A P2P and SOA Infrastructure for Distributed Ontology-Based Knowledge Management

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

Lab. Of Distributed Multimedia Information Systems
Technical University of Crete (MUSIC/TUC)
University Campus, Kounoupidiana, Chania, Greece
[nektarios,nikos,fotis,ganest,stavros]@ced.tuc.gr

Abstract. The Digital Business Ecosystem Integrated Project funded by the EU aims at the development of a Europe-Wide virtual economic environment where European SMEs will advertise themselves and their service offerings with the purpose of competing, collaborating and co-evolving with other SMEs. One of the main strategic goals for this project is to model, design and implement such an environment based on observations made on biological ecosystems where several species live together, compete and co-evolve as the physical environment of the ecosystem changes. From a technical point of view, the DBE will provide a Distributed Open-Source Infrastructure which will be the glue that will bring together the European SMEs and will enable knowledge sharing and intelligent discovery of partners and services. The focus of this paper is on the Knowledge Management Infrastructure of the DBE. We firstly discuss the problem, we then present the proposed approach, and finally we identify some of the research issues that we will concentrate on.

1 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 is the enabling technology for the Business Ecosystem. A Digital Business Ecosystem is defined as "evolutionary self-organising system aimed at creating a digital software environment for small organisations" that support regional and local development by empowering open, distributed and adaptive technologies and evolutionary business models for the growth of small organisations.

The Digital Business Ecosystem vision is in fact based on two fundamental ideas: those of self-organisation and biological evolution. The concept of self-organisation implies intelligent behaviour and the ability to learn on a short time scale, whereas evolution implies an ability of the system to optimise 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.

Thus, a fundamental objective of DBE is that it aims to build an infrastructure for supporting Knowledge-Based Business Communities. The Knowledge Base 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. Within the DBE the purpose of Knowledge Base is to provide a common and consistent description of the DBE world and its dynamics, as well as the external factors of the biosphere affecting it. The Knowledge Base will provide a consistent Knowledge Representation Model for storing Business and Service Ontologies, Business and Service Descriptions, Regulatory Framework, and usage history regarding the day-to-day activities in the DBE. 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 "living" 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. The knowledge Base will store the suggestions made by the recommender, in order to provide feedback for the evolutionary procedures, which will lead to the long-term study and adaptation.

This paper presents ongoing work done in the DBE project. Its focus is on the knowledge modelling and management requirements and the architectural approach envisioned for DBE Knowledge Base Infrastructure.

2 The DBE Knowledge Base Requirements

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, among others, representations of the SME's Business Models and Ontologies, the Service Models and Ontologies, the SME views of the ecosystem, the user models, and models for gathering statistics. The KB is being used in order to provide a consistent knowledge model and input for the Service Description / Business Modeling Language, the recommendation process and the Service Composition process.

An important aspect of the KB management is to support the sharing of its content. For that it should follow a scalable design/implementation that can support virtually any number of DBE users as well as organization-wide transactions and/or cooperation. The KB provides a variety of advanced content services (and tools) that include the following:

- **Content management services:**
  these services provide functionality for storing, finding, retrieving and presenting any type(s) of the aforementioned content that it hosts.
- **Content personalization services:**
  the services of this type utilize DBE user profiles and content access patterns to create and/or present customized content to each DBE user.
- **Content group services:**
  such services provide scheduled delivery of content according to predefined business rules, as well as packaging the content for individual DBE users (e.g. services for DBE users that participate to a special cooperation schema).
- **Content aggregation services:**
  These services include automatic combination of content from a variety of content sources and (probably) formats (i.e. other KBs and/or legacy systems).

Another important aspect of the KB management is to support the sharing of its content in a P2P environment. At first glance, many of the challenges in designing P2P systems seem to fall clearly under the banner of the distributed systems community. However, upon closer examination, the fundamental problem in most P2P systems is the placement and retrieval of data. Indeed, current P2P systems focus strictly on handling semantics-free, large-granularity requests for objects by identifier (typically a name), which both limits their utility and restricts the techniques that might be employed to distribute the data. These current content sharing systems are largely limited to applications in which objects are large, opaque, and atomic, and whose content is well-described by their name; for instance, today's P2P systems do not emphasize content-based retrieval of text files or fetching only the abstracts from a set of text documents. Moreover, they are limited to caching, pre-fetching, or pushing of content at the object level, and know nothing of overlap between objects.

These limitations arise because the current P2P world is lacking in the areas of semantics, data transformation, and data relationships, yet these are some of the core strengths of the data management community. Organizational Information is complex and has complex relationships in the real business world, and the DBE seeks to capture and manage such information. Queries, views, and integrity constraints can be used to express relationships between existing objects and to define new objects in terms of old ones. Complex queries can be posed across multiple sources, and the results of one query can be materialized and used to answer other queries. Data management techniques can be used to develop better solutions to the data placement problem at the heart of any P2P system design: data must be placed in strategic locations and then used to improve query performance. On the other hand businesses in the real world are independent and they may cooperate, but at the same time compete. The environment is clearly a P2P environment where some, but not all the information of the business is revealed to the outside world. From the above, it becomes clear that the use of relational database management technology for managing the DBE KB content is very important , and it can leverage the increased scalability, reliability, and performance of a successful P2P architecture for the DBE project.

The DBE KB supports the Recommendation Services that will act as an autonomous processes that manage SME preferences (either business preferences or service preferences) and matches theses preferences with available business descriptions and service descriptions. An important assumption and investigation focus of DBE which is used in the design and implementation of the Recommendation mechanisms is the existence of powerful business and service Ontologies that capture the semantics of business models and service descriptions. These Ontologies are used to define the corresponding preferences for businesses and services. These preferences are attached to the business model of each SME in the DBE environment. The Ontologies within DBE are described in an XML based language (XMI) and mapping mechanisms will be used to ensure the transformation from different Ontology languages to the standard Ontology descriptions used in the DBE. The DBE Ontologies can thus be used by the recommendation mechanisms regardless of the particular models used by each SME in the DBE environment.

3 A Preliminary Knowledge Management Architecture

In this section we present a preleminary conceptual architecture (i.e. a set of logical layers and components and relationships) for the DBE Knowledge Base Infrastructure and other system components that directly or indirectly exploit this infrastructure. The purpose of this architecture is to show how the DBE Knowledge Base will be perceived by the rest DBE System Components and how it will be exploited by them.

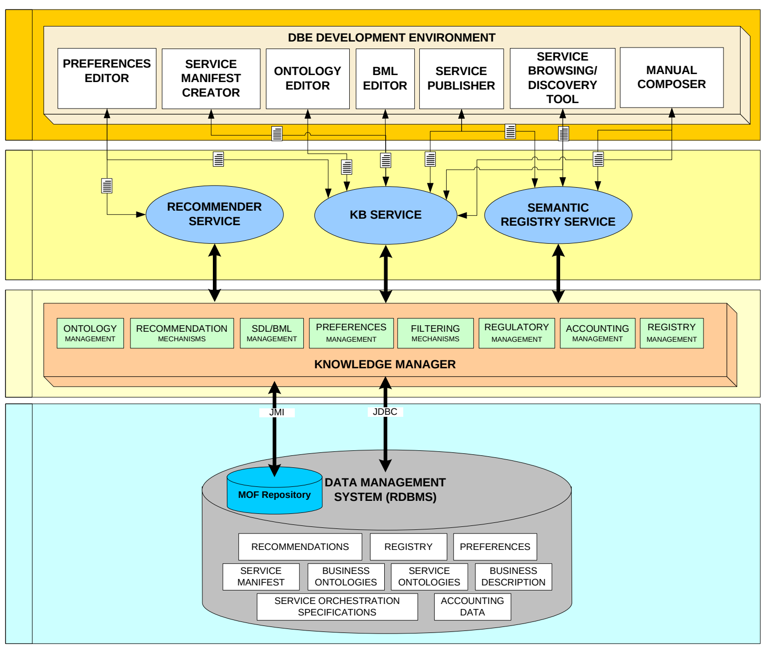

Figure 1 : Conceptual Architecture of the DBE KB Infrastructure

The abstract architecture of the DBE Knowledge Base illustrated in Figure 1 clearly distinguishes between the client modules (front-end applications that provide

GUIs with which one can create and access knowledge), the enterprise-level services (a set of core services that expose the functionality of the KB to the client modules), the middle-tier management modules (which implement the core logic of the Knowledge Base) and the back-end persistency infrastructure that permanently stores the DBE Knowledge.

1.1 System Components

In this section we briefly describe the main components that appear in Figure 1 and in particular their use with respect to Knowledge Base Infrastructure. We start by describing the Core DBE Services that are used to export the functionality of the DBE Knowledge Base Infrastructure to the rest System and then we describe several components that exploit the aforementioned services and are used to create, store, modify, and present Knowledge which is finally stored in the DBE Knowledge Base.

From a technical point of view, the DBE KB will implement a MOF-Based Repository utilizing the Java Metadata Interface (JMI) [15] specification. This approach will provide a standard and flexible mechanism for handling different knowledge models. The persistency of the KB contents will be achieved through a java-based relational database management system providing this way a fully platform independent solution. Thus, meta-data browsing can be done through standard interfaces (JMI) while advanced recommendation mechanisms (requiring powerful querying features) can be implemented on-top of the RDBMS where the real contents are stored.

- KB SERVICE:The KB Service is a Core DBE Service that encapsulates all the functionality of the DBE Knowledge Base and provides a standard interface through which the other DBE components can use this functionality in a coherent and consistent manner. Each higher level application that wants to exploit the Knowledge Base has to invoke the required KB Service operations. The core logic of the KB service is implemented by the Knowledge Base Manager which is responsible for semantically organizing and managing the various KB Data as well as handling the requests for accessing those data that are coming through the KB Service. In the distributed DBE environment, although the KB Service will be perceived as a single service, it will actually be a set of KB Service Instances (each instance will be hosted in an SME node) that will be self-coordinated and will operate as a whole over a P2P network following this way a SOA approach.
- RECOMMENDER SERVICE:The Recommender Service is a Core DBE Service responsible of handling SME preferences that specify the needs of each independent SME in the DBE environment in terms of partnerships with other SMEs as well as services needed to compose more complex services for its customers. It provides a standard interface that hides the technical details of the underlying Recommendation Mechanisms that implement powerful matching algorithms using the DBE Business and Service Ontologies. The produced recommendations are permanently stored in the Data Management System (Persistency Tier). As with the KB Service, in the distributed DBE environment, this service will be a set of Recommender Service Instances over a P2P network.
- SEMANTIC REGISTRY SERVICE: It provides a standard interface capable to implement a Semantic Registry Service functionality required in the DBE envi-

ronment. The Semantic Registry Service provides standard representation hierarchies and query facilities provided by the standard registries. It should be noted that this service simulates the functionality of a Semantic Registry and it is not a repository itself. On the contrary, it exploits the Data Management System (Persistency Tier) for the storage and retrieval of its contents. At the Middle-Tier a Registry Manager is used to implement the advanced functionality of this Service and to handle the requests for retrieving and accessing the stored content. The Semantic Registry Service is a Core DBE service and from a technical point of view is a set of Semantic Registry Service instances that are hosted on special peers (super peers) over a P2P network.

- ONTOLOGY EDITOR: The Ontology Editor is a component that provides a GUI to the end user and it is used for the management (creation/update/retrieve) of the Ontologies (Service and Business Ontologies) that are developed and used in the DBE environment. In order to import a new ontology in the DBE an extension of the DBE Knowledge Base is required. Such a task is under the responsibility of the Ontology Editor. It accesses the KB by "talking" to the (local instance of the1) KB Service for manipulating the various Ontologies (store/retrieve/update Ontologies). The mechanisms that perform the real manipulation of the Ontologies are implemented at the Middle-Tier by the Ontology Management Module. The final modifications are stored in the persistent Data Management System.
- BML EDITOR: It is a Tool that resides on each DBE node (at least on the Nodes that provide DBE Services) and it uses the Business Ontologies (created with the Ontology Editor) in order to describe SMEs according to their business models, policies, assets, competencies, partners, etc. It also uses the Service Ontologies in order to describe the services that are offered by the SMEs. It interacts with the KB Service for accessing Business Ontologies and storing/updating/retrieving Business Descriptions and Service Descriptions.
- SERVICE MANIFEST CREATOR: It is a Tool that resides on each DBE node (at least on the Nodes that provide DBE Services) and it is used to integrate the BML Description and SDL Description of DBE Services producing their Service Manifest. It interacts with the local instance of the KB Service to store/update/retrieve Service Manifests.
- SERVICE PUBLISHER: It is a tool that is used to publish a Service Manifest to the Semantic Registry. It "talks" with the KB Service in order to retrieve Service Manifests and with the Semantic Registry Service in order to publish/update/remove Service Manifests (the published ones). The Semantic Registry Service is based on the Data Management System for the storage of this content.
- SERVICE BROWSING/DISCOVERY TOOL: It is a Tool that resides on each SME node and it is used to contact the Semantic Registry Service through its local instance in order to browse and retrieve the contents of the Semantic Registry of the DBE. This tool has two modes: The simple one, where it just browses the registry for Services, and the advanced one where using Business and Service Ontologies it provides guidance for query formulation when a user looks for particular services. In this mode, the discovery of services is actually an iterative process

---

[1] Local Instance means that both the KB Service Client (e.g. the Ontology Editor) and the Service Instance are running on the same host.

through which the user is guided to form the query request that expresses his/her needs in the best possible way. After the query formulation the user submits the query and gets a ranked set of results (Published Service Manifests). Beyond the interaction with the Semantic Registry Service, this tool needs to interact with the KB Service in order to retrieve (parts of) Business and Service Ontologies for guiding the user in formulating queries.

- PREFERENCES EDITOR: Using Service and Business Ontologies an SME describes its preferences for desired services. SMEs are also able to describe business requirements that must be satisfied by the providers of those services (e.g. business policies, accounting policies, etc.). The SME may ask explicitly for recommendations according to its preferences or may browse the recommendations that the Recommender Service had automatically produced. The Preferences Editor is a tool that provides a GUI and uses the KB Service for storing/retrieving and updating preferences, and accessing Ontologies. It also uses the Recommender Service in order to ask for explicit recommendations according to the defined preferences. At the middle-tier special management modules implement the advanced functionality that is required for the realization of the aforementioned behavior.
- MANUAL COMPOSER: The Manual Composer component provides a GUI through which the end-user can compose new services. It interacts with the Semantic Registry Service or explores the suggestions made by the Recommender Service in order to discover candidate constituent services and to publish the newly created (composite) services. Also it interacts with the KB Service in order to store/retrieve/update Orchestration Specifications and Service Manifests of Composite Services and to retrieve the Required Ontologies for performing the required Service Matchmaking.

4 Knowledge Management in P2P Environment

In the DBE environment a large number of SME nodes will be pooled together to share their resources, information and services while keeping themselves fully autonomous. The SMEs will share data that will have rich structure and semantics, and thus advanced mechanisms for querying and discovering knowledge from these data are needed.

Taking these into account above, the DBE knowledge base should be built upon an efficient P2P data management system. The tasks during the development of the DBE KB include conceptual modeling, exhaustive study of current P2P data sharing technologies and investigation of strategies for deploying data management systems over P2P networks.

The design of the DBE KB should take into account the benefits of the P2P networks that can be concluded in adaptation, self-organization, load balancing, fault tolerance and high availability. Despite these benefits the deployment of P2P networks can introduce in many cases performance and consistency problems. Also current P2P solutions do not support advanced mechanisms for data management.

The challenge in designing and developing the DBE KB is to enable the effective support of rich semantics, data transformation, data relationships, data constrains and complex queries across multiple sources in the DBE network. KB will have to place data in strategic locations in order to improve the performance of the retrieval mechanisms, since the data sharing of enormous amounts of data is useless without advanced search mechanisms that will allow SMEs to quickly locate and use the desired information.

Current service oriented architectures and standards almost completely ignore the modeling of the business environment including business models, business processes, business environment models, domain specific specializations, instantiations etc. In addition many of the service oriented standards proposed are too close to a centralized or "one-owner" environment, and not a real P2P environment. The objective of the DBE KB is to organize, store, and efficiently retrieve description metadata about businesses and services. The KB will have to handle a lot of metadata that follow different schemas, since a great variety of business and service models will exist. Thus, techniques and strategies for message and query routing in schema-based P2P systems (should be investigated and deployed([8],[9],[10] The KB is responsible for matching and accessing semantic data across the DBE network nodes by offering to the SME users a set of Core Services for transparently searching and accessing DBE Knowledge.

The search mechanism to be deployed determines the behavior of peers in three areas: Topology (how peers are connected to each other), Data placement (how data or metadata are distributed across the network of peers) and Message routing (how messages are propagated through the network) [2].

While designing the search mechanism and defining the peer behavior, the query language, the returned results and the autonomy of peers (e.g. an SME may wish to define what data wants to store, or share data only within a group of SMEs) should also be considered.

A P2P model can either be pure or it can be hybrid. In a pure model, there is no centralized server (e.g. Gnutella and Freenet). In a hybrid model, a server is approached first to obtain meta-information, such as the identity of the peer on which some information is stored, or to verify security credentials. From then on, the P2P communication is performed. Examples of a hybrid model include Napster, Groove, Aimster. There also intermediate solutions with Super Peers, such as KaZaa. Super Peers contain some of the information that others may not have. Other peers typically lookup information at Super Peers if they cannot find it otherwise.

Pure P2P systems tend to flood the network with query messages and the limited capabilities of some peers cause performance bottlenecks. So, pure P2P systems present efficiency weaknesses that make them unsuitable for the DBE. On the other hand hybrid systems perform better but they are not scalable and fault tolerant.

Super peer systems are trying to combine the two previous approaches to take advantage of the efficiency of a centralized search and the autonomy, load balancing and robustness of a distributed search. Super-peer based P2P infrastructures usually exploit a two-phase routing architecture, which queries first in the super-peer backbone, and then distributes them to the peers connected to the super peers. Super-peer routing is based on different kinds of indexing and routing tables [6], [7] The DBE KB architecture follows the super peer network paradigm for the efficiency benefits that it presents and its capability of taking advantage of the heterogeneity of the peers by assigning greater responsibility to those peers that are more capable to handle it [3]. However the choice of the super peer model does not solve all the problems that the DBE KB will face, since the design should consider several of challenging issues like: dynamic self-organization of peers and super peers, performance trade offs, load-balancing among equivalent peers and among simple peers and super peers, avoidance of single-point of failure in the super peers, search performance using super peers, data placement and indexing across super peers and other research issues.

In the DBE environment many IT systems will have to be integrated in a common network. These systems will be developed, managed and integrated using a range of methodologies, tools and middleware. We have witnessed during the last few years, especially as a result of efforts at OMG and W3C a gradual move to more complete semantic models as well as data representation interchange standards. OMG's Model Drivel Architecture (MDA) provides an open, vendor-neutral approach to the challenge of interoperability, building upon and leveraging the value of OMG's established modeling standards: Unified Modeling Language (UML), Meta-Object Facility (MOF), XML Metadata Interchange XMI etc. [12,13,14].

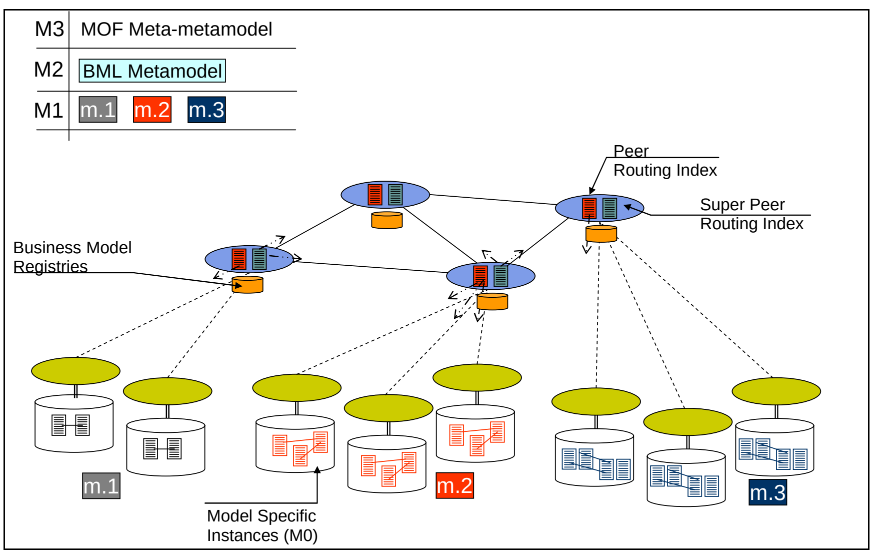

Figure 2: The DBE KB P2P Infrastructure

One of the most challenging objectives for the DBE is to enable common business understanding and at the same time to provide for differentiation between businesses and organisations. From a technical point of view, this means that while each SME may have its own model for service or business description, at the same time it needs to communicate with other SMEs and to understand each other. We have decided to follow a multi-layered approach in the Knowledge Modelling and Representation where each level in a top-down fashion adds more semantics into knowledge modelling concepts through specialization. For that, we are adhering to the rules of the MDA technology [12]. The top layer (M3) in MDA is the MOF meta-meta-model [13] that provides a mechanism for defining meta-models (meta-meta-data) which are actually residing at the M2 layer. M2 meta-models are used to provide specific models (meta-data) of representing knowledge (data). Thus the various types of content included in the KB are represented via XMI [14] so that multiple participants can query the stored content in a standardized and consistent manner. The KB Manager is responsible for the implementation of the appropriate structures for managing contents described with XMI. The manager uses the metamodel descriptions to determine how to manage contents that are actually instances of these meta-models. Figure 2 shows the DBE approach to metadata layering.

The P2P Knowledge Base infrastructure is shown in Figure 3. Three instances of core DBE services run at each SME node that belongs to the DBE network: the Recommender Service (RS), the Semantic Registry Service (SRS), and the KB Service (KBS). These services cooperate with the other instances in the peer network to perform the distributed knowledge update or retrieval tasks. These core services export a transparent common interface to DBE applications that need access in the DBE KB and hide the complexity of the necessary P2P communications between the service instances.

As we have described above, the distributed KB will follow the super peer network paradigm. Some SME peers that have enough bandwidth and processing capabilities will have more responsibilities than other SME peers. These SMEs will also run the three core service instances (of the KB infrastructure) that will behave differently than the service instances of the simple SME peers. The SME super peers will form a network among them and will be aware of the topology. These super peers act as servers to a group of SMEs are accepting recommendation, knowledge extraction and service discovery requests from their clients and based on the information distribution strategy that will be applied, forward the requests to the appropriate super peers that are able to satisfy or can reference to simple nodes that can satisfy the submitted requests.

For example the Service Publisher application (as shown in Figure 3) publishes the Service Manifest Advertisement (probably a summary of the Service Manifest that contains the most valuable information of a Service Description) accessing the local instance of the Semantic Registry Service. The local instance communicates with the respective service instance of the super peer that the SME is connected to and sends the advertisement. The super peer Service processes the Service Manifest Advertisement and based on the distribution strategy either it stores the service description in the local registry node (data management system) or it forwards the advertisement to the appropriate super peer Semantic Registry Service where it should be stored. The KB service acts in the same way as the Semantic Registry service. For example, BML descriptions produced by the BML Editor are pushed to the local KB-Service. This Service processes the description and forwards the public description to the super peer KB-service.

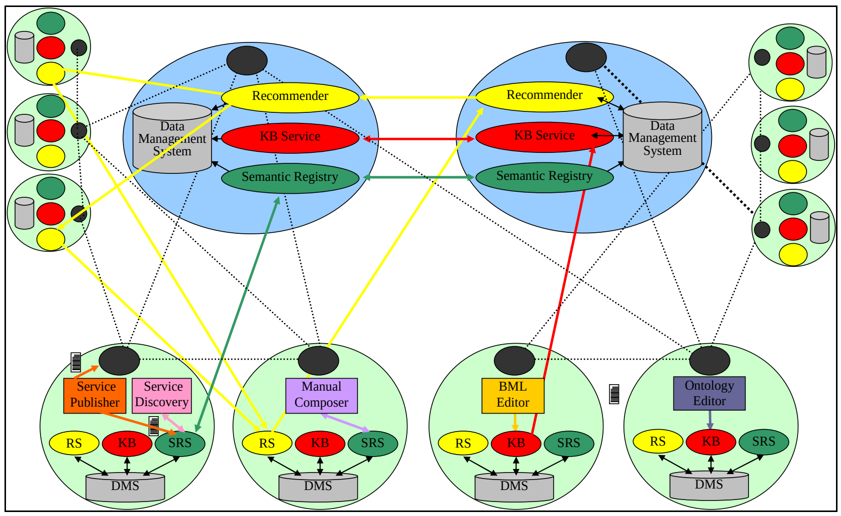

Figure 3: The Core Service Components of the DBE KB Infrastructure

The super peer KB service identifies the appropriate super peer that will handle the storage of the public information to the data management systems contributed by other SME peers. On the other hand the Recommender service acts autonomously. Periodically the local Recommender Service Instance asks the super peer Recommender Service for new system recommendations. The super peer Recommender service based on declared or automatically produced SME preferences in cooperation with other super peers identifies a group of SMEs that might contain information or services that should be recommended to the SME. The recommendation request is forwarded then to the Recommender Service instances of the qualified SMEs where is being processed and the recommendation results are sent (in P2P manner) back to the SME Recommender Service that requested them.

5 Conclusions and Open Issues

The previously described infrastructure for knowledge management will be an ontology-based P2P meta-data management system that its architecture can be used in many different environments. Some of the research issues related to the P2P Knowledge Management that will be examined during the project are the following:

- Ontology management (insertion, maintenance, conflict resolution and utilization) in P2P systems with the use of Relational Databases at each peer following a Service Oriented Architecture.
- Business model ontologies, business process ontologies, environmental ontologies, domain specific ontologies and their use and interplay in a dynamic service environment.
- Distributed Semantic Recommendation and Service Composition mechanisms

- Self-organization of the P2P network.

6 Acknowledgements

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

7 References

1. J.F Moore, The Death of Competition, 1996, pag.6-7
2. Neil Daswani, Hector Garcia-Molina, Beverly Yang "Open Problems in Data-Sharing Peer-to-Peer Systems" International Conference on Database Theory, ICDT 2003
3. Beverly Yang, Hector Garcia-Molina "Designing a Super-Peer Network" IEEE International Conference on Data Engineering 2003
4. Steven Gribble, Alon Halevy, Zachary Ives, Maya Rodrig, Dan Suciu. "What Can Databases Do for Peer-to-Peer?" WebDB Workshop on Databases and the Web, 2001.
5. W. Nejdl, W. Siberski, M. Sintek" Design Issues and Challenges for RDF- and Schema-Based Peer-to-Peer Systems" Technical Report, July 2003
6. B.Yang, H. Garcia-Molina "Improving search in peer-to-peer systems". Proceedings of the 22nd International Conference on Distributed Computing Systems 2002
7. Crespo, H. Garcia-Molina "Routing indices for peer-to-peer systems". Proceedings of the International Conference on Distributed Computing Systems 2002
8. W. Nejdl, B. Wolf, C. Qu, S. Decker, M. Sintek, A. Naeve, M. Nilsson. M. Palmer, T. Risch "EDUTELLA: A P2P Networking Infrastructure Based on RDF". Proceedings of the International World Wide Web Conference 2002
9. W. Nejdl, M.Wolpers, W.Siberski, C. Schmitz, M. Schlosser, I. Brunkhorst, A. Loser "Super-peer-based routing and clustering strategies for RDF-based peer-to-peer networks. Proceedings of the International World Wide Web Conference 2003
10. M.Schlosser, M.Sintek, S. Decker, W. Nejdl "HyperCup – Hypercubes, Ontologies and Efficient Search on P2P Networks. International Workshop on Agents and Peer-to-Peer Computing 2002
11. OWL Services Coalition: OWL-S:Semantic Mark-Up for Web Services", http://www.daml.org/services/owl-s/1.0/owl-s.pdf
12. Object Management Group: "Model Driven Architecture a Technical Perspective", http://www.omg.org/cgi-bin/apps/doc?ormsc/01-07-01.pdf
13. Object Management Group: "Meta Object Facility Specification", http://www.omg.org/cgibin/apps/doc?formal/02-04-03.pdf
14. Object Management Group: "XML Metadata Interchange Specification", http://www.omg.org/technology/documents/formal/xmi.htm
15. Java Community Process "Java Metadata Interface Specification", http://www.jcp.org/aboutJava/communityprocess/review/jsr040/"