PeopleCloud for the Globally Integrated Enterprise

Maja Vukovic[1], Mariana Lopez[2], and Jim Laredo[1]

[1] IBM T.J.Watson Research, Hawthorne, NY 10532, USA
[2] Carnegie Mellon University, 5000 Forbes Ave, Pgh, PA 15213
{maja,laredoj}@us.ibm.com, nanalq@gmail.com

Abstract. Crowdsourcing has emerged as the new on-line distributed production model in which people collaborate and may be awarded to complete a task. While many existing services enable enterprises to employ the wisdom of crowd, there is no existing practice defined for integration of crowdsourcing with the business processes. We propose PeopleCloud, as the (1) mechanism to enable access to scalable workforce on-line, connecting it to the enterprise and (2) an interface to services required for crowdsourcing tasks. We define requirements for PeopleCloud, based on our experiences in employing wisdom of crowd to source business and IT information within the enterprise.

Keywords: Collaborative intelligence, Crowdsourcing, Globalization.

1 Introduction

As the physical and digital worlds are becoming universally connected, and computational resources and data are available beyond their immediate owner, it is now possible to effortlessly reach out to the masses, and externalize the "*function once performed by employees and outsourcing it to an undefined … network of people in the form of an open call",* the process which Howe [1] defines as crowdsourcing.

Leveraging the wisdom of crowds in the form of open calls and through marketplaces is not an entirely new paradigm. In the past companies have run competitions to engage their customers and end-users to contribute towards certain enterprise functions, such as advertising campaign design and specific problem resolution challenges [2].

Numerous purpose-built crowdsourcing solutions are becoming available, allowing enterprises to reap the benefits of scalable, global workforce, while lowering the cost of execution. For example, uTest is the marketplace for software testing services, oferring real-world QA services through their community of more than 14,000 professional testers from 151 countries around the globe. Mob4Hire Inc reaches out to 86 countries, 2000 handsets and more than 130 mobile operators, helping mobile application developers access variety of testing platforms and testers in real field conditions. TopCoder.com is a community of over 140,000 skilled software engineers.

Despite some of the success stories, such as The Goldcorp Challenge [3] and Threadless.com, the actual realization of the promising advantages for enterprises from crowdsourcing are far from being well-achieved and pose an extensive range of interesting challenges along social, legal, economical and technical dimensions [4]. One of the key barriers for enterprises to effectively employ crowdsourcing is how to

A. Dan, F. Gittler, and F. Toumani (Eds.): ICSOC/ServiceWave 2009, LNCS 6275, pp. 109–114, 2010.

integrate this process with the existing enterprise operations, as well as provide the incentives encouraging (honest) contributions. Tapscott and Williams [5] discuss how businesses can harness collective capability to facilitate innovation, growth, and success. In contrast, our research investigates applicability of crowdsourcing methodology within the enterprise.

Xerox's Eureka system [6] is one of the early examples of a crowdsourcing within the enterprise, where wisdom of crowd was employed to enrich enterprise support knowledge base. Enterprises address their business goals by formulating them as tasks, typically addressed by the subject matter experts. This approach is often inefficient because related know-how might be scattered among different teams and as a result hidden behind the teams' organizational structure. This problem becomes more significant as the enterprise's size and geographic coverage increase.

In this paper, we propose PeopleCloud, a model for encapsulating work tasks and people activities as a service. It provides a framework for business processes that span geographical and organizational borders. By defining a set of roles and their activities in the PeopleCloud, we take the first step in exploring the governance model for crowdsourcing-enabled processes. Finally, we describe the findings from deploying the PeopleCloud to source IT and business information within an enterprise.

2 Overview of PeopleCloud

Building on the requirements for a general-purpose crowdsourcing service, presented in [8], and our experience with employing wisdom of crowd within the enterprise, we define the PeopleCloud as an on-demand service system that spawns and manages scalable virtual teams of knowledge workers by building on the wisdom of crowds within an enterprise or across a value chain. PeopleCloud are design to increase effectiveness in building the transactive knowledge networks within an enterprise to execute complex and transformative knowledge-intensive tasks.

We envision enterprises being enabled to create highly customizable requests to access the PeopleCloud, while defining the business process standards and crowd requirements, as well as services required for their execution. PeopleCloud concept necessarily builds upon state of the art in CSCW, CHI, organizational theory and service design. In addition to the crowdsourcing features discussed in [8] PeopleCloud provides mechanisms for:

a) Effective and sustainable expert discovery
b) Supporting teamwork
c) Task creation, precedence rules and integration to the enterprise systems
d) Provisioning services required to execute tasks

Figure 1 depicts the main roles and their corresponding activities in the PeopleCloud.

Requestors initiate PeopleCloud process by defining the request, specifying task goals, task sequence and the crowd requirements. The requestor searches for crowd members, manages notifications, accepts bids and reviews submissions. A requestor may also cancel work demand or extend deadlines.

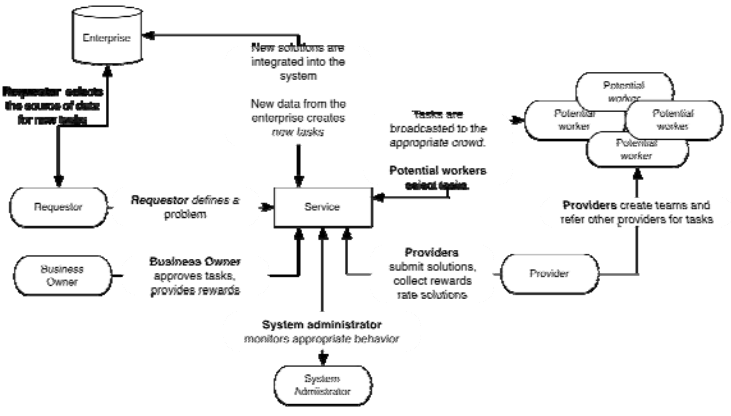

Fig. 1. Key Roles and activities in PeopleCloud

Business owners work jointly with the requestor, and are responsible for defining integration points, validating the successful completion and compensating the appropriate parties. They handle the selection of appropriate compensation models depending on the task complexity and the level of crowd expertise required.

The PeopleCloud is divided into two categories: the potential worker and the service providers. Potential providers may search amongst tasks or be suggested tasks. Once potential worker accepts a task, the requestor or the system must approve this. The **provider** executes the tasks they have selected. Each task has a specific space for the worker to provide their results and comments. Providers may return tasks that cannot be executed, or suggest another provider. On completion the provider submits the task in exchange for rewards, and may also rate other solutions.

The **collaborator** is responsible for forming teams, dividing tasks between team members, commenting on open tasks, suggesting providers and rating submissions. A collaborator could be either a provider or a requestor performing their roles in teams; and is crucial to expanding the knowledge network and to aiding the process of finding the correct providers for the task [9]. Thus, the PeopleCloud becomes a self-sustainable and self-discovery mechanism. In the enterprise context often we find existing repositories of task owners, based on their business responsibilities. This can in addition be used to jump-start the crowdsourcing process.

The **facilitator** is responsible for providing mediation if conflicts arise in crowdsourcing process, whether between requestors and providers as well as issues between team members. The facilitator monitors the behavior of the crowd, ensuring that requests are valid, ratings and rewards are fair and legitimate.

3 PeopleCloud for IT Optimization

In this section, we present the results from deploying a PeopleCloud to manage and control enterprise knowledge, specifically to identify the business use of the underlying physical IT infrastructure. Complexity of crowdsourced tasks, specifically in the

knowledge transfer domain [8] may vary, from the microtasks – requesting a single piece of information or activity to be performed, to crowdsourcing large-scale projects. In this work, we focus on discovery of "extreme information", that is, information that only a few possess, out of a pool of thousands of candidate information holders. IT inventory management is a domain where such information is hidden in the knowledge of individual team members, yet there is little transparency on who knows what.

Background. At present, information about the enterprise applications is typically stored in a centralized repository. This data is often outdated and incomplete, with limited bindings to the server infrastructure. At the same time, there is a lack of global knowledge within the enterprise (e.g. the different stakeholders, such as production control and application support experts possess partial information that can be leveraged to enrich and co-create the enterprise-wide view of the business capabilities of the infrastructure). Furthermore, interaction among employees exposes the inherently embedded social networks, which can be further exposed and employed to control the enterprise knowledge.

The goal of this crowdsourcing exercise was to discover, integrate and manage the knowledge about the business applications' infrastructure. Our PeopleCloud achieves this by harvesting Business-IT knowledge through wisdom of crowd and rewarding participants for the knowledge contributions (server to application mappings) and knowledge seeking behavior (e.g. finding other knowledge experts), which captures the global enterprise knowledge. Our PeopleCloud embeds capabilities to build, discover and maintain linkages among user communities to enable collaborative intelligence.

PeopleCloud deployment. We have deployed People Cloud to gather Business and IT information of more than 4,500 business applications within and enterprise, owned by more than 2,300 unique business application owners within the global enterprise. The initial seeding e-mail was sent to existing application owners inviting them to contribute their knowledge. Once they accessed the service, contributors had the option to refer their tasks to other experts in the enterprise. There was no access control, thereby allowing anyone from the group of application owners to complete the crowdsourced tasks. The system, however, had the capability to provide full audit trails, thereby enabling transparency of the community contributions. Following the seeding e-mail, a reminder for pending tasks was sent to the currently known application owner each week for the first 6 weeks of the run. After the sixth week, once we reached 90% task completion, the reminder frequency was increased to twice weekly.

Effectiveness. Within the first four days of deployment we have collected 50% of the targeted mappings between applications and servers. Table 1 shows the distributions of responses over the time.

The overall experiment increased the efficiency of the process by 30x, in comparison to the traditional approach (that involved two full time experts). They shared a collaborative virtual team-room, in which they documented the servers, applications and their owners, as they discovered them. They typically contacted the application owners through instant messaging, email or phone, and have then updated the master spreadsheet. Using this approach, two full time experts would take 2-3 months to collect the information about 140 applications and roughly 700 servers.

Table 1. Distribution of crowd responses over the time

% Completed	Duration in days
96	70
87	56
50	4
10	1

During this experiment, we have discovered that 5% of the original population of business application owners has changed. 220 referrals resulted and were completed from the access to the crowdsourcing service, providing an insight about the social network and dynamics within the enterprise that govern the flow of the knowledge that we seek. The feedback from the application owners has demonstrated the value of the knowledge that was collected. They were interested in the inferred information and the various communities that were formed, such as:

1. `I would like to see the list of servers for my application.'
2. `Which other applications are hosted on the same servers?'
3. `On which servers is my application running?'
4. `Who are the Application Owners of these applications?'
5. 'Which applications will get affected if I consolidate this data center?'

Incentives. Incentives play the key role in crowdsourcing efforts [10], both within and outside of the enterprise. Numerous incentive mechanisms exist, which may introduce inconsistency in crowdsourcing process and make collaborative production challenging [11]. Traditional award schemas are presently employed by enterprises, e.g. salary, performance bonuses. Incentives for crowdsourcing raise new legal challenges (e.g. compliance and taxation). Furthermore, existing HR policies may further impose limitations on participation of permanent and contract employees. The deployed PeopleCloud did include the capabilities allowing participants to collect virtual points for crowdsourcing task, as well as for any successful referrals. The points, however, at this stage were not exchangeable for a tangible, material award. The users had the ability to see their rating, compared to other participants. Finally, many users found the access to the consolidated knowledge that was a result of the crowdsourcing task, and incentive on its own.

4 Summary

In this paper we outlined the challenges for enterprises to reap the benefits promised by employing wisdom of crowd approach in their business processes. We proposed PeopleCloud, as the (1) mechanism to enable access to scalable workforce on-line, connecting it to the enterprise and (2) an interface to services required for crowdsourcing tasks. Finally, we described our experience in deploying PeopleCloud within the enterprise to execute IT inventory management exercise.

Acknowledgments. We would like to thank Kumar Bhaskaran for insightful discussions on PeopleCloud.

References

1. Howe, J.: The Rise of Crowdsourcing. Wired 14(6), http://www.wired.com/wired/archive/14.06/crowds (accessed October 23, 2008)
2. Surowiecki, J.: The Wisdom of Crowds. Anchor (2005)
3. GoldCorp Challenge, http://www.goldcorpchallenge.com/ (accessed November 14, 2008)
4. Brabham, D.C.: Crowdsourcing As A Model For Problem Solving: An Introduction And Cases. Convergence. The International Journal of Research into New Media Technologies 14(1), 75–90 (2008)
5. Don, T., Williams Anthony, D.: Wikinomics: How Mass Collaboration Changes Everything. Portfolio Hardcover (December 2006)
6. Galegher, J., Kraut, R.E., Egido, C. (eds.): Intellectual Teamwork: Social and Technological Foundations of Cooperative Work. Lawrence Erlbaum Assoc., Inc., Mahwah (1990)
7. Vukovic, M.: Crowdsourcing for Enterprises. In: International Workshop on Cloud Services at 7th IEEE International Conference on Web Services (2009)
8. Argote, L.: Organizational Learning: Creating, Retaining, and Transferring Knowledge, 1st edn. Kluwer Academic Publishers, Dordrecht (1999)
9. Kittur, A., Chi, E.H., Suh, B.: Crowdsourcing user studies with Mechanical Turk. In: Proceeding of the Twenty-Sixth Annual SIGCHI Conference on Human Factors in Computing Systems, CHI 2008, Florence, Italy, April 05 - 10 (2008)
10. DiPalantino, D., Vojnovic, M.: Crowdsourcing and all-pay auctions. In: Proceedings of the Tenth ACM Conference on Electronic Commerce, EC 2009, Stanford, California, USA, July 06 - 10 (2009)
11. Bartlett Christopher, A., Ghoshal, S.: Transnational management: text, cases, and readings in cross-border management, 3rd edn. Irwin/McGraw Hill, Boston (2000)