

Competence Management Within and Between Organizations

Mikko Laukkanen and Heikki Helin

TeliaSonera Finland

P.O. Box 970 (Teollisuuskatu 13), FIN-00051 Sonera, Finland

{mikko.laukkanen,heikki.j.helin}@teliasonera.com

Abstract. We introduce a Semantic Web-based competence manager system called Competence Manager, which allows managing employee competencies and other relevant resources about a set of topics. The Competence Manager provides flexible operations for both querying the knowledge base, and managing the ontologies. The usage scenarios of the Competence Manager range from intra-organization use to interoperability between multiple organizations.

1 Introduction

For companies with intellectual property it is crucial to have an environment, where the knowledge can be shared efficiently within the company. In this paper we describe an ontology-based tool called Competence Manager, with which the members of an organization (e.g., a company) are able to manage, publish, disseminate, and search for competencies and other relevant resources, such as documents, projects and customers.

Our system allows supporting multiple ontologies for different domains, which can be imported and switched at runtime. In addition to maintaining multiple simultaneous domain ontologies, the interoperability—for instance from one organization to another—does not require changes to the implementation.

The difference between the Competence Manager and other typical competence management tools is that in the case where the best candidates (i.e., persons having competence directly on the searched topic) are not available, the candidates having competence on semantically nearby topic(s) can be found. This means that we can look up persons that are the best possible candidates for the searched competence area. In addition, we can search for the best candidate(s) from several domains with different emphasizes (e.g., candidates that have excellent competence on ontologies and can speak Finnish, but the language criterion with lower priority).

The benefits of the Competence Manager are clear for companies with intellectual properties. In addition, because of formally defined semantics of the domain ontologies, the system is useful when disseminating competencies between organizations. In doing so, the Competence Manager can be used as tool for finding organizations based on their real competencies.

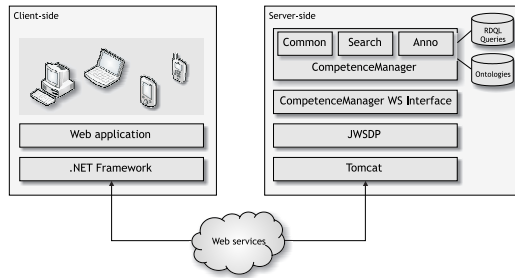


Fig. 1. Competence Manager architecture

2 Architecture and Ontologies

In this section we shortly outline the intra-organizational usage of the Competence Manager. The details are described in [2]. The architecture is depicted in Figure 1. The clients access the system through a front-end server, which connects to the back-end by using Web services. The requests are then handled by the Competence Manager, which is implemented on top of Jena [3].

Competence Manager employs OWL for the ontologies, which are divided into general and domain-specific parts, see Figures 2 and 3, respectively. The general ontology defines the mandatory classes and properties for the competence topic hierarchy. The domain ontologies define the domain-specific resources, and must conform to the general ontology. This asserts that the query engine works correctly with the domain ontologies.

The topic hierarchy in the domain-specific ontology is realized as instances of the **Topic** class, and parent-child relations are defined using the **hasParentTopic** property. The instances of the **Relevance** class specify the relevance level of some resource (e.g., a person's competence) to a certain topic.

The domain ontology defines the classes and their instances that have competencies or are related to the topics. The resources can be freely defined as domain-specific classes, as long as they are subclasses of the **Resource** class. Furthermore, the relationships (properties) between the domain-specific resources can be freely defined, but the properties must be defined as a subproperty of the **hasResource** property.

In a case where direct matches are not found, the people with competence on semantically nearby topic(s) can be searched for. This is crucial today, when for instance new products hit the markets rapidly. Let us consider the following example. At a software company C_s a person has competence on a mobile phone model α , which is manufactured by phone vendor C_v . Then, C_v releases a new mobile phone β , which is based on the model α , but on which nobody at C_s has yet competence. However, at some point when competencies for the new mobile phone β are needed at C_s , in our approach the persons having competence on the phone model α are found. Although this is not an exact match, it is a very close one, because usually the given phones are similar in functionality. Therefore,




Fig. 2. General ontology

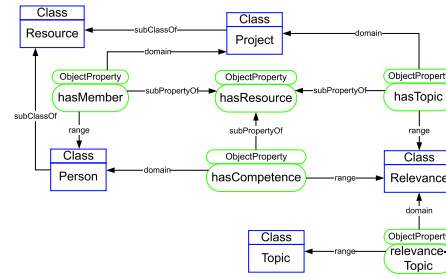


Fig. 3. Domain ontology

adding or changing the competence area structure does not necessarily require corresponding updates to the competencies of the persons.

The distance calculation between competence topics is based on formally defined semantics and an algorithm called Object Match [4], which calculates the distance to the neighboring topics, and returns the resources related to the closest ones. One of the most powerful features of the object match algorithm lies in its ability to take into account the depth of the object hierarchy. That is, the objects deeper in the hierarchy are also semantically closer to each other.

3 Using Competence Manager Within and Between Organizations

Within a single organization, the Competence Manager is used for managing and disseminating employee competencies. Firstly, competent persons can be searched for within the organization. Secondly, the Competence Manager can be used as a tool during performance appraisal between superiors and their employees. Thirdly, the Competence Manager can serve as a tool in strategic coordination activities. The overall competence levels on various areas can be analyzed, and in defining company's strategic goal, the areas needing more or less competence can be identified. As a result, the aggregated competencies of the organization can be defined, and made available for other organizations. In doing so, the Competence Manager can be used as tool for finding organizations based on their real competencies.

Usually when companies are formed, they have to be registered to some national authority using for instance an industrial classification schemas. However, especially for small companies, it may be hard to find a suitable classification for the company. Furthermore, if the company refocuses itself at some point later, changing the classification is cumbersome. This may lead to a situation, where the company has competencies that are not directly related to their registered core business. The Competence Manager can be used to overcome this problem by letting the companies publish (and update) their competencies at any time. Therefore, the Competence Manager acts as a directory service—Competence Di-

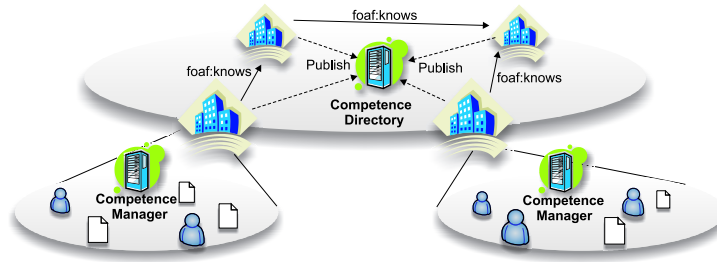


Fig. 4. Competence Directory aggregating multiple organizations

rectory, from which companies can search for instance possible partners based on their competencies, see Figure 4.

Because companies may use their own domain ontologies, they cannot be uploaded as such to the Competence Directory. Instead, the companies publish their domain ontology (topics hierarchy) and other resources—competencies, documents, projects—using FOAF [1]. The amount of information published depends on the company, but at least the competencies have to be uploaded to the Competence Directory. If Competence Manager is used within a company, it can be used to automatically export the domain ontology and selected resources as FOAF description to the Competence Directory.

Trust plays an essential role in systems like Competence Directory, where companies may state their own competencies. The essential challenge is to assure that the information published is accurate. We believe that in this kind of collaborative systems, the abusers will eventually be exposed. Reputation among other companies is very important, and it would be unwise to ruin it by publishing false competencies. Further, the `foaf:knows` property can be used to define one-way links between companies. For instance, a company C_s is searching for a partner, and the search results multiple companies, C_{r1} , C_{r2} , and C_{r3} . The C_s may have an existing and successful partnership with a company C_p , that is, C_s `foaf:knows` C_p . Now, if C_p has an existing partnership with some of the companies in the search result, they are certainly preferred over the ones that the C_s knows nothing about.

References

1. D. Brickley and L. Miller. FOAF Vocabulary Specification. Apr. 2005. Available at: <http://xmlns.com/foaf/0.1/>.
2. M. Laukkanen and H. Helin. Using Semantic Web Technologies in Competence Management. In *XML-Finland 2005*, Pori, Finland, Mar. 2005.
3. B. McBride, D. Boothby, and C. Dollin. An Introduction to RDF and the Jena RDF API. Nov. 2004. Available at: http://jena.sourceforge.net/tutorial/RDF_API/index.html.
4. N. Stojanovic, A. Maedche, S. Staab, R. Studer, and Y. Sure. SEAL: A Framework for Developing SEMantic PortALs. In *Proceedings of the International Conference on Knowledge Capture*, pages 155–162, Oct. 2001.