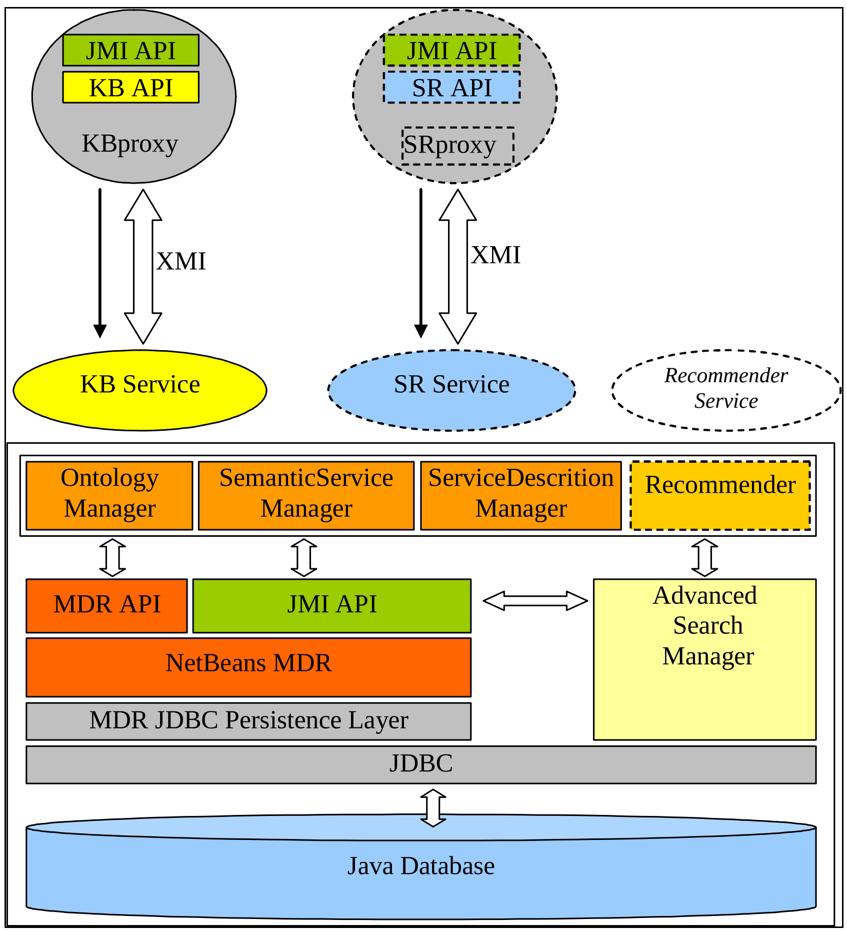

Fig. 4: The KB Metadata Management Infrastructure

As described above, the DBE Knowledge Base infrastructure follows the OMG's MOF and MDA approaches. The metadata management component of the knowledge base is able to handle metamodels and models providing MOF Repository functionality and uses XMI documents for metadata interchange.

The KB Architecture is based on Netbeans/MDR[10] along with a data management system (Java database) and is compliant with the JMI.1.0[11] specification. The so far developed metamodels are imported and appropriate functionality for processing; storing and retrieving metadata is already supported.

The basic manager infrastructure contains the Ontology Manager, Semantic Service Manager, Service Description Manager, and Advanced Search Manager. These modules are supported by the MDR and use the JMI interfaces. The Advanced Search Manager provides advanced search features exploiting both the JMI and the search functionality of the underlying data management system.

On top of this basic infrastructure the DBE core services related to knowledge base, namely the Knowledge Base Service (KB Service) and the Semantic Registry Service (SR Service) are running. It is possible during the installation phase to configure the system so as to enable only the features that are necessary to operate the SR Service or the KB service respectively.

Using suitable proxies that provide functionality for storing, retrieving and searching for models, front-end applications can create new models, search for or browse existing models that are stored in KB.

Appropriate relational database models have been developed for each of the aforementioned metamodels in order to allow efficient storage and manipulation of their instance models and corresponding data. The design of these relational models have taken into account involved requirements of the specific metamodels (e.g. OWL object properties and hierarchies).

Conclusions and Future Work

This paper describes the design and preliminary implementation of a MOF-based metadata middleware for supporting knowledge sharing among networked businesses. Although the current implementation is a centralized one, the architecture of the envisioned system is a general P2P meta-data management middleware and could be used in many application domains.

It is anticipated that during the implementation of the described functionality in a P2P environment we will face challenging research issues related to:

- Ontology Management (creation, maintenance, conflict resolution) and exploitation in P2P environments with the use of both MOF Repositories and Relational Databases at each peer following a Service Oriented Architecture
- Model Management (creation, maintenance, versioning, validation) in P2P systems.
- Distributed Semantic Recommendation of services exploiting the metadata stored in the metadata middleware (part of the DBE Knowledge Base).

---

[10] http://mdr.netbeans.org/

[11] http://jcp.org/aboutJava/communityprocess/final/jsr040/index.html

Acknowledgements

This work is being carried out in the scope of the DBE Integrated Project (IP:507953) funded by EU under FP6.