Collaboration Environment for ITIL

Sven Graupner, Sujoy Basu, Sharad Singhal

Abstract:

This paper describes a knowledge-based collaboration environment for ITIL [1] which is aimed at making work of IT consultants dealing with planning, designing and implementing ITIL processes easier. The collaboration environment organizes work of consultants around built-in ITIL concepts and allows them to map and refine concepts for a particular customer engagement. The collaboration environment is designed as a so called domain wiki that has, in contrast to other wikis, a pre-populated knowledge base built-in, which actively guides further information gathering and refinement as collaborative processes by consultants. In this paper we present the design of a domain wiki for ITIL. Common ITIL knowledge is modeled from a subset of the ITIL publications and represented as an RDF knowledge base. Inferences on this knowledge base allow users to determine the aspects that need to be refined by consultants for particular engagements. Provided facts are then added to the knowledge base. In ongoing research we investigate how such a knowledge base can be established for the domain of ITIL. One approach we take is a manual knowledge engineering effort; another investigation is how text analysis techniques can be used to automatically learn facts from ITIL publications.

Collaboration Environment for ITIL

Sven Graupner, Sujoy Basu, Sharad Singhal

Hewlett-Packard Laboratories
1501 Page Mill Rd
Palo Alto, CA 94304, USA
{sven.graupner, sujoy.basu, sharad.singhal}@hp.com

Abstract

This paper describes a knowledge-based collaboration environment for ITIL [1] which is aimed at making work of IT consultants dealing with planning, designing and implementing ITIL processes easier. The collaboration environment organizes work of consultants around built-in ITIL concepts and allows them to map and refine concepts for a particular customer engagement.

The collaboration environment is designed as a so-called domain wiki that has, in contrast to other wikis, a pre-populated knowledge base built-in, which actively guides further information gathering and refinement as collaborative processes by consultants.

In this paper we present the design of a domain wiki for ITIL. Common ITIL knowledge is modeled from a subset of the ITIL publications and represented as an RDF knowledge base. Inferences on this knowledge base allow users to determine the aspects that need to be refined by consultants for particular engagements. Provided facts are then added to the knowledge base.

In ongoing research we investigate how such a knowledge base can be established for the domain of ITIL. One approach we take is a manual knowledge engineering effort; another investigation is how text analysis techniques can be used to automatically learn facts from ITIL publications.

1. Introduction

The IT Infrastructure Library (ITIL) [1] presents a comprehensive set of guidelines for defining, designing, implementing and maintaining management processes for IT services from an organizational (people) as well as from a technical (systems) perspective. The ITIL methodology is generally applicable in IT across industries and hence suited for establishing reusable practices.

Introducing and implementing ITIL-based processes for IT management is a people-intensive task involving domain expertise from ITIL consultants. It furthermore requires detailed knowledge of the industry in which a client is active and of the specific circumstances of a particular client's business and IT environment. IT service providers have established practices to help clients introduce ITIL-based processes in IT environments.

2. ITIL Practice and Technology

IT vendors have traditionally focused technology (tools and systems for management) around operational domains of ITIL such as service delivery and service support. A variety of tools and systems exist for service request management, incident management, problem management, configuration and change management, release management, and software asset management. Ideally, all supporting tools and systems are integrated and share a common configuration and management database.

An ITIL best practice is a proven set of documented activities and processes that have been successfully used across multiple organizations and can be reused with predictable outcome [2]. ITIL itself is a best practice, but at a rather abstract level which requires refinement and specialization for particular verticals (industries, geographies) and further down into particular clients' business and IT environments.

Consultants are trained to map abstract ITIL guidelines into the context of different clients. They refine, extend and complement common ITIL concepts into more specific concepts and the processes surrounding them. Reusing once established mappings is of key interest for IT service providers employing large numbers of consultants to make their work more effective.

However, while operational aspects of ITIL are relatively well supported by technology as mentioned before, pre-operational stages (particularly in service strategy and service design) are not. Modeling tools such as ProVision [3] and Aris [4] allow the definition of processes and even simulation, but the documentation they produce (files) is exchanged,

shared and used by traditional means (email, file shares) often losing the context in which they were produced and used. When projects end, and consultants move to other engagements, remaining documentation frequently becomes unusable due to the lack of contextual information.

3. Wikis as Collaboration Environments

Wikis have become popular as open collaboration environments on the web. They enable the creation of web content easily over the web itself and allow collaboration between large communities of volunteers. Wikipedia is the largest instance of a successful open wiki environment.

Wikis are structured around topic pages which can be linked from text in other topic pages. Each topic page is identified by a unique URL. Some contextual information can easily be captured by the wiki such as who created a page, when, who changed it, etc. Linking topic pages can capture further context such as the context within which a page is produced and used. Wikis, it appears, are an ideal collaboration environment. But, technically, since everyone can contribute anything to a wiki, they require some rigor in order to remain usable. Hence public wikis are frequently moderated by a person or a group. However, corporate project wikis are rarely moderated and hence quickly develop sprawl that is hard to navigate. Current and obsolete information cannot be distinguished, documents and pages are hard to find when not properly linked, etc.

It is thus desirable that *wikis themselves guide the process of information collection* around a topic domain such as ITIL. In order to do that, wikis need an information model which would allow them to:

- already know base-level concepts and relationships about a topic domain rather than starting from scratch;
- support reusability of base-level knowledge;
- support refinement of base-level concepts and relationships as a structured process derived from built-in domain knowledge;
- maintain methodology over the process of creating the information of the topic domain.

4. Wikis and Information Models

While common wikis make it easy to author web content at the front-end, they lack a formal model in the back end. Wiki implementations store content in files or compressed in databases. Logic to access content is coded into the wiki server, e.g. as PHP code. TWiki has been an early perl/cgi implementation [5] storing content as files in the back end. MediaWiki is a free wiki implementation which was originally written for Wikipedia, but is more broadly used today. It uses PHP and a SQL database in its backend [6].

Semantic Wikis aim to establish an explicit model in the backend which can be authored and changed through a wiki front-end, but which is also inferable by queries in the back-end. Semantic wikis draw on the Semantic Web, with the Resource Description Format (RDF) at its core [7]. Consequently, semantic wikis use an RDF store to represent information for topic pages. Each page contains facts about a topic, which is represented as a RDF resource in the RDF store with a unique URI. Content of topic pages can be annotated to represent resource properties in the RDF model. References to other topic pages can be turned into reference properties in RDF. RDF inferences can be executed on the backend RDF model and produce new content such as all topics that are related to a page.

Semantic Media Wiki [8] is a prominent semantic wiki implementation by the AI Lab at the University of Karlsruhe developed with the purpose to create a Semantic Wikipedia [9]. IkeWiki is a similar semantic wiki created by Salzburg Research [10]. OntoWiki [11] and KaukoluWiki [12] are further semantic wikis.

A limitation of semantic wikis is that resource properties need to be explicitly entered into the text of a topic page as assertions. Research is being conducted to learn properties from textual page content [13, 14].

Another approach is to utilize the domain model to explicitly guide the information gathering process.

5. Domain Wikis

A *Domain Wiki* is a wiki that has base-level knowledge about a topic domain (such as ITIL) already built-in. One question is how the base-level domain knowledge is defined. A second question is how it can be used to guide the information gathering and refinement process in order to become more complete. For example, a concept may be defined for a domain, but isn't detailed yet. The domain wiki in this case can explicitly ask users to refine this concept and prompt them to create additional topic pages.

It appears logical to build a domain wiki based on a semantic wiki since it requires an explicit model in the back end. In common with semantic wikis, domain wikis require:

- explicit RDF model in the back-end;
- inspection and manipulation via topic pages.

In addition, a domain wiki has:

- pre-populated base-level knowledge about the topic domain as a RDF model;
- concept of roles and responsibilities for knowledge and content creation by refinement; and
- observance of content creation as a controlled process based on the domain model.

6. Domain Wiki for ITIL as Collaboration Environment for ITIL Consultants

An ITIL domain wiki can theoretically be built based on any semantic wiki such as Semantic Media Wiki [8] or IkeWiki [10]. We have not chosen one of these wikis for implementation yet. For a domain wiki it is important to understand the relationships between the knowledge it starts with and the knowledge it wants to acquire. Logic in the semantic wiki server must be adapted to actively prompt the right person for input rather than relying on voluntary creation.

6.1. Knowledge Modeling and Representation

The ITIL publications represent a large body of literature with five main volumes about Service Strategy, Service Design, Service Transition, Service Operation and Continual Service Improvement [1].

As a starting point, we have chosen a subset of this documentation for knowledge modeling: "Supplier Management" from Vol. 2, "Service Design". This section is the starting point for the service outsourcing practice, which is of particular interest to us.

In order to illustrate the knowledge modeling approach, we consider a text fragment from the before mentioned publication:

"The goal of the Supplier Management process is to manage suppliers and the services they supply, to provide seamless quality of IT service to the business, ensuring value for money is obtained."

The important fact here is the concept of a Supplier Management process. Other facts are that it is the main artifact to manage relationships with suppliers and the services they provide. The remaining text is non-factual and can be ignored.

Aspects of the Supplier Management process are explained in more detail further down in the text:

The Supplier Management process should include:
- *Implementation and enforcement of the supplier policy,*
- *Maintenance of a Supplier and Contract Database (SCD),*
- *Supplier and contract categorization and risk assessment,*
- *Supplier and contract evaluation and selection,*
- *Development, negotiation and agreement of contracts,*
- *Contract review, renewal and termination,*
- *Management of suppliers and supplier performance,*
- *Agreement and implementation of service and supplier improvement plans,*

- *Maintenance of standard contracts, terms and conditions,*
- *Management of contractual dispute resolution,*
- *Management of sub-contracted suppliers.*

The important concepts here are the 11 aspects that constitute a Supplier Management process. (Note that the notion of a process in ITIL does not necessarily refer to an executable process, which means the steps above being executed in their order. It is rather meant to represent essential aspects.) All of these aspects are of an abstract nature as well and need to be refined by consultants by producing further content (topic pages and documents residing on those pages) reflecting how a concept is implemented in a particular context.

Different roles of consultants are involved contributing the different aspects of supplier policy, contract negotiation and agreements, reviews, etc.

The knowledge representation of the text fragment in RDF/OWL is:

```
:SupplierManagementProcess
  a                 owl:Class ;
  rdfs:subClassOf   :ITILActivity ;
  :aspectOf         :SupplierManagement ;
  :aspects (
        :ImplementationAndEnforcementOfSupplierPolicy
        :MaintenanceOfSupplierAndContractDatabase
        :SupplierAndContractCategorizationAndRiskAssessment
        :SupplierAndContractEvaluationAndSelection
        :DevelopmentNegotiationAndAgreementOfContracts
        :ContractReviewRenewalAndTermination
        :ManagementOfSuppliersAndSupplierPerformance
        :AgreementAndImplementationOfServiceAndSupplierImprovem
        :MaintenanceOfStandardContractsTermsAndConditions
        :ManagementOfContractualDisputeResolution
        :ManagementOfSubcontractedSuppliers
  ) .
```

The aspect of maintenance of a Supplier and Contract Database is then detailed as:

```
:MaintenanceOfSupplierAndContractDatabase
  a                 owl:Class ;
  rdfs:subClassOf   :ITILActivity ;
  :aspects (
        :SupplierAndContractDatabase
        :MaintenanceProcessOfSupplierAndContractDatabase
  ) .
```

The concept of a SupplierManagementProcess is defined as a class (rather set) with no members (yet), which means individuals (supplier management process instances) need to be produced as members of this class during the process of refinement.

The domain wiki can easily identify member-less concepts by simple inference and create stubs for topic pages for them. Users are then prompted to fill in detail creating the individuals for topic classes and enriching the knowledge base from provided input.

6.2. Knowledge Extraction by Text Analysis

Engineering knowledge using ITIL publications is an intense manual effort due to its volume. It is hence desirable to seek automated support for extracting knowledge from ITIL publications. The approach we chose for this effort is a text analysis approach [14].

The key concepts present in the ITIL documents are listed in the Glossary as single words, noun phrases and acronyms. This suggests that a noun phrase parser can do a reasonable job of marking these phrases in the text. Techniques from the information retrieval literature [15] can then be used to identify the important ones. We are currently in the process of evaluating alternative approaches for this step. Subsequently, the labels of the arcs in the RDF graph must be identified. The main intuition here is that the verbs in the sentences connecting the concepts will provide these labels. We use WordNet [16], a lexical database of the English language, to handle synonyms.

7. Prototype and Early Findings

A prototype is being developed to validate the approach and explore use cases. The prototype uses a light-weight java-based wiki as a front-end and the Jena toolkit for the RDF information base in its back end.
The use case of ITIL's supplier management has been modeled and implemented. The prototyping work has shown, already at the early stage, that there is a strong need for exploration capabilities of the models in the back end, which relates back to user interface design and user interaction models. A user is not satisfied with a system that is instructing him or her, although it might be the logical conclusion is information that is in the information base. Known issues of older AI-based systems surfaced quickly such as an ability to explain why the system suggested a certain action. In order to address the issues, further investigation of explanation techniques is needed and pursued in our research.

8. Related Work

Besides semantic wikis, which are domain agnostic, the work here relates to ITIL Ontologies, which are primarily considered in context of management system and SOA integration [17], primarily related to service delivery and service support systems including schema and data transformation aspects.

The interpretation of ontologies and RDF models as knowledge for inference and reasoning is mainly investigated in the AI community using a number of reasoning and inference engines [18]. A number of agent-based inference systems use RDF/OWL or DAML/OIL [19]. We combine AI and wiki aspects for a particular domain, ITIL.

9. Summary

This paper presented research about an ITIL domain wiki as a collaboration environment for consultants. The collaboration environment has a pre-populated ITIL knowledge base built-in, which actively guides the information gathering and refinement process. The result is a structured knowledge base in the back end, which can be explored and extended using a wiki front-end. A second part of research is directed towards learning knowledge (facts) from the ITIL documentation using text analysis. Research in the project is ongoing.

References

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

[2] ITIL Best Practice, http://www.best-management-practice.com.

[3] MetaStorm, ProVision, http://www.metastorm.com/products.

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

[5] TWiki, http://twiki.org.

[6] Media Wiki, http://www.mediawiki.org.

[7] Resource Description Framework (RDF), http://www.w3.org/RDF.

[8] Semantic Media Wiki, http://semantic-mediawiki.org/wiki/Semantic_MediaWiki.

[9] Krötzsch, M., Vrandecic, D., Völkel, M., Haller, H., Studer, R.: *Semantic Wikipedia*. In Journal of Web Semantics 5/2007, pp. 251–261. Elsevier 2007.

[10] Ike Wiki, http://ikewiki.salzburgresearch.at.

[11] OntoWiki, http://ontowiki.net.

[12] KaukoluWiki, http://kaukoluwiki.opendfki.de/.

[13] Cimiano, P., Völker, J., Studer, R.: *Ontologies on Demand? - A Description of the State-of-the-Art, Applications, Challenges and Trends for Ontology Learning from Text*, Information, Wissenschaft und Praxis 57 (6-7): 315-320. Oct 2006. http://www.aifb.uni-karlsruhe.de/WBS/pci/Publications/iwp06.pdf.

[14] Blohm, S., Cimiano, P., Stemle, E.: *Harvesting Relations from the Web -Quantifiying the Impact of Filtering Functions*, In Proceedings of the 22nd Conference on Artificial Intelligence (AAAI-07), pp. 1316--1323. Association for the Advancement of Artificial Intelligence (AAAI), July 2007.

[15] Furnas, G.W., Deerwester, S., Dumais, S.T., Landauer, T.K., Harshman, R.A., Streeter, L.A., Lochbaum, K.E.: *Information Retrieval using a Singular Value Decomposition Model of Latent Semantic Structure*, Proceedings of the 11th annual international ACM SIGIR conference on Research and development in information retrieval table of contents, Grenoble, France, Pages: 465 – 480 (1988).

[16] WordNet, http://wordnet.princeton.edu.

[17] Siebert, H.-G.: SOA and ITIL Integration using Ontologies, EuroCMG, 2007.

[18] Gomez-Perez, A., Fernandez-Lopez, M.: Ontological Engineering, 2nd printing, Springer, London, 2004.

[19] http://www.ksl.stanford.edu/people/dlm/webont/wineAgent.