Enabling a Semantic Wiki to Drive Business Interactions

Sven Graupner, Sharad Singhal, Sujoy Basu, Hamid Motahari

HP Laboratories
HPL-2009-193

Keyword(s):
business interactions, business process, best practice framework, outsourcing, cloud, services, knowledge repository

Abstract:
This paper describes a Semantic Wiki which enables business interactions among people who are collaborating in the context of business activities. The Wiki has domain knowledge (e.g. about ITIL) and defines and implements domain independent semantics which allows it to drive related business interactions among people. We believe that such an environment can be used to capture best-practices that can be customized quickly, and evolved to meet specific business needs, while retaining the repeatability and work efficiency necessary within business contexts.

External Posting Date: September 6, 2009 [Fulltext] Approved for External Publication
Internal Posting Date: September 6, 2009 [Fulltext]

Enabling a Semantic Wiki to Drive Business Interactions

Sven Graupner, Sharad Singhal, Sujoy Basu, Hamid Motahari

OST/HP Labs

{sven.graupner, sharad.singhal, sujoy.basu, hamid.motahari}@hp.com

Abstract

This paper describes a Semantic Wiki which enables business interactions among people who are collaborating in the context of business activities. The Wiki has domain knowledge (e.g. about ITIL) and defines and implements domain-independent semantics which allows it to drive related business interactions among people. We believe that such an environment can be used to capture best-practices that can be customized quickly, and evolved to meet specific business needs, while retaining the repeatability and work efficiency necessary within business contexts.

Introduction and Problem Statement

Business support systems such as SAP organize business processes in a very strict way in their implementation. While certain activities require this rigidity, most business interactions occur informally as conversations, coordination, and exchange of information among people. Furthermore, human interactions occur within different contexts, which can range in complexity from an ad-hoc purchase of an item seen in an advertisement to establishing a formal outsourcing contract between two large enterprises. This context provides information to people about the importance of different tasks that need to be performed, as well as enabling them to quickly adapt what they do in order to meet those requirements. While information technology (IT) systems such as SAP are good at achieving predictable behavior for repetitive mechanized processes, they do not provide the flexibility or agility required for business activities that largely rely on human interactions.

Although human interactions have few constraints, within an enterprise setting they are not random. Best practices have been captured in a variety of frameworks to guide business activity. Frameworks such as the IT Information Library (ITIL) [1] capture practices that are widely applied by HP in customer environments for defining and managing IT processes. However, ITIL largely defines issues that need to be addressed by people, and does not provide specific detailed processes for executing actions. Because these frameworks are generic and are usually interpreted and customized before being used, they serve as best-practice guides for business and IT processes, but are difficult to implement in traditional business support systems.

Many tools have been developed to capture plans and organize projects; estimate budgets and spending; support and manage schedules [2]. Other tools can support virtual teams working on shared projects [3]; manage and organize documents and their workflows such as MS SharePoint [4], [5] to name some. However, these tools lack awareness of the best-practices that need to be followed. For example, a document management system can manage documents and workflows around them, but it does not estimate budgets or organize schedules. Sharepoints or Wikis support creation of workspaces that allow documents related to a given task be shared, but lack any information about the context within which those documents are being created, or the interdependencies between those documents. It is commonly observed that collaboration environments such as Sharepoint or Wiki quickly turn into wastelands where it is hard to maintain oversight, with no context remaining about which documents are important, or even useful, which ones need to be archived, and which ones are current.

Consultants typically use a variety of mechanisms ranging from plain-text documents and spreadsheets to formal specifications using process modeling tools [6], [7] to define how work within a business setting can be organized according to best practice frameworks. However, since the actual people who are doing the tasks are not "process experts," the specifications are not always followed, and chaos ensues. Furthermore, since the context within which these specifications are created rarely stays constant, such specifications quickly become outdated, and become difficult to reuse.

On the other hand, developing a single tool that can organize all the information spread across the plethora of

tools and systems used by groups of people in a business context seems to be a hopeless endeavor. No technology claims to have solved this problem. People often feel lost in all their tools, and spend considerable time feeding data from one tool to another. Tools often have difficulty answering simple questions, beginning with "why," because the relevant information resides elsewhere, often in the head of some person other than the tool user.

In this paper we address the question of how business activity among *people* can be better supported and guided using modeling frameworks and collaboration environments. We specifically consider how a collaboration environment that is aware of the context within which people work can be constructed, and how it can be used to both organize the information necessary, as well as to orchestrate business activities. We identify methods that can be used to incorporate domain frameworks within that environment. Finally, using an example outsourcing scenario, we consider the business implications of such an environment.

Our Solution

We observe that collaboration environments [8] and content management systems [5], [10] are often used to manage content and documentation related to any business task. Almost all tasks that require human interaction also require that information (documentation) be created and shared among the participants involved in the tasks.

We hypothesize that if the collaboration environments could be augmented with knowledge about best-practice frameworks, and be enabled not only to organize task-related content based on the frameworks, but also to direct and guide the human activities occurring around that content in accordance with the frameworks, then those activities could be made significantly more efficient, while retaining both the flexibility and agility needed by the business. This would enable the business to follow best practice rules and guidelines in a repeatable manner. It would also enable rapid sharing of knowledge and experience by capturing best practices within the business as templates constructed using such environments.

Solution Architecture

In this section, we present concepts, the major building blocks and operational principles that enable the realization of such a system. We optimistically call our solution a ***Business Operating Environment*** (BOE).

The overall conceptual architecture of the BOE is shown in Figure 1. The solution is loosely organized around three building blocks: a) an information base, b) an activity manager, and c) an interactive work environment that supports collaboration between people.

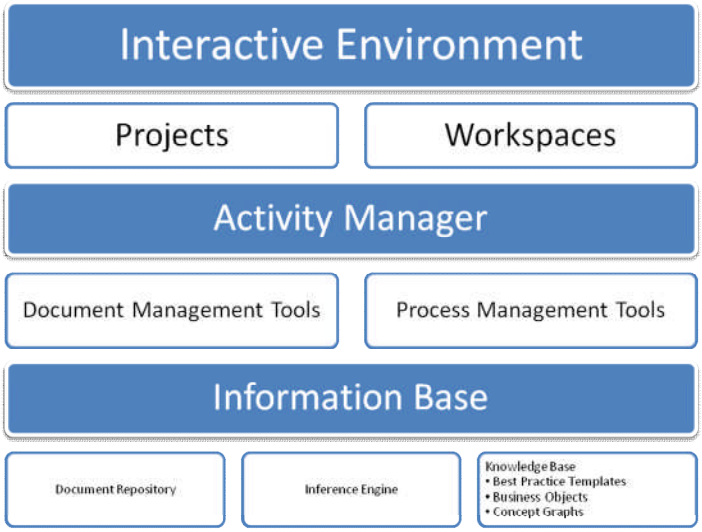

Figure 1: Overall Structure of the Business Operating Environment

The Information Base

Underlying the BOE are one or more information bases. Each information base includes knowledge bases that provide "context" information to the BOE. Business frameworks such as Zachman [11], OMG BMM [12] and

ITIL [1] are used to create best-practice templates within the knowledge base, and are used to organize the information within the BOE. Individual information elements are abstracted as business objects within the knowledge base, while relationships between the objects are captured as concept graphs. An inference engine is used to provide query and inference capability over the templates.

In addition, the information base includes repositories containing the documents being managed within the context of that information base. The business objects hold references to these documents. This enables the inference engine to reason about dependencies between the documents, to create and delete dependencies between documents as the context requires, and to organize the documents within the context of the domain framework.

The Activity Manager

The activity manager facilitates the flow of activities between people. It guides the evolution of information being held in the information base in accordance with the best practice templates. The templates are configurable to account for new contexts and projects. Thus, rather than only relying upon people to properly generate or manipulate the required documents, the activity manager orchestrates the flow of activities between people in accordance with the domain templates. It also allows the definition and execution of ad-hoc activities. In addition, it can link into domain-specific tools that augment, correlate, and in some cases generate documents intended for the people working in the BOE.

The Interactive Environment

The interactive environment provides collaboration tools that enable people to interact among themselves, as well as with underlying systems. The underlying activity manager guides these interactions, provides the relevant information to people, manages the documents visible to them, and supports them as they perform their activities. It is organized around two concepts: *projects* and *workspaces*.

Projects represent short-lived groupings of people and activities for a business purpose. Within the information base, projects may be created by creating the relevant templates and documents in the information base. Projects may be also created by copying the content from previous projects. The templates may allow interdependencies to exist within projects. The activity manager treats each project independently of others.

Workspaces represent *views* into the projects from the perspective of an individual or group of individuals. Thus all information relevant to a project (or individual) can be presented within a workspace set up for that project (or individual). Unlike views offered by standard content or document management systems, workspaces show items that are actionable by the workspace owner—the activity manager changes the workspace as tasks are assigned, acted upon, or completed by the workspace owner or others.

Solution Realization

Clearly, many existing collaboration, process management, or knowledge management tools can be integrated to realize the above architecture. In our research, we have built our solution around a popular collaboration environment: Wiki, which is commonly used today by a growing community of people in businesses to exchange and share information. In particular, we extend Wikis to link collaboration, document management, and knowledge management capabilities using concepts from the Semantic Web [13].

Knowledge base

The core innovation in our solution is the ability to drive the interactions between people within the collaboration environment using a knowledge base. We initially populate the knowledge base with the following concepts:

- *roles* and *responsibilities* of people participating in the collaboration. This enables the underlying Wiki environment to recognize the relationships and dependencies between a given user and other users of the collaboration environment.
- *business activities* of interest within the domain. These allow the documents and information being managed within the Wiki to be linked and organized in accordance with the underlying domain framework, and allow workflows to be created between the collaborating participants.
- *work items* to be accomplished (e.g. a ROI analysis or a risk assessment for a service engagement). These items may be assigned by people, as in a traditional Wiki, or may be generated by the underlying activity

manager as a consequence of work being done by someone else.

- *document annotations* to capture information relevant for linking business tasks. Unlike standard Wikis, where all content is generated by the people, we add annotated document templates that are used by participants. The underlying tools can extract tagged information from the annotations in the documents, enabling the activity manager to automatically update (and in some instances, generate) documents being held elsewhere in the wiki.
- *events* as means to coordinate activities. As documents are updated or changed within the wiki, events are generated to enable communication with underlying tools or with people around the Wiki. In particular, we use *time*-based events to control progression of tasks within the wiki.

These concepts are represented within the knowledge base in an RDF [14] store and drive functionality implemented in the back-end of the Wiki. For instance, when a document is uploaded to the Wiki by some person, a work item can be triggered to require approval of release for that document by another person. This work item then can be associated with a timer indicating a time line for approval. A request can then be routed to the approver with either a URL link to the document, or the corresponding document can be added to the approver's personal workspace for review. This toolbox of common, domain-independent concepts can be overlaid by *blueprints* or *templates* defining business activity for business domains such as ITIL, Six Sigma or others.

In our research, we have chosen ITIL as a basis for experimentation to assess the effort required to create and maintain domain-specific templates. In order to establish an ontology, we manually examined ~20 pages of ITIL documentation and created a mind map reflecting the most important concepts in a particular section of the ITIL Volume on Service Design dealing with "Supplier Management". A part of the graph, which contained 524 nodes, is shown in Figure 2.

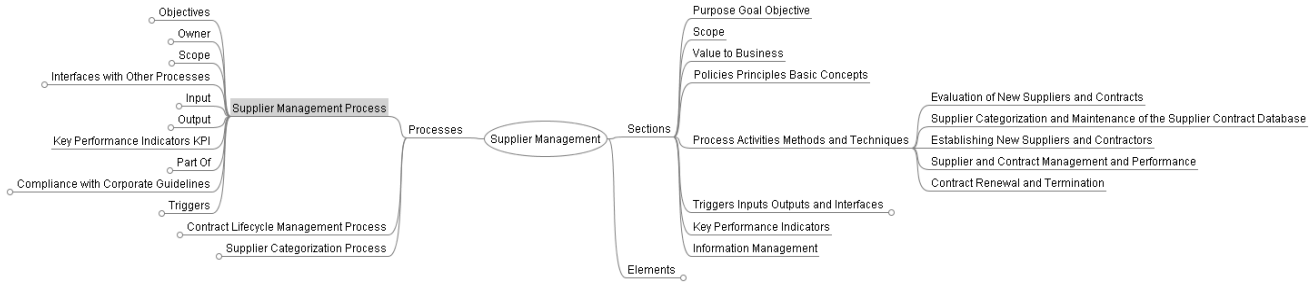

Figure 2: Segment of domain knowledge about Supplier Management obtained from the ITIL Service Design document.

The concept graph was then programmatically converted into RDF. For example, the *SupplierManagement* node above produced the following RDF statements:

```
:SupplierManagementProcess :a :MMConcept ; :conceptOf :Processes; :name   "Supplier Management Process";
   :has :Triggers, :Output, :CompliancewithCorporateGuidelines, :InterfaceswithOtherProcesses, :PartOf, :Scope, :Owner, :KeyPerformanceIndicatorsKPI, :Objectives, :Input ;
   :mmid   "Freemind_Link_1343844995" .
```

This concept graph forms the backbone of knowledge of the Wiki, and is loaded into its RDF model store. The Semantic Wiki permits execution of SPARQL queries on the graph. To support inferences over the graph, we use the Jena [15] toolkit.

Automated knowledge base creation

The complexity of this manual task (interpreting only 20 pages from about 1350 pages of text) convinced us to investigate ways of automatically extracting concept information from documentation. Under normal circumstances, we expect that domain experts will refine and extend the concepts and relationships over time as necessary. Thus, we are not concerned with extracting every concept that a domain expert would label as important. Additionally, if the knowledge base does contain some concepts deemed "unimportant" by users, those concepts will simply not be used by people in their work. Thus it is important that any automated technique capture most of the important concepts (high recall), but it is acceptable if the technique also generate some concepts that are deemed irrelevant by experts (low precision).

Again, we used the ITIL documentation for our experiments. We divided the documents into their sections, and treated each as an independent document. Next, we applied an information-retrieval (IR) technique known as Latent Semantic Indexing (LSI) [16] to this set of documents. Traditionally, LSI has been used in IR tasks such as finding matching documents for query terms or similar documents for a query document. If terms such 'car' and 'automobile' have been used synonymously in the set of documents, a query containing 'car' will also retrieve documents containing 'automobile'. In our work, we deviate from LSI's goal of finding results for a query. Instead, we employ the underlying mathematical technique, the Singular Value Decomposition (SVD) [17], to find important terms (concepts) and their relationships across the set of documents. The technique allows us to assign a numerical value indicating how strongly two terms are correlated across the set of documents, and rank-order the term pairs based on this value.

We seeded the technique with 385 known concepts present in the ITIL glossary. Each complex concept was treated as a single term. The output of the SVD then generated rank-ordered term pairs for the document set. For each term pair (e.g., 'service' and 'management'), we looked for trigrams that consisted of the term pair with another term in between. Thus if 'continuity' occurred frequently in the documents with 'service' and 'management', then 'service continuity management' was produced as a trigram. This allowed us to create more complex concepts (that were also rank-ordered). Based on a threshold, we added many of the important complex concepts discovered to our list of known concepts. The augmented list served as the next seed, and the process repeated. Relationships between concepts were identified using part-of-speech information obtained from an online dictionary for the associated word. Thus, if two concepts (nouns) were accompanied by an action (verb) in the trigram, the action was labeled as a relationship between the concepts as opposed to a complex concept. This allowed us to automatically extract both the important concepts, as well as relationships between them.

For the ITIL documentation, the first iteration of our technique produced 2072 rank-ordered trigrams. Starting at the top of this list, the trigrams were 'service level management', 'service catalogue management', and 'service capacity management' which appear to be relevant concepts. Starting at the bottom of this list, we find the trigrams 'dependency defining countermeasure', 'stakeholder exception estimation', and 'help optimum throughput', which are not meaningful concepts or relations. For the second iteration, we included all trigrams up to rank 2000. While this arbitrary threshold included some trigrams that did not make sense, the second iteration again brought the important concepts to the top and eliminated some of the trigrams that did not make sense.

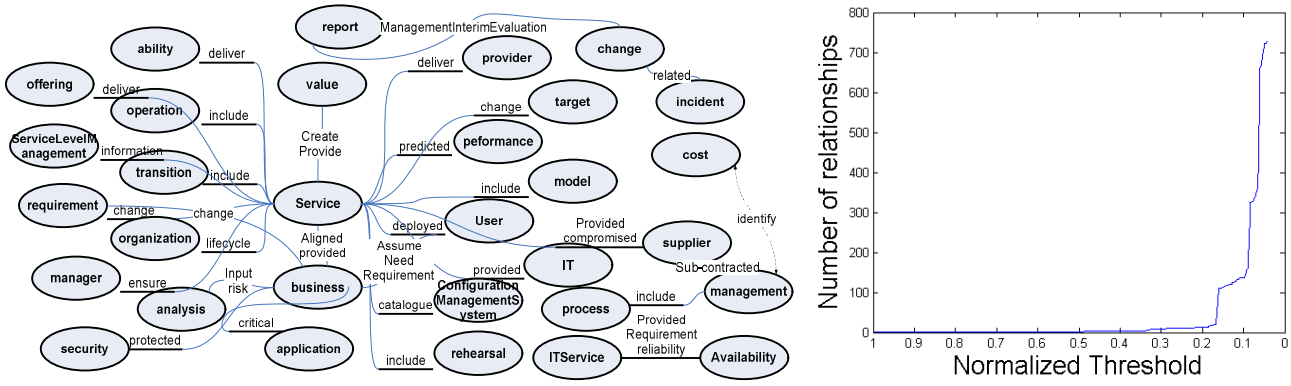

Figure 3: Automatically extracted concepts and relationships from ITIL documents. As the normalized threshold (with 1 corresponding to highest value) is reduced, additional concepts and relationships are created.

Figure 3 shows a snapshot from the second iteration. In this graph on the left, some edges have multiple labels. For example, the edge between 'service' and 'supplier' has the labels 'provided' and 'comprised'. It indicates that there were two trigrams 'service provided [by] supplier' and 'service comprised [of] supplier[s]'. Note that the Figure only shows a small fraction of the concept graph to keep it legible. The plot on the right in Figure 1 shows that a total of 727 relationships were discovered in this iteration. There are several knees in this plot, indicating that concept relationships are picked up in clusters as the threshold for selection is varied. The technique discovered 20 relations above a threshold of 0.17, with larger clusters being created as the selection threshold was reduced further.

Note that it is not always the case that the middle term in the trigram is the relationship between the remaining two concept terms. This ambiguity is due to different ways in which natural-language sentences can be constructed. By incorporating parts-of-speech tags generated by a natural-language parser, we were able to disambiguate relationships between concepts from compound. For example, if a verb and two nouns make a trigram, the verb describes the relationship irrespective of its position in the trigram.

Using the knowledge base for driving the environment

To provide users with the ability to organize documents using this framework, as an initial step, Wiki pages were auto-generated for each node of the graph with its properties and link structure to other pages. Figure 4 shows one of the generated pages, containing the navigation structure around the given node. This enabled anyone using the Wiki to browse the concept graph as regular Wiki pages, and upload and edit documents relevant to the corresponding concept.

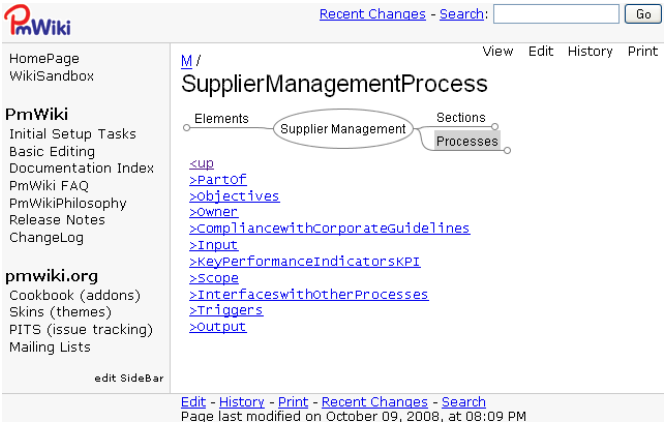

Figure 4: Auto-generated page for one node of the supplier management process

Thus for example, suppose a project to select a supplier is created. The tools generate all relevant pages for each relevant concept as part of project creation as shown in Figure 4. They also populate the pages with document templates that have been created by domain experts from the document repository. This enables people working on the project to quickly navigate to, for instance, document templates related to service level agreements (e.g., definitions of key performance indicators), retrieve them or update them. Because the documents stay with the pages corresponding to the concept, the knowledge base can use them to drive other tasks. In addition, because users can create other pages and link them to existing pages; upload new content onto the pages, or modify links between the pages, it is easy for them to customize the project to their own needs. This allows the project team members to rapidly organize project information aligned with the domain template, rather than using arbitrary hierarchies. Because the relevant documents are automatically organized within the collaboration environment, project information created by one group of people stays accessible to others who may not otherwise know about it, increasing efficiency within the project.

Furthermore, as these changes are made, the activity manager captures updates made to the Wiki structure within the RDF graph representing the domain concepts, thus evolving the knowledge base as the users work with it. Additional templates may be created from documents being uploaded by the users. This enables the knowledge base to expand and become customized as users work with it, rather than requiring domain experts at each stage of modification.

Finally, it becomes straightforward to duplicate the entire project hierarchy for a new project that is "similar" to the previous one, enhancing the ability to reuse documents and templates across projects, mine information across related projects, or evolve the best practice templates across projects.

Our solution in practice: An illustrative example

We were recently made aware of an example which can provide an illustration for our approach within HP's current practice. We stress that we are simply using this as an illustration for our approach—we do not imply that

we have implemented this example. We consider HP's outsourcing business within EDS. In an effort to streamline HP's services offerings, HP has developed a *Transformation Services Catalog* [18] as a repository of guidelines for engaging and managing outsourced services transformation engagements with customers. The catalog is structured by a number of project categories such as consolidation projects, network optimization projects or process projects. Each category has sub-project categories: for example the process category has sub-categories of change management, configuration management, incident management, problem management and SLA early warning. Each sub-category contains a listing of template documents to be used by engagement managers working on a particular outsourcing activity. Templates provide detailed guidance and assistance to scope the engagement, to develop a cost and schedule model, deployment guide, precedence diagrams, process guide, task check list and technical documentation. The templates for the various project categories follow a common pattern for tasks and thus represent a model of a best practice under which outsourcing engagements should be conducted in HP. During an engagement, these documents can provide a "living" environment documenting the assumptions and status of the engagement.

Currently, the Transformation Services Catalog is used by engagement managers as a project guide. The manager downloads the set of documents from an appropriate category from the catalog and uses them for a particular engagement. The content of the downloaded documents forms the foundation for engagement. The Transformation Services Catalog is a web-based repository through which template documents are distributed. After distribution, no information automatically flows back into the catalog. Once documents are distributed from the catalog, they exist on their own. It is left up to the engagement manager to interpret the documents, organize the project teams, track the relevant tasks, and ensure that all documents are properly filled out as the engagement proceeds.

Under our approach, content for such ongoing activities would not only exist in the document templates (e.g., as Microsoft Office Project documents), they would also be captured in the knowledge base directly. When a new engagement is started, a new project workspace would be created using an appropriate category. This would create the project workspace for the engagement manager, and the appropriate documents and templates would be created for the engagement. If a particular template recommends that a customer workshop be conducted as part of requirements gathering for the engagement, pages would be created for the workshop, and populated with the information necessary to the staff responsible for running the workshop. The documents generated at the workshop would simply be uploaded into the same workspace, and become part of the repository for later use within the project, or to serve as examples for other projects of a similar nature. Action items associated with the workshop could be defined and propagated to individuals responsible for them. Time lines used within one project could be used by later teams who can browse documents from earlier engagements.

These initial sets of templates would be provided by the same domain experts who are currently providing the document templates in the Transformation Services Catalog. Conceivably all activities currently described in the template documents could be mapped into actionable templates represented in our knowledge base; however, our framework does not require completeness. In the simplest case, the current Transformation Services Catalog could be created as the repository from which templates are distributed as documents. In this case, only the information about the documents would be known to the knowledge base, but not the activities described in them. As people use the templates and modify the Wiki structure in their work, information can be captured by back-end tools to help factor out information that is useful for task orchestration and made part of the knowledge base.

Depending on how the initial templates in the knowledge base are modeled, they can guide the engagement manager (an initiating role) through the overall steps of the engagement, which is referred to as bootstrapping an activity. The starting steps include creating a project workspace and choosing a domain category for it. The domain category determines the set of initial material copied to the project workspace. Additional steps include adding people to the project and assigning them roles and tasks. This would trigger the creation of new workspaces for newly appointed individuals, eventually pre-populating their workspaces with material which is relevant for them. Initially, this activity could be done manually by the engagement manger. As the Wiki content is changed, back-end tools could capture these changes as "macro-recordings", which could then be examined by the domain experts to automate the tasks for engagement managers of later projects. Similarly, as work items are created by the managers and assigned to individuals, they would appear within the appropriate individuals' workspaces. Material that is produced during the processes could be uploaded by them into their workspace, and

automatically shared as necessary based on the relationships defined in the domain templates.

The ability to observe and record actions within workspaces, and enable creation of templates that augment existing templates within the knowledge base can provide a powerful mechanism for creating repeatable engagements and increasing operational efficiency. Since people receive and conduct their work through workspaces on our system, our system would be able to intercept work activities and mirror current status back into the knowledge base, avoiding additional effort for maintaining current information in the knowledge base. While in our solution, we still cannot ensure that content within documents is correct; we can observe if documents are acted upon and updated.

Since the initial set of domain templates is copied for the engagement workspace, changes can be made within that workspace by the people involved. The scope of those changes affects only the actual engagement in which they occur. Broader changes made by domain experts to the "master" templates will become effective for each newly started project, but will not affect existing ones.

This illustration highlights the following advantages of our approach over existing approaches:

- Procedures described in the set of templates documents can be mapped into actionable ontologies in the knowledge base and guide the targeted business activities.
- Work environments can be created in the form of workspaces for specific engagements through which people's activities could be coordinated. Documents and updates can be delivered to the people through their workspaces.
- Material produced as result of an engagement can be copied into a subsequent engagement and serve as examples to be used within the new engagement. Engagement managers may modify, store and re-use their own templates if they wish.
- The back-end knowledge base can maintain a true "current state" of the engagement, since work is organized and orchestrated within the workspace. The effort needed to track the engagement could be minimized by reducing additional data entry simply for tracking purposes.
- Domain experts responsible for maintaining the templates can retrospectively analyze engagements from the project data, and capture new practices or processes in the master templates, which would become available for the next engagement.

Under our approach, the Transformation Services Catalog could turn into a living and evolving working environment for HP's outsourcing practice.

Related Work

Many existing requirement capture and management tools [19] and business process analysis tools such as ProVison [6] simplify the tasks of gathering, documenting, tracking and managing requirements and process definitions in an enterprise. Typically these tools help document requirements and processes, and in some instances simulate the impact of changes. They are geared towards implementing and executing projects and processes in IT systems, not among people.

Business process management tools such ARIS or SAP [20] allow definition of rigidly structured business processes. However, many processes in the enterprise involve human interactions, and are ad-hoc or semi-structured. Our solution differs from business management systems in that it is much more "lightweight" and relies on supporting a collaborative environment rather than executing hard-coded processes in software.

Our solution differs from document management systems [4] such as Documentum [21] in that it is targeted at not just managing documents or document workflows, but business interactions between participants. A major hurdle to wide-spread adoption of knowledge management tools [22] today is the poor linkage between them and the surrounding human processes. Our work addresses this hurdle. The other key differentiator of our work from knowledge management systems or business process management tools is that our tool does not have to be used "after the fact." Often knowledge bases or process design tools are used to craft processes, but people rarely go back to update information in them. This reduces the ability of enterprises to "reuse" information and processes as a result of ad-hoc changes needed by people. Our approach differs from collaborative business process definition tools such as Caramba [3] in that the knowledge base in our environment is a "living and interactive" knowledge

base that evolves as people use it.

The proposed Wiki-based platform differs from Semantic Wikis (e.g. Semantic Media Wiki [9], IkeWiki [23], OntoWiki [24] or KaukoluWiki [25]) because we incorporate domain-independent semantics into the Wiki to organize people's activities by defining concepts in the RDF model, which relate to functionality implemented in the Wiki back end. We also differ in the ability to upload RDF models templates with domain knowledge into the Wiki.

Battle, Booth, Jahn, Lawson and Peltz [26] proposed a Semantic Wiki for ITIL. However, their work was restricted to providing query capabilities from RDF generated manually from the ITILv3 glossary. Our work extends that by both automatically expanding the generated RDF, auto-generating pages for each concept in the RDF model, and providing back-end capabilities that can make the Wiki environment interactive.

The social computing concepts, approaches and tooling for knowledge sharing and collaboration incorporated in the enterprise systems, collectively known as Enterprise 2.0 [27], are complementary to our work. Our platform benefits from leveraging concepts and methods from this area for knowledge capture and sharing between business users.

Status and Next Steps

We have defined the ontologies necessary for domain-independent concepts, and have incorporated them into our prototype. In addition, we plan to add support for the following additional concepts: *business objects* and *states, schedules, lifecycle, resources, allocations, tasks, subtasks, task lists, services, contracts, metrics, control points, processes* and *workflows*. Introduction of these concepts requires additional work to define concepts in the domain-independent RDF model, and to provide the corresponding functionality in the Wiki back-end. Our framework allows capturing knowledge of business activities as templates, organizing the activities using simple and intuitive structures (mind-maps), and making them available for reuse by other business users.

We used PmWiki [28] as the initial Wiki implementation platform since it allows easier integration of functionality in the back-end. We are currently integrating an event system and jBPM [29] engine into the Wiki to create the Activity Manager. However, it remains an open question whether the Wiki provides the optimal environment for mediating user interactions. Nevertheless, the focus of this work is on the fundamental principles and building blocks of providing an actionable knowledge base for conducting business activities among people. This relates to the identification of core concepts and their interpretation. The same concepts and methodology are applicable if another collaboration environment (such as Sharepoint) is selected.

For the longer-term, richer templates need to be created for other domains than ITIL. We hope to leverage the concept extraction technique. Furthermore, we need to gain experience how users would make use of such a system and define metrics to measure its usefulness.

References

[1] Information Technology Infrastructure Library (ITIL), Version 3, 5 Volumes, http://www.itil-officialsite.com.

[2] Wikipedia, List of Project Management Software, http://en.wikipedia.org/wiki/List_of_project_management_software.

[3] S. Dustdar, Caramba — A Process-Aware Collaboration System Supporting Ad hoc and Collaborative Processes in Virtual Teams. Distrib. Parallel Databases 15, 1 (Jan. 2004), 45-66.

[4] HCi Journal, List of Document Management Software, http://www.hci.com.au/hcisite3/journal/List of document management software.htm.

[5] Microsoft SharePoint, www.microsoft.com/Sharepoint/default.mspx.

[6] Metastorm, ProVision, http://www.metastorm.com.

[7] IDS Scheer, Aris, http://www.ids-scheer.com.

[8] Wiki, http://en.wikipedia.org/wiki/Wiki.

[9] University of Karlsruhe, Semantic Media Wiki, http://semantic-mediawiki.org.

[10] List of content management systems, http://en.wikipedia.org/wiki/List_of_content_management_systems

[11] Marc Lankhorst, Enterprise Architecture at Work: Modelling, Communication and Analysis, Springer, 2005.

[12] OMG. Business Motivation Model (BMM), www.omg.org/technology/documents/br_pm_spec_catalog.htm.

[13] W3C. Semantic Web FAQ. http://www.w3.org/2001/sw/SW-FAQ
[14] W3C: Resource Description Framework (RFD), http://www.w3.org/RDF.
[15] HP: Jena – A Semantic Web Framework for Java, http://jena.sourceforge.net.
[16] Dumais, S. T., Furnas, G. W., Landauer, T. K. and Deerwester, S., Using latent semantic analysis to improve information retrieval. In Proceedings of CHI'88: Conference on Human Factors in Computing, New York: ACM, 281-285, 1988.
[17] Furnas, G. W., Deerwester, S., Dumais, S. T., Landauer, T. K., Harshman, R.A., Streeter, L.A. and Lochbaum, K.E.: Information Retrieval using a Singular Value Decomposition Model of Latent Semantic Structure, In Proceedings of ACM SIGIR'88. Pages 465 – 480, 1988.
[18] HP Transformation Services Catalog. EDS/HP Services http://intranet.hp.com/tsg/ams3/services/001/Pages/TDS_catalog.aspx.
[19] Volere. List of Requirement Managemnt Tools. http://www.volere.co.uk/tools.htm.
[20] J.-N. Gillot (2008). The Complete Guide to Business Process Management. BookSurge Publishing, 2008.
[21] EMC. EMC Documentum, http://www.emc.com/products/category/subcategory/collaboration-and-document-management.htm.
[22] Philippe, M.: Knowledge Representation and Translation in RDF+OWL, N3, KIF, UML and the WebKB-2 languages, http://www.webkb.org/doc/model/comparisons.html.
[23] Ike Wiki, http://ikewiki.salzburgresearch.at.
[24] OntoWiki, http://ontowiki.net.
[25] KaukoluWiki, http://kaukoluwiki.opendfki.de.
[26] Battle, S., Booth, D., Jahn, K., Lawson, R., Peltz, C.: RDF, Ontologies and ITIL – Reducing the pain of IT information management, http://battle-s-1:8080/IkeWiki/index.jsp.
[27] T. Davenport Enterprise 2.0: The New, New Knowledge Management? Harvard Business Online, Feb. 19, 2008.
[28] PmWiki, http://www.pmwiki.org.
[29] JBoss jBPM. JBoss Workflow Engine, www.jboss.org/jbossjbpm.