

Social Computing and Governance in an Enterprise Service for Managing Business Processes

Melissa J. Bucu, Hani Jamjoom, Mark Podlaseck, Huiming Qu

IBM Thomas J. Watson Research Center
Hawthorne, NY, USA

e-mail: {mjbuco, jamjoom, podlaseck, hgu,} @us.ibm.com

Abstract— The enterprise has historically been ambivalent about Social Computing. Although dynamic communities, collaboration, innovation, and the wisdom of crowds were recognized as valuable, doubts about content quality and standards frequently overshadowed the benefits that Social Computing had already demonstrated outside the enterprise. In this paper, we describe how we have been challenged to add governance to an existing enterprise wiki designed for capturing, collaborating on, and evolving best practice business process assets. With the addition of governance, our goal has been to transform a best practice process wiki into an authoritative source for process assets for the enterprise, a repository for registering usage, tracking variations, and linking process assets to one another and to other assets (e.g. accounts) and a vehicle for continual improvement of process assets through annotation, rating, and collaboration. The extensive scope of the content in this project dictates that both the creation and the governance of the content be distributed to the appropriate level of the community. It also requires functions to assist users with leveraging potentially significant amount of user feedback. Overall, the focus of the project has been on maintaining the quality of the content while fostering its continuous improvement and relevance to the community.

Keywords-social computing, process governance, business process management)

I. INTRODUCTION

Enterprises are becoming increasingly process focused as they recognize that their processes are their business. Standardizing on best practice processes enables an enterprise to improve performance, consistency, quality, and efficiency - all of which are critical success factors in today's economic climate. Moreover, adherence to standards is not only good practice but in many cases is legally mandated. There is no shortage of industry standards and government regulations to which an enterprise may need or choose to conform. For example, there is the ITIL [1], ISO 9000 [2], eTOM [3], and Sarbanes Oxley [4] to name just a few.

Standardization of processes, whether based on one or many standards, is a top down activity. It is up to the leadership to select the standards and policies to which an organization will adhere. In order to align process assets with

the organizational strategy as well as manage risk, conformance, and quality; process governance is required. Without a governance structure for business processes, it is very difficult to achieve consistent business process improvement. [5]

One can represent the processes which define a business as a triangle (see Fig. 1). At the top of the triangle, one can think of a process called "Run the business". That top level process can be repeatedly broken down to describe the processes involved in "running" the business. This triangle of processes becomes very wide at the bottom indicative of the growing number of processes and users as the level of detail increases. Optionally, one can differentiate process levels into processes, activities, tasks, procedures, and work instruction. As an example for a Service Delivery organization, an activity might be "Fulfill service request" and a work instruction might be "Decommission Linux server for Acme account".

As Fig. 1 shows, customization or variations can occur at any level of process. There are generally industry standards at the higher levels of process, organizational standards at the middle levels, and local standards at the lower levels. Capturing enterprise processes and organizing them into such a structure becomes increasingly difficult as one moves to lower levels of detail. However, this is necessary in order to drive standardization to these lower levels where so much of the work is performed.

The capturing of processes used in a large enterprise is an enormous undertaking. The ongoing management and continual improvement of those processes is an equally onerous task. Such tasks are best handled by those who are closest to the performance of the work. In a large enterprise,



Figure 1. Enterprise Processes

these subject matter experts may have no occasion to interact with one another. This is precisely why an enterprise process repository is an ideal candidate for social networking. The ability to leverage the wisdom of the crowd is essential for the vitality of an enterprise process repository.

In this paper, we explore how we met the challenge of augmenting a social networking service with process governance to drive standardization and maintain quality while encouraging participation and innovation. In doing so, we have been able to achieve our goal of providing an authoritative enterprise process repository. The paper is organized as follows. Section II describes the Cyano Process Wiki. In Section III, we outline the governance model which has been developed. Section IV is a discussion of related work. The paper is summarized in Section V.

II. CYANO PROCESS WIKI

Cyano is a social networking service that was designed for IBM's IT service delivery organization to capture, link, rate, and socialize best practice processes. Process assets which have been modeled in a supported process modeling tool or described in a web form can be imported into Cyano.

Although a major purpose for our repository is to drive process standardization across the organization, we also recognize that there is a need for some level of customization to accommodate unique business needs and government regulations. Therefore, users are encouraged to annotate processes with comments, procedures, and process variations. Users may also annotate other annotations. Our goal is to track variations, eliminate unnecessary ones, and when appropriate incorporate useful ones into the standard processes.

The key design point in Cyano is that users are encouraged to define a profile that captures their 'operational context'. This includes what organization they belong to, what role they play, what account they support, etc. This context is used to restrict their view of the underlying knowledge (e.g. *Alice* will only see information relates to the accounts she supports for her particular). Annotations authored by users can be tagged by their context such that only users of overlapping context (e.g. supporting the same account) can view their annotations.

Cyano currently contains approximately 600 IT delivery process assets and has a user community of 13,000 who have added roughly 75,000 annotations. Our goal is for Cyano to become (1) the authoritative source for IT delivery process assets used across the enterprise, (2) a repository for linking process assets to one another and to other organizational assets (e.g. people, accounts, etc.) - necessary to perform impact analysis for change management, and (3) the vehicle for collecting user feedback and ratings. To achieve this goal, we are extending the current functionality to include process governance.

III. THE GOVERNANCE MODEL

Given the scope of the undertaking, we deemed it essential to be able to decentralize ownership and administration responsibility, distributing it across the organization to the appropriate level authority and expertise. Because Cyano is intended to be an authoritative repository for processes for service delivery, conformance with standards such as ITIL and ISO 9000, as well as internal quality and privacy standards was also essential. In order to decentralize the responsibility for ensuring that the content is conformant to the relevant standards, we added the ability to designate owner(s) for each object (e.g. process) in the repository. For processes, it is the owners' responsibility to understand relevant standards, ensure conformance to those standards, and create links between processes and relevant standards. An owner is encouraged to leverage the feedback provided from other members of the community but has ultimate responsibility for the content.

We also added ownership to other objects in the repository (e.g. organization objects). Owners of an organization (e.g. an account) can act on behalf of an organization. For example, only the owner of an account can link that account to a process indicating that the account uses that process.

Support for versioning, change management, access and document control features (e.g. revision history and frequency) were also added. The application was enhanced to allow the owner(s) to designate the reviewers and approvers as well as those with read, write, and administrative privileges for an object. A lifecycle state (see Fig 2.) was added for each object in the repository, as well as the ability for the designated owner(s) to manage the lifecycle of the object. How to assist process owners with leveraging what could be a considerable amount of feedback from the user community was yet another challenge we faced.

In developing a process governance model, our approach is to (1) organize the user generated content to effectively support future consolidation, (2) surface only the most important changes (or the thread of the changes) periodically - allowing time for ideas to mature prior to being considered for standardization, (3) distribute the change management responsibilities to the appropriate level within the organization in order to improve the quality of the process assets and expedite the review process.

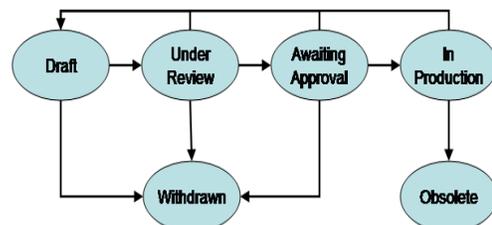


Figure 2. Lifecycle for a process asset

A. Process Information Management

As a Cyano user, an IT professional can review the most up-to-date processes, procedures, work instructions, and tools. More importantly, s/he can also contribute her/his empirical experience that relates to her/his job roles as well as rate existing content. Fig. 3(a) shows a logical view of a process with various user messages. In this example, the process consists of four tasks, one of which has been tagged with five branches of messages. These messages might be suggestions on how to improve the task. One message could be followed by series of discussions and may also be branched out to other topics as well, resulting as a deep sub-tree rooted at the task. The messages may also be scope definitions, such as restricting this message to a set of accounts only. Such sub-trees are usually wide and shallow since there are few overlaps among supported accounts.

For the data model of messages, we need to capture not only the structure of the tree with different type of instances, but also important attributes for rating the messages. Fig. 3(b) shows a view of the data model for messages. The *parent type* of a message is needed since all instances (processes, tasks, user messages, etc.) can be tagged by user messages. The *status* of a message could be *Deleted*, *Open*, *Under Review*, *Approved*, or *Rejected*.

B. User Content Review

Due to the massive amount of user-generated content with varying qualities, it is prohibitive for the review team to constantly monitor all the changes. Some level of automation guided by content popularity is needed to select only the most important changes for periodic review. As a first step, the system allows users to rate any individual message on how influential or helpful it is (e.g. indicated by the number in each message box in Fig. 3(a)). The ratings are further weighted according to the user level (e.g. high profile users' ratings are valued more). We also have found that proper aging of user feedback is critical. Exponential aging, with a slow decaying constant seems to work well.

Another challenge lies in how to choose the granularity of the messages. Since all messages are inter-connected or connected with part of the processes, we cannot pick only the top-rated messages to review. Currently, the messages are grouped into sub-trees that are rooted at tasks or other instances that have already been reviewed and approved to be included in the process. The rating of the sub-tree is the accumulative sum of the ratings of included messages. Only the sub-trees with top accumulative ratings will bubble up to the review team. In this way, not only the review overhead is significantly reduced, but also it allows more time for the new topic of changes to mature and grow from the collective knowledge in the community.

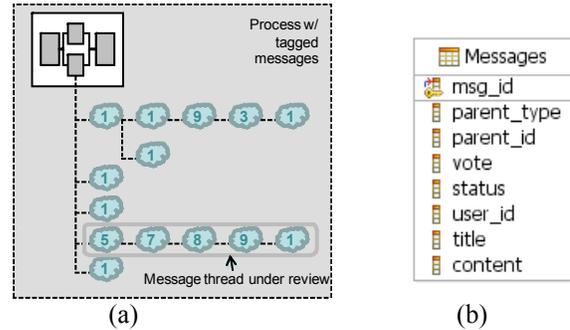


Figure 3. Process with user-tagged messages

IV. RELATED WORK

Governance of socially created knowledge is not new as it overlaps with reputation-based systems, collaborative filtering, and recommender systems [6, 11, 12]. There, several models have seen wide adoption. To the best of our knowledge, all successful models have not been in enterprise or mission-critical environments. The most popular is probably Amazon's product review, where users are encouraged to review any product listed by Amazon [9]. They are also encouraged to review resellers. What contributed to Amazon's success are four simple ideas. First, users can rate other users' reviews. Second, helpful reviews bubble to the top. Third, both the most helpful positive and most helpful negative review are highlighted. Fourth, authors of consistently helpful reviews are identified as "top reviewer". The system is remarkably effective: some analysts attribute its reviews as the primary reason that Amazon.com is the origin of more than 30% of all e-commerce traffic.

Another well-known model is that of Wikipedia [7], with its increasingly decentralized and complex system of negotiated policies and guidelines, all of which have evolved from historical circumstances and conflicts, rather than originating in a grand plan by Wikipedia's founders. One study compared Wikipedia to Encyclopedia Britannica and concluded that Wikipedia's accuracy is comparable to that of Encyclopedia Britannica [8]. When introduced content on Wikipedia is inaccurate or controversial, one of the many regulatory mechanisms is engaged and the issue is resolved, often within minutes. Viewed from a governance perspective, these studies are pointing to slow convergence speeds for content in terms of accuracy.

There are other wiki implementations that attempted to challenge Wikipedia. Most recently is Knol by Google [10]. Knol added three design points to muster interest in its service. First, users' identities are public and can be verified. The intuition is that a person's background (e.g. a medical doctor) can add a level of authority to the content s/he generates. Second, document owners control editorial access by other users. This allows the formation of well-defined

sub-communities with clear leaders. Third, a user can embed AdSense in their documents. This allows revenue generation and sharing for popular documents. At first glance, these three ideas should create a platform for generating high-accuracy content. Interestingly, this is not the case. While there is no definitive answer, some cite that Knol's lack of success is attributed to authors' greed. Basically, authors are discouraged from allowing others to edit their documents for fear of losing ad revenues, which leads to stale information.

V. SUMMARY

In summary, the addition of governance to our existing social networking service for managing best practices has enabled the service to provide considerably more value to the enterprise. Unlike self-organizing communities (e.g. Wikipedia), an enterprise generally has an inherent hierarchy for resolving conflicts, policies can be mandated, and decisions need not be made through consensus. Our challenge has been to strike a balance between controlling the quality of the content and encouraging innovation and collaboration. The proper balance will result in a service which is recognized as both vital and valuable to the enterprise and community.

REFERENCES

- [1] <http://www.itil-officialsite.com/home/home.asp>
- [2] <http://www.iso.org/iso/home.htm>
- [3] TM Forum, <http://www.tmforum.org/browse.aspx>
- [4] Sarbanes Oxley, <http://www.sarbanes-oxley.com/>
- [5] J. Jetson, and J. Nelis, "Management by Process: A Practical Road-map to Sustainable Business Process Management," *Elsevier Science & Technology Books*, May 2008.
- [6] G. Linden, B. Smith, and J. York, "Amazon.Com Recommendations: Item-To-Item Collaborative Filtering," *In proceedings of IEEE Internet Computing*, Jan 2003.
- [7] A. Forte and A. Bruckman, "Scaling Consensus: Increasing Decentralization in Wikipedia Governance," *Hawaii International Conference on System Sciences, Proceedings of the 41st Annual*, Jan 2008.
- [8] J. Giles, "Special Report: Internet Encyclopedias Go Head To Head," *Nature*, no. 438, pp. 900-901, Dec 2005.
- [9] R. Spector, "Amazon.com: Get Big Fast," *Collins Business*, 2002.
- [10] Knol, knol.google.com
- [11] G. Adomavicius and A. Tuzhilin, "Toward the next generation of recommender systems: a survey of the state-of-the-art and possible extensions," *In IEEE Transactions on Knowledge and Data Engineering*, Vol. 17, Issue: 6, pp. 734-749, June 2005.
- [12] C. Basu, H. Hirsh, and W. Cohen, "Recommendation as Classification: Using Social and Content-Based Information in Recommendation," *In Proceedings of the Fifteenth National Conference on Artificial Intelligence*, 1998.