To appear in: Proc. of X Brazilian Symposium on Human Factors in Computing IHC'11. Porto de Galinhas, Brazil. October 25-28 2011.

Knowledge and Information and Needs of System Administrators in IT Service Factories

Cleidson R. B. de Souza, Claudio S. Pinhanez, Victor Cavalcante
IBM Research - Brazil
Rua Tutóia, 1157 – São Paulo – SP - Brazil
cleidson.desouza@acm.org, claudio@pinhanez.com, victorfc@br.ibm.com

ABSTRACT

In this paper we describe the results of an empirical study about the information needs of system administrators in large scale Information Technology (IT) service delivery organizations. This study is based on interviews, non-participant observation, and an electronic survey with more than 200 systems administrators working on incident management, covering their daily work including knowledge and information needs in complex situations and routine work. Although previous work has uncovered the knowledge-intensive and collaborative nature of system administrators' work, the results described in the paper detail a much more complex picture: (i) low usage of knowledge management tools; (ii) high usage of personal communication and collaboration tools; and (iii) need of gathering information about customers from a specific set of stakeholders. We also found that the most useful sources of information in handling complex situations are: the customer; the customer account team; and other employees who were experts both in the customer and in particular aspects of the delivery of services. The results of this study indicate that knowledge management in IT service factories is very challenging and possibly should evolve from the current passive model of knowledge management to a dynamic style emphasizing both reusability through information technologies and sharing through face to face and computer-supported discussions among employees.

Categories and Subject Descriptors

H.5.2 [**Information Interfaces and Presentation**]: User Interfaces – interaction styles; Group and Organization Interfaces – collaborative computing; K.6.4 [**Management of Computing and Information Systems**]: Systems Management

General Terms

Management, Documentation, Design, Experimentation, Human Factors

Keywords

Information seeking, information needs, system administrators, knowledge management, information technology management, IT services, service factories.

1. INTRODUCTION

While the in-house, aquarium-like datacenter was the dominant provider structure of Information Technology (IT) from the inceptions of the computer age through the end of the 1980s, the 1990s saw the emergence of outsourcing of IT services as an efficient and often more cost-effective way for organizations to get their IT needs satisfied. This new business model where the care and control of most elements of the IT infrastructure is outsourced to specialized IT providers was made possible by the appearance of an IT service delivery model based on what we call *IT service factories*. Those organizations, often employing thousands of specialized IT workers, have sprung all over the world, but especially in India and other emerging countries like Brazil. The use of the term *factory* reflects the fact that those organizations borrow a great deal of structural resemblance to traditional manufacturing factories.

In IT service factories, some of the most critical functions necessary to maintain the customers' IT infrastructures running well and efficiently are performed by highly specialized and trained professionals often referred to as *system administrators*, a.k.a. *sysadmins* or *SAs*. They are the ones responsible for determining the cause of problems affecting IT systems and fixing them; for preventively taking actions to keep the systems running without downtimes; for installing, updating, and upgrading the hardware and software components installed in the machines; for protecting the system against attacks and other threats; to manage thousands of user accounts; and to manage, protect, and backup all the data in the systems. Although the term system administrator is also used to designate people providing IT support in all kinds and sizes of companies, in IT service factories sysadmins often work in large workgroups with specialized goals and competencies, sometimes supporting systems from multiple customers. These are very knowledge-intensive jobs [1], where the SAs have to find, understand, interpret, and apply not only technical knowledge, but also knowledge which is specific about each customer and its IT infrastructure, and interact with users and IT workers from the customer.

Only recently researchers have started to pay attention to the work of system administrators in the context of service factories [2, 19, 24]. Examples of previous work include the study by Magglio and colleagues in which they illustrate the collaborative [3] and knowledge-intensive [4] nature of sysadmins' work. Another example is the work of Gonzalez and colleagues [1] which clearly suggests that IT workers' activities are varied, fragmented, and overlapped, what forces people to limit the focus on each activity for a short period of time.

While previous work has illustrated the knowledge-intensive nature of IT management, to the best of our knowledge, previous research has overlooked the information and knowledge needs of sysadmins, including which type of information they seek, how often this happens, and how much time they spend doing so. This paper reports on an empirical study, a survey, conducted in a large-scale organization which provides IT services for a variety of other large organizations through outsourcing contracts. Such customer organizations have hundreds of IT components (hardware and software) and hundreds, often thousands of users. The IT services provider, on the other hand, delivers its services also through hundreds of system administrators whose main goal is to support the information infrastructure of the customer organizations.

This paper is based on a large scale survey where professional sysadmins from an IT service factory were asked questions about their needs when working with the maintenance of large complex IT systems. Questions in the survey were designed to gather information about the type of information needed by the sysadmins, as well as the resources in which this information was available (such as in documentation, from co-workers from the same or different departments, etc.), the effort to gather such information, and the frequency in which sysadmins needed to seek information.

The survey provided two major insights about the nature of the knowledge needed by SAs in IT service factories. First, knowledge about the customer organization is absolutely critical for the successful delivery of IT services. We found out that in complex situations which are relevant for the customer, the most important sources of information for employees are: (i) the customer; (ii) the group of employees responsible for representing this customer inside the organization; and (iii) other employees who were experts both in the customer and in particular aspects of the delivery of services. We saw that the most relevant information were directly associated with the customer. The knowledge necessary for handling complex situations includes information not only about the customers' IT infrastructure and critical business processes and their mapping into the IT systems, but also information about people to be contacted, decision-makers, and business models. If, for instance, a customer's IT component is problematic, it is important that the service organization understands how this component affects the business of the customer organization and who to work with to properly correct the problem, while maintaining the customer informed about the progress of the work.

The second idea uncovered by our survey is related to the collaborative nature of the work of IT sysadmins. Our results corroborate previous work [2, 3, 20, 24] indicating that the work of sysadmins is highly collaborative. However, it provides initial evidence that there is a potential imbalance in the work of sysadmins: while some professionals are overwhelmed because they are important sources of information for their colleagues, others are not invoked by their colleagues during their work. Our data also suggests that expertise is not an explanation for this difference between IT professionals. Understanding this aspect is important because it might help us to design better tools for the "different" types of IT professionals.

The paper is organized as follows. Given that the IT service factory model of delivering IT services to customers is still not known by many researchers, we start the paper by describing IT service factories and how work is performed there (sections 2 and 3). We then present the survey and its methodology (section 4), the data collected (section 5), and an analysis of its results (section 6). Section 7 presents our conclusions and ideas for future work.

2. IT SERVICE FACTORIES

As mentioned in the previous section, this study was conducted in a large-scale organization, here called the *BSF* for the *Big Service Factory*, an organization which provides IT services for a variety of other large organizations through outsourcing contracts. In this paper we call the organizations which outsourced their IT infrastructure to the BSF simply as the *customers* of the BSF.

A fundamental aspect of IT service factories is their large-scale. In our case this is reflected in two ways. First, the number of employees responsible for dealing with the delivery of services: at the BSF: there are several thousand employees working solely with this aspect. And second, the number of customers, i.e., a service factory provides services to dozens of customers, which often are not collocated with the service factory. In particular, the service factory we studied, BSF, delivers services to customers in the U.S., Europe, Latin America, and Brazil.

It should be noted that there is a huge difference in scale between customers: while some of them have thousands of IT components (servers, network equipments, etc), others only have dozens of them. In general, smaller numbers of IT components also mean smaller number of problems, although there are sometimes very problematic small-scale systems. The large number of IT components being managed by a service factory is another interesting aspect which influences how service factories operate. Furthermore, often IT service factories work 24 hours per day, 7 days per week, although this level of support is often not mandatory for all of customers' systems.

A key sub-organization of an IT service factory is the one which handles *incidents.* An *incident* can be defined as: *"Any event which is not part of the standard operation of a service and which causes, or may cause, an interruption to or a reduction in, the quality of that service*." [5]

The incident management department performs some of the key tasks needed to maintain the customers' IT infrastructure. This includes restoring normal operations as quickly as possible with the least possible impact on the customers' business. Incidents are one of the smallest units of work in IT service factories: each employee working in the delivery of services usually works on incidents from any of the different customers which have contracts with the service factory. In most IT service organizations, the most important information related to one specific incident is collected in what is called a *ticket*, which corresponds to a record of one single incident in the incident management database.

3. THE BIG SERVICE FACTORY (BSF)

Service factories, as any other organizations, can be organized in multiple ways [6]. In particular, in this section we describe the organizational structure of the BSF we studied.

From an organizational point of view, system administrators at the BSF are assigned to work in departments in such a way that each department is based on common skills, competencies, and activities performed. Examples of departments include: UNIX, responsible for dealing with aspects of UNIX-based operating systems; security, responsible for accountability, security updates and similar issues in different types of operating systems; etc. At the BSF there are about 40 different departments responsible for the delivery of IT services for the different customers. Some of the departments are much bigger than others, reflecting the high number of incidents they have to address. Note that the structure of the BSF's organization in competence areas means that in a given day, a sysadmin might work with incidents from completely different customers.

This BSF basically implements the principles proposed by the *ITIL* [5] framework for IT organizations, the most important standard for the industry. The departments mentioned before (UNIX, security, etc) are departments responsible for the *Incident Management* of services delivered, i.e., organized according to competence areas. In addition, there are departments responsible for managing the delivery of other services, following the structure proposed by *ITIL*: *Change Management*; *Process Managemen*t; among others. As an example, the *Process Management* department is responsible for establishing, documenting, enforcing, and maintaining the different organizational processes used to deliver additional services for the customers which are not related to incidents.

Within each department of the BSF, employees are classified in three levels according to their expertise: Level 1 (L1), Level 2 (L2) and Level 3 (L3). Level 1 employees have limited technical knowledge on the competence area of their department and handle technically simple incidents which can (or should) be handled in a short time-frame. Level 2 and 3 employees handle progressively more complex incidents with associated longer resolution times. Finally, some of the Level 3 employees are also *CEs*, or *customer experts*, meaning that they are responsible for understanding in more details the IT environment of particular customers. Given the large number of customers, some BSF employees are CEs for up to 5 different customers.

Finally, within the BSF it is also possible to identify an organizational entity responsible for acquiring new customers, coordinate the outsourcing process, and selling additional services. Within this organizational unit, there is the so-called *customer account team*, the group of employees whose job is to manage the customer account and the overall relationship with the customer. In other words, they are the employees who interact with both the BSF employees and the customer, negotiating prices, contracts and conditions, and, more importantly, discussing the quality of the services delivered by the BSF. There is a clear separation of responsibilities in the customer account team. One of the roles, referred here as the *Account Technical Manager*, or *ATM*, is responsible for interacting with the BSF's delivery personnel. Another role is the *Account Leader*, or *AL*, responsible for interactions with the customers. A third role, referred here as the *Account Relationship Manager*, or *ARM*, acts as a bridge between those two roles. This does not mean, however, that employees in the ATM role are not required to interact with the customers. In general, the customer account team can be understood as representing a particular customer while interacting with the BSF teams; and representing the BSF as a whole while interacting with customers.

Finally, it is important to mention that most of the BSF employees are located in two different sites 100 kilometers apart from each other. Furthermore, some members of the customer account teams sometimes work collocated with their customers, so that members of the service delivery, management and customer account teams are all spread out in different places. This means that face-to-face communication is often very difficult, especially in emergency situations. It is also worthy mentioning that the BSF studied has first-class service quality levels and work practices which position the BSF among the top IT providers in the world.

4. METHODOLOGY

To study the knowledge needs of the SAs in the BSF, we are conducting a series of studies. This paper reports the first large scale study we performed in this organization, where we used a sequential exploratory methodology characterized by collection and analysis of qualitative data followed by collection and analysis of quantitative data. The main study activities of this first large study are described in the next subsections.

4.1 Qualitative Data Collection and Analysis

The qualitative approach used in this study consists of semi-structured interviews and non-participant observations conducted over a period of about four months. In this period, employees from different departments of the IT service organization and from different organizational positions were interviewed [7] and observed [8]. We selected informants based on our personal contacts and on recommendations from key informants.

We conducted semi-structure interviews [7] using an interview guide. This guide included open questions that allowed the interviewees to comment on the most important problems they identified in their departments and in the delivery of IT services. Interviews lasted about one hour and were not recorded, therefore we wrote down notes for later analysis. Interviews were conducted during a period of about two to three months and were intertwined with analysis of the data. The analysis of the qualitative data we collected was loosely inspired on Grounded Theory methods [23], aimed to develop emerging themes for research. One of the themes that emerged was the need to better record, understand, and manage the customer context in the deliver of IT services.

Based on this insight, we used non-participant observation [8] of employees who worked solely with incidents to further investigate this aspect. During the observations, we wrote field notes which were later integrated with the notes from the interviews. The main idea of non-participant observation is to observe IT professionals in their workplace performing their usual daily activities. More specifically, observation consisted of writing notes about the informants' activities, events, interactions, tool usage, and any other phenomena. To collect this information, the first author sat together with the informants, distant enough to not disturb them but close enough to be able to observe the contents of physical or digital objects which the informants were handling. Information collected during the observations was recorded without distracting the informants. In some occasions, the first author asked clarification questions during the observations, but not without ensuring that such questions would not disturb the informants' activities. If it was not possible to ask clarification questions during the observations, those questions were asked at the end of the day or during work breaks.

Non-participant observation was an important part of this research because it allowed us to understand the overall context of work during the daily delivery of services. For instance, it helped us to understand which tools the informants used and, more importantly, how they used them, including the limitations of the tools. Moreover, the results from the observation were a source of valuable insights to be explored during the survey performed later. As an example, it allowed us to identify which options should be made available for each question in the survey. In addition, the survey allowed us to test some of our hypothesis derived from the qualitative analysis. For instance, our observations lead us to believe that the Knowledge Management style adopted in the BSF organization was passive [10] as presented in the discussion section, so we framed some of the questions to identify the individuals who were the more common and helpful sources of information (as later shown in Figures 1 and 2).

4.2 Survey Design

After the interviews and observational work, it became clear that we needed to collect quantitative data about the information and knowledge needs of IT delivery personnel. More specifically, we were interested in professionals who worked in the direct delivery of services, the system administrators. Therefore, given the organizational structure of the BSF (described in section 3), we restricted the survey to BSF employees who worked in technical departments and, more specifically, in incident management. Managers and other support personnel were not included in the sample as well as employees with managerial functions or members of the customer account team, namely the ATM, the AL and the ARM.

Respondents were invited to answer a 45-question electronic survey designed to cover two main aspects of the employees' work: their overall daily work and their work handling complex tickets which demanded significant additional effort from them. Here, we use the term *effort* as the amount of mental and physical energy spent while working in the ticket. We also collected demographics (gender, age, education, employment status, etc.) about the informants as well as the context in which the respondents were embedded: work shift, job role, experience playing this role, and their overall experience with IT service delivery.

The survey was reviewed by important stakeholders in the organization and by some sysadmins from different departments to validate questions, eliminate misunderstanding, and align the set of options available in the answers with those relevant and adequate for the respondents. Again, the qualitative data collection and analysis performed before the survey had already provided valuable input for this, but the survey review was necessary to avoid any problems.

Employees were selected based on a stratified random sampling [15] approach based on the following *stratum*: department, expertise level, and gender. By doing so, we wanted to make sure that we covered the information needs of different types of workers from different IT departments (UNIX, security, databases, etc.), from different levels of expertise, work shift, and gender. The selection of the respondents was done by a project manager in the BSF organization to guarantee that participants' names remained unknown for the researchers.

An initial e-mail message about the survey was sent to all service delivery personnel in the Incident Management organization. Two days later, invitations to fill the survey were sent by e-mail to all selected employees, i.e., our respondents. About a week later a reminder was sent to the respondents and then, one day prior to the end of the survey, another e-mail message was sent to remind them again. Overall, respondents could fill the survey during a 14-day period during December of 2010. All e-mail messages were sent by the high level executive in charge of the delivery organization to attract respondents and all messages contained information about the anonymity of the survey. The survey had a 61% response rate with more than two hundred informants[1].

The data from the survey was extracted from the web-based survey tool and imported into a standard tool for statistical analysis (*SPSS*). We conducted statistical tests in the data, mainly chi-square, to calculate the descriptive statistics. After analyzing the data, a presentation with the main results, which includes the results reported in this paper, was created and presented to several important stakeholders in the BSF organization as a way to validate our results. The results discussed in this paper have all been discussed before with the BSF team members.

5. SURVEY RESULTS

We organize our results in three major parts. Initially, we describe overall descriptive statistics of the data. Following, we look at factors which influence the occurrence of dependency changes, such as project duration, team size, team distribution, CM usage, etc. Finally, the third part describes information collected about the respondents who have already faced dependency changes in the project.

5.1 Descriptive Statistics

Before conducting our analysis we removed missing, redundant, and wrong information. For instance, we had some open questions in the survey in which respondents could write additional comments or concerns. In a few cases, respondents who seem to not work directly with incidents answered the survey, in spite of all our care with the selection process and criteria. Therefore, we eliminated those respondents from this analysis. The data reported below is based on the data after this filtering process.

Of the respondents, 25% had an expertise level L1, 33% had a L2 level, while 23% had a L3 expertise level. However, about 19% of the respondents did not report their expertise at all. Respondents' experience working at the BSF and prior to it varied as presented in Table 1.

As for work shifts, 59% of respondents worked in the Morning shift, 23% worked in the Afternoon shift, and 18% worked during the Night shift, which is representative of the overall number of BSF employees working in different shifts. Finally, we also had good coverage regarding highest level of education and age. In general, the numbers above illustrate that the survey covered well the different levels of expertise and experience of the sysadmins,

[1] We cannot provide the exact number of respondents, since we already provided the response rate, to preserve our confidentially agreement with the BSF.

as well as other factors which could potentially influence our results.

5.2 Information Needs in Complex Tickets

In the survey, sysadmins were first asked to answer a set of questions about incidents in which the sysadmins had worked and which demanded a resolution effort above normal. As mentioned before, by *effort*, we told them to consider the amount of physical or mental energy and concentration level required of them while working with the incident.

Table 1. Respondents' working experience at the BSF.

	Experience at the BSF	Prior Experience
Less than a year	22.36%	5.83%
Between 3-1 years	34.18%	28.33%
Between 6-4 years	24.89%	16.67%
Between 10-7 years	7.59%	16.67%
More than 10 years	10.97%	32.50%

5.2.1 *Quality of the Information in Incident Tools*

As mentioned in section 2, work is assigned to sysadmins as single incidents, i.e., events that might disrupt the service in the customers' IT infra-structures. These incidents are managed as tickets in an Incident Management Tool (a.k.a, IPC tool) which stores information about when the ticket was created, the event, the IT component, the customer, and the person within the customer who reported the incident, among other pieces of information. So, one of our initial interests was to which extent the IPC tool provided enough information for the sysadmins to successfully work with an incident, i.e., for them to make sure the event associated with the incident did not cause any interruption in the quality of the service provided. Respondents answered that, in only about 33% of the cases the information from the IPC tool was enough, i.e., in 67% of the cases they needed to consult additional sources of information. We explored whether particular levels of expertise, shifts, and experience within the BSF were more associated with the need for additional information, but we did not find any significant relationship.

5.2.2 *Most Helpful Sources of Information*

In the following question, we asked sysadmins which other sources of information they used to gather additional information. Among the possible answers, we included the customer him/herself, the customer account team, sysadmins from the same department, among other options[2]. In particular, one of the options is the so-called *duty manager*, who is responsible for monitoring and handling crisis, i.e., high-severity incidents which require special attention from the corporation. Results describing the most used additional sources of information are presented in Figure 1. In this Figure, respondents could select more than one source of information. The median value of the number of information sources was 2.5; the minimum was 0 (i.e., no other source of information required); and the maximum was 9 (i.e., at least one sysadmin in our sample needed to consult nine different sources of information).

The following question asked respondents to indicate which information source they considered the most helpful for them while working in the incident. Results are presented in Figure 2. In this case, a sysadmin could only select one option. The results show that the most helpful sources of information are people who have knowledge about the customer, i.e., either the customer him/herself (21.33%), the customer account team (14.67%), or the customer expert (14%) from the same department from the respondent. These 3 sources account for 50% of the most helpful sources of information in complex tickets. In contrast, BSF's internal tools and sources of customer information are considered the most helpful by only 9.33% of the respondents.

A respondent could select "Other sources of information" in the question above, which represented 10% of the answers. So we included an open question asking them to detail which sources of information they considered more helpful. These open questions were then manually classified and aggregated in 7 different categories. The results of the 3 most frequent categories are: 44.44% of these respondents answered that some type of technical information was the most helpful source of information; 14.81% answered additional information about customers; and another 14.81% commented about generic sources of information (e.g., Internet, Google, etc). These last sources seem to indicate the means used by respondents to *find* the information instead of the *sources* themselves, which seems to be mostly technical. Since the survey was anonymous, we can not validate this hypothesis. In any case, the results still indicate – not surprisingly given the technical nature of the sysadmins work – the need for technical information as something relevant for the work of sysadmins.

We aggregated information from the two previously mentioned questions about the most helpful sources of information in Figure 3. To better visualize the main sources of information we created "categories" of information sources, namely: information sources related to customers account (54.99% of the answers)[3], artifacts and tools for technical and knowledge management (18.96%), other personnel (17.48%) and finally, what we classified as "facilitators", i.e., BSF employees whose only work was to facilitate the delivery of services (8.66%).

5.2.3 *Time to Access Sources of Information*

The survey included a question about the time sysadmins spent looking for their sources of information in the complex incident. More specifically, we asked sysadmins to estimate how much time, in minutes, the respondents needed to find and access the source of information that they considered the most helpful. Results from this question are presented in Figure 4 below.

In general, it is possible to notice that 71.77% of the sysadmins were able to access their most helpful source of information in less than 15 minutes. In contrast, 28.24% of the sysadmins spent more than 16 minutes in this same situation. We looked for

[2] This is a concrete example of how the qualitative part of the study influenced the survey.

[3] In this case we got a percentage of answers related with customer information that were in the "Other sources of information" option. We manually classified them as customer-related and added their percentage to the previous percentage previously reported. The result of this addition is 54.99%.

differences between the fastest sysadmins to find and access information (less than 2 minutes) and the slowest (between 31 and 60 minutes and more than 60 minutes). As for the technical expertise, we found no significant difference, i.e., there are sysadmins with all levels of expertise that are either the fastest or the slowest to find and access information. However, our analysis did not reveal any sysadmins from the Night shift among the fastest ones. Since the most helpful sources of information are people, this result is not surprising, since people with customer knowledge are not readily available in the Night shift.

5.2.4 *Additional Desired Sources of Information*

The survey had an open-ended question in which respondents could describe additional sources of information they needed. In this particular case, the question was framed in a way that respondents could focus on information needs about the customer. We used the same process as described before, i.e., we read all the answers from the respondents and then the first author classified them in different categories. These categories and their percentages are described in Figure 5.

The top most information required is better structure of information in current tools, i.e., in this case, respondents provided feedback on the current BSF tools including the incident management tool. As for the business and IT relationship, the second most requested source would be documentation "mapping" the IT components managed by BSF and the customer business. For instance, when a particular server is down, what does it mean for the customer? Similarly, when the customer has an incident in a particular business system, how does that business systems maps to the very large set of IT components being managed by BSF? The third most requested information by sysadmins are details about the customer contacts. In other words, a BSF customer might have hundred, or even thousands, of users which use their IT components. When there is an incident with one of those components BSF employees need to understand and often contact the user (among the hundreds of them) who are being, or might be, affected by the incident.

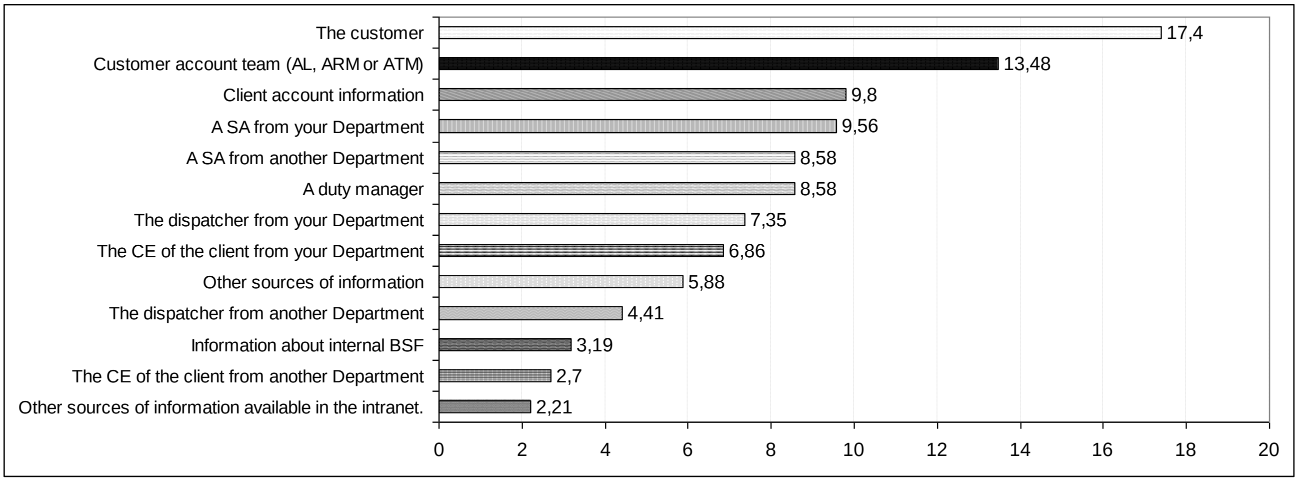

Figure 1 – Additional sources of information used by sysadmins.

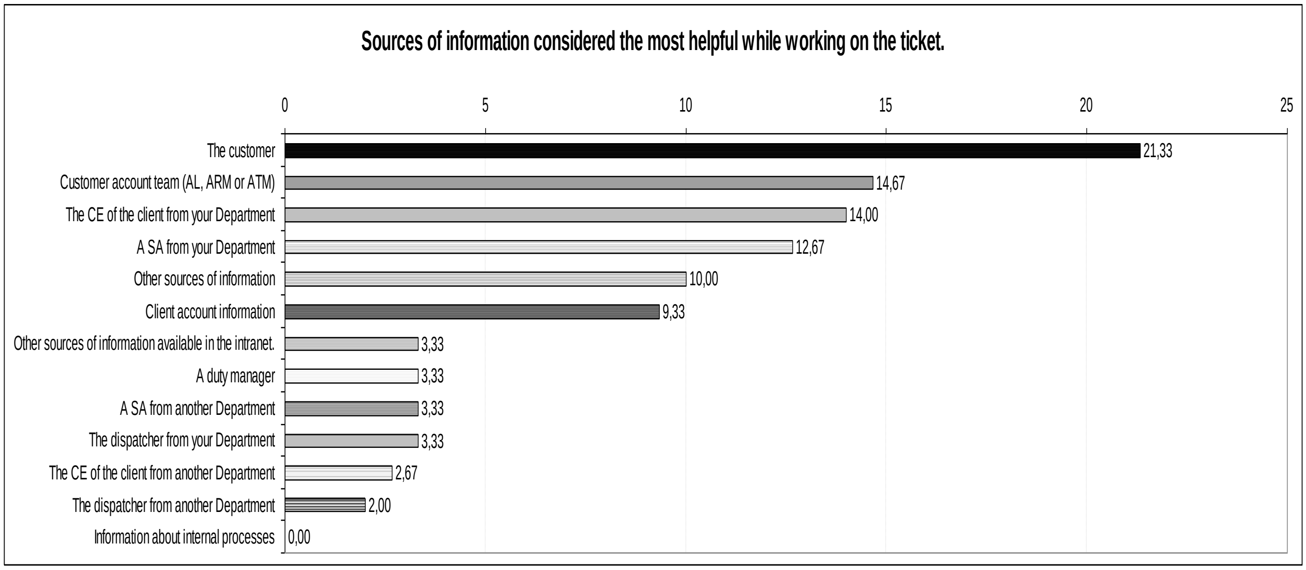

Figure 2 – The most helpful source of information used by sysadmins.

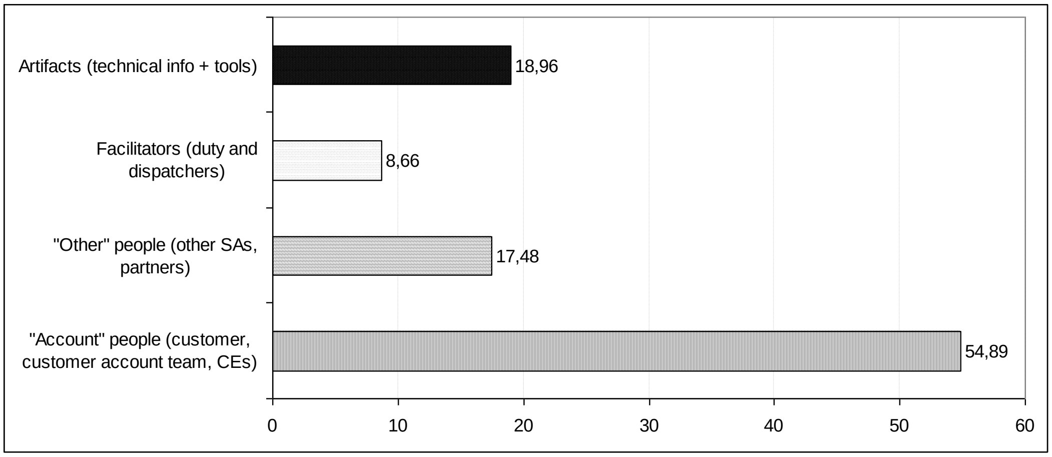

Figure 3 – Classification of the categories of the most helpful sources of information used by sysadmins.

5.2.5 Tool Usage

We asked respondents to inform all tools they used while working with a particular incident. Based on this question, as shown in Figure 6, we observed that the tools used to store customer information are barely used, i.e., only in 4.81% of the cases. This is in contrast with collaboration and communication tools like (telephone, instant messenger, etc) which were used by 59.36% of the respondents. This result is consistent with previous results which suggest that people, as opposed to tools and artifacts, are the most used and helpful sources of information for sysadmins [3, 24] working with complex tickets. On the other hand, instant messaging is a tool widely used at the BSF. Given the BSF's geographical dispersion (see section 3), these results are hardly surprising.

5.3 Information Needs in Routine Work

The survey had a set of questions asking respondents their daily work, asking them to rate the information sources used by the sysadmins in their routine work with incidents. In this case, the focus was not in complex tickets, but in sysadmins' daily work. The information sources evaluated in the survey were: sysadmins from their own departments, from other departments, members of the customer account team, and the tools and technical documentation about the customers. The questions were structured based on a 5-point Likert scale ranging from to "Strongly Agree" to "Strongly Disagree". Most respondents assessed positively these information sources with more than 90% of positive evaluations (either Agree or Strongly Agree). The only exception was the tools and technical documentation that had about 14% of negative and neutral evaluations. In this case, we observed that the sysadmins who not rated this information source positively were predominantly sysadmins who worked in the Night shift with the level of expertise L1.

Finally, we asked respondents how many hours per week they spent dealing with requests for information about the customers from their colleagues. The results are presented in Figure 7. We looked for differences in age, gender and previous experience in service delivery among the respondents and found no significant result. On the other hand, we observed that a set of sysadmins spent more than 4 hours a week providing this information for their colleagues, i.e., more than 10% of their weekly time with these requests. This amounts for 27.27% of Level 3. This is a somewhat expected result when we take into account that Level 3 employees are those who have the highest levels of technical expertise. L2 employees have the second highest technical expertise in the organization and 21.74% of them spent more than 4 hours a week with requests for information. At the same time, 17.39% of these L2 respondents disclosed that they are not requested for help at all (the column in the far right of Figure 7). This suggests an imbalance in the efforts of L2 employees, i.e., while some of them are potentially overwhelmed with requests for information, other are in a more favorable situation not being disrupted with requests for help.

6. DISCUSSION

Recently, due to the failure of several knowledge management (KM) systems [9], the KM research community has started investigating styles of knowledge management [10] aiming to find out whether the failures and/or successes were related to these styles. Furthermore, some researchers [16] recognized that in order to be successful, KM should also emphasize human aspects: cognitive, social, cultural, and organizational. Researchers who adopted this approached called it *expertise sharing* to differentiate it from previous approaches in KM [16]. In this case, rather than focusing on the management level of an organization, expertise sharing focuses on the self-organized activities of the organization's members. In short, modern KM researchers recognize the existence of two main approaches: *"the first approach focuses on explicit knowledge and, thus, emphasizes the capability to help create, store, share, and use explicitly documented knowledge, while the second focuses on tacit knowledge and emphasizes knowledge sharing by interpersonal interaction."* [10]. Based on that view, Choi and Lee [10] identified three main KM styles. The first one, called *system-oriented* is characterized by the "explicit" knowledge approach, the second focuses on "tacit" knowledge and interactions, and is called *passive* approach, and finally, the third approach, called *dynamic*, equally emphasizes both the "explicit" and "tacit" knowledge. Not surprisingly, Choi and Lee found out that companies which adopt this dynamic approach have better performance.

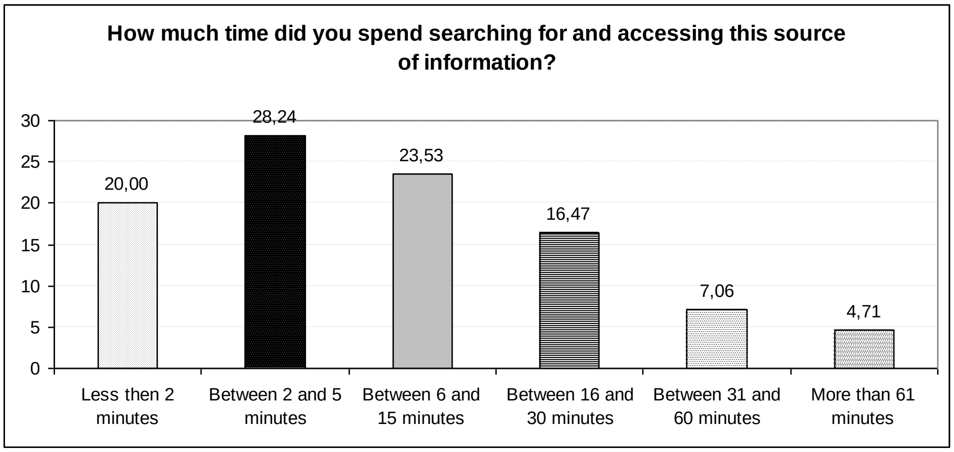

Figure 4 – Time spent finding and accessing the most helpful source of information

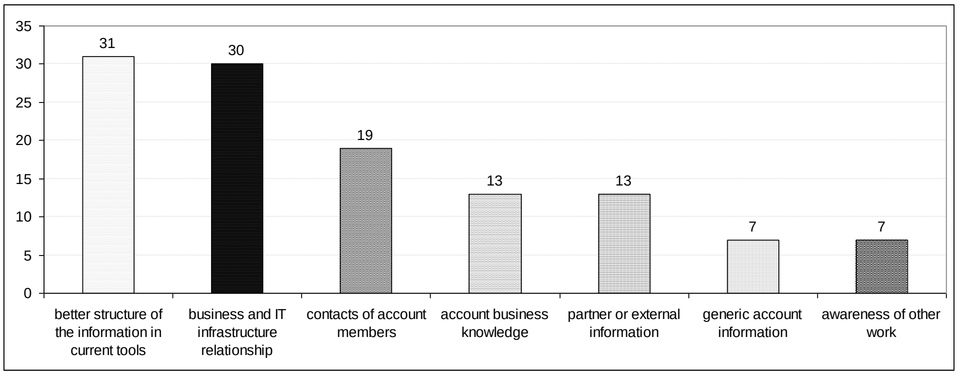

Figure 5 – Sources of information about the customer who were desired by the respondents.

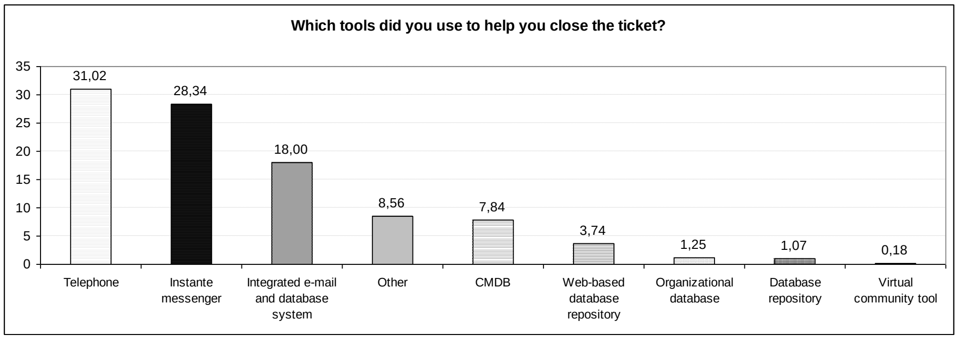

Figure 6 – Usage of tools by sysadmins working in complex tickets at BSF.

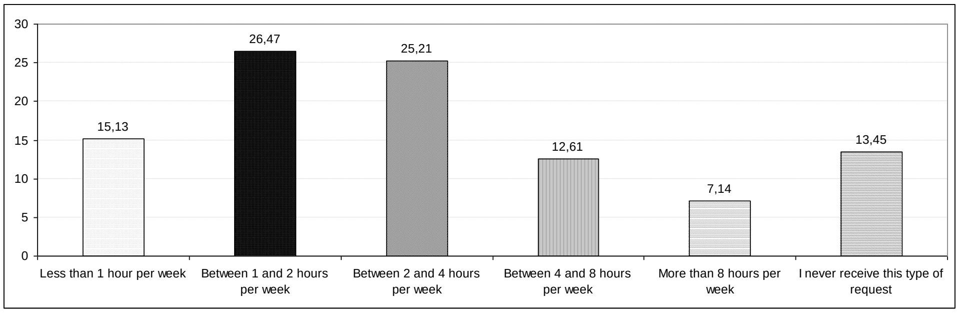

Figure 7 – Number of hours spent per week by sysadmins dealing with information requests about customers.

In the previous section, we found the following situation in the BSF: (i) low usage of KM tools, (ii) high usage of communication and collaboration tools, and (iii) sysadmins' need of gathering information about a customer from a specific set of stakeholders. These findings suggest that the BSF adopts a passive approach for knowledge management. Furthermore, according to our respondents, information about the customer is absolutely critical for the successful delivery of IT services. Again, this is based on the fact that the most useful sources of information for employees are: (i) the customer, (ii) the customer account team, the group of employees responsible for representing this customer inside the organization, and (iii) other employees who were experts both in the customer and in particular aspects of the delivery of services.

As pointed out by other researchers [9, 10] and mentioned by stakeholders in the BSF during the validation of the study results, the passive approach for knowledge management is limited because the costs of maintaining such approach usually increase as the number of different customers and employees involved increases. Furthermore, while in general the time to find and access helpful sources of information does not seem to be a problem, the increased time that is required by sysadmins during the Night shift seems to indicate that this is, at least, something the BSF needs to be cautious about. Therefore, a dynamic knowledge management style emphasizing both knowledge reusability through information technologies and knowledge sharing through informal discussions among employees seems to be a better fit for the BSF and other IT service factories.

If we take a point of view from Service Science, this need to obtain information from the customer during the delivery of the service is not a surprising result, since value is co-created [11], i.e., in services, the value the customer obtains from a provider is created during the interaction between customer and service provider. What is surprising, however, is the lack of support in the BSF tools, approaches and processes for this aspect: most of them (including the incident management tool used at the BSF) provide only limited information about the customer. Their entire focus is on the technical aspects of work, e.g., providing technical information for sysadmins. However, often this information is difficult to be contextualized for a particular customer and its specificities [18]. Because of this lack of support, sysadmins must rely on other people as sources of information. Note that we are *not* arguing that sysadmins' need to consult additional sources of information while working with incidents is not relevant: the work of sysadmins is highly technical so we expect them to consult *technical* sources of information. Instead, we are arguing about the need to provide *contextualized* technical information since the details, configurations, settings and other information from each customer do impact how technical work is performed.

It also should be noted that the complexity required to handle such situations seems to be quite beyond of traditional Customer Relationship Management tools which tend to focus mostly on the marketing relationships between companies. Overall, our results point that customer knowledge management in IT service factories is an important area for research, albeit hardly explored.

Other results of this research might also be used to inform the design of tools for IT service factories. In particular, our results indicate that the information necessary for handling complex incidents includes information not only about the customers, their IT infrastructure and critical business processes, but also about the *mapping* of these business processes to IT components. In this case, approaches based on collective intelligence [12] like crowdsourcing [13, 14] are potentially suitable given the scale of the problem, i.e., the number of IT components and the number of sysadmins who have, or want to have, information about these components.

In addition to solely technical and contextualized technical information, sysadmins also need information about customer contact people and decision-makers. If, for instance, a customer's IT component is problematic, it is important that the sysadmins and the entire IT service factory understand how this component affects the business of the customer organization (i.e., the mapping), but also who in the customer organization to work with to properly correct the problem, while maintaining the customer informed about the progress of the work. And, again, despite this need, the BSF tools provide limited support for that. In this case, while the work for finding out information about other people is not the Nardi's networking [17], some of the tools proposed to do so might still be relevant. We believe that all these aspects are relevant to collaborative tool builders who are designing systems to be used to provide services.

Finally, our results suggest that while information management work is indeed collaborative [2, 3, 20], there are significant differences between sysadmins: while some are potentially overwhelmed being an important source of help for colleagues, others are not as helpful. This result is similar to recent results in software development [21, 22] which also point out that there is a difference between the coordination efforts of professionals. This raises the question of how current tools for sysadmins should be designed: currently they are designed assuming that the collaborative effort of individuals is very similar, which is not the case. That is an interesting research aspect that we plan to explore in our future work.

7. CONCLUSIONS AND FUTURE WORK

IT service factories are organizations that often employ hundreds or thousands of employees to deliver Information Technology services for customer organizations. Such factories arose in the past decades and face several challenges due to their scale. As mentioned, past research uncovered the knowledge-intensive and collaborative nature of the work performed by sysadmins. In this paper, we extend previous work with data collected from an empirical study conducted at a large-scale IT service factory.

First, we report that part of the knowledge necessary for sysadmins to perform their work is related to the customer to which they are providing services. In other words, sysadmins seek information from stakeholders who have knowledge about the customers. In fact, they seek different types of information about the customers; information that is not currently available in the tools, approaches and methods used at the IT service factory. Second we report in this paper is about the collaborative nature of sysadmins work: while there are some sysadmins potentially overwhelmed with requests for information, others do not face this challenge. These results can and should be used to inform the design of tools to be used by sysadmins in IT service factories and similar environments.

8. REFERENCES

[1] Gonzalez, V. M., Galicia, L., & Favela, J. (2008). Understanding and supporting personal activity management by IT service workers. *ACM Symposium on Computer Human Interaction for Management of Information Technology* (p. 1). New York, New York, USA: ACM Press.

[2] Barrett, R., Kandogan, E., Maglio, P. P., Haber, E., Takayama, L. A., & Prabaker, M. (2004). Field studies of computer system administrators: analysis of system management tools and practices. *Proceedings of the 2004 ACM conference on Computer supported cooperative work* (pp. 388-395).

[3] Haber, E. M., Kandogan, E., & Maglio, P. P. (2011). Collaboration in system administration. *Communications of the ACM*, *54*(1), 46.

[4] Kandogan, E., Haber, E. M., Bailey, J. H., & Maglio, P. (2009). Collaborative Work: The Case of IT Service Delivery. *Proceedings of the 13th International Conference on Human-Computer Interaction* (pp. 504-513).

[5] http://www.itlibrary.org/index.php?page=Incident_Management ITIL Incident Management - The ITIL Open Guide.

[6] Scott, Richard. ORGANIZATIONS: RATIONAL, NATURAL, AND OPEN SYSTEMS. 5. ed. New Jersey: Prentice Hall, 2003. 430 p.

[7] McCracken, G., *The Long Interview*. 1988, Thousand Oaks, CA: SAGE Publications.

[8] Jorgensen, D.L., *Participant Observation: A Methodology for Human Studies*. 1989, Thousand Oaks, CA: SAGE publications.

[9] Carlile, P. R. A Pragmatic View of Knowledge and Boundaries: Boundary Objects in New Product Development, *Organization Science* 13 (4) (2002): 442-455.

[10] Choi, B., & Lee, H. (2003). An empirical investigation of KM styles and their effect on corporate performance. *Information & Management*, *40*, 403-417.

[11] Vargo, S. L., & Lusch, R. F. (2004). Evolving to a New Dominant Logic for Marketing. *Journal of Marketing*, *68*(1), 1-17. doi: 10.1509/jmkg.68.1.1.24036.

[12] Malone, T. W., Laubacher, R., Dellarocas, C. N., Harnessing Crowds: Mapping the Genome of Collective Intelligence (February 3, 2009). MIT Sloan Research Paper No. 4732-09.

[13] Vukovic, M. Laredo, J., Rajagopal, S.: Challenges and Experiences in Deploying Enterprise Crowdsourcing Service. International Conference on Web Engineering: 460-467, Vienna, Austria, 2010.

[14] Lopez, M., Vukovic, M. Laredo, J.,: PeopleCloud Service for Enterprise Crowdsourcing. IEEE International Conference on Services Computing, 538-545, Miami, FL, USA, 2010.

[15] Wild, C., Seber, G. Chance Encounters: A First Course in Data Analysis and Inference, John Wiley & Sons, NY, 1999.

[16] Ackerman, M., Pipek, V., Wulf, V. Sharing Expertise: Beyond Knowledge Management, MIT Press, 2003.

[17] Nardi, B., Whittaker, S., Isaacs, E., Creech, M., Johnson, J., Hainsworth, J. (2002). ContactMap: Integrating Communication and Information Through Visualizing Personal Social Networks. *Communications of the Association for Computing Machinery (CACM)*. April, 2002.

[18] Velasquez, N.F., Weisband, S. P. System administrators as broker technicians. In *Proceedings of the Symposium on Computer Human Interaction for the Management of Information Technology*. ACM, New York, NY, USA, 2009.

[19] Velasquez, N.F., Durcikova, A. Sysadmins and the need for verification information. In *Proceedings of the 2nd ACM Symposium on Computer Human Interaction for Management of Information Technology*, ACM, New York, NY, USA, 2008.

[20] Haber, E. M. 2008. System administrator teamwork: evidence from the SAGE salary survey. In *Proceedings of the 2nd ACM Symposium on Computer Human Interaction for Management of Information Technology*. ACM, New York, NY, USA, 2008.

[21] Cataldo, M., Wagstrom, P., Herbsleb, J. and Carley, K. (2006). *Identification of Coordination Requirements: Implications for the Design of Collaboration and Awareness Tools*. In Proceedings of the Conference on Computer Supported Cooperative Work, Banff, Alberta, Canada.

[22] Costa, J.M.R, Cataldo, M. and de Souza, C.R.B (2011). The Scale and Evolution of Coordination Needs in Large-Scale Distributed Projects: Implications for the Future Generation of Collaborative Tools. in Proceedings of the International Conference on Human Factors in Computer Systems, Vancouver, Canada.

[23] Strauss, A., & Corbin, J. (1998). Basics of qualitative research: Techniques and procedures for developing grounded theory. Thousand Oaks: SAGE.

[24] Werlinger, R., Hawkey, K., Muldner, K., & Jaferian, P. (2006). The Challenges of Using an Intrusion Detection System : Is It Worth the Effort ? In Proceedings of the Symposium on Usable Privacy and Security, Pittsburgh, PA, July 23-25, 2008, 107-116.