Information Needs of System Administrators in Information Technology Service Factories

Cleidson R. B. de Souza, Claudio S. Pinhanez, Victor F. Cavalcante
IBM Research Brazil
Rua Tutóia, 1157
São Paulo, SP, Brazil, 04007-005
cleidson.desouza@acm.org, csantosp@br.ibm.com, victorfc@br.ibm.com

ABSTRACT

In this paper we describe the results of an empirical study about the information needs of system administrators. This study is based on an electronic survey with more than 200 systems administrators, or sysadmins, working on incident management in a large scale IT service delivery organization. The survey covered their information needs in both complex and routine situations. The results of the survey described in this paper go beyond previous work on system administrators by presenting a much more complex picture, suggesting that sysadmins make low usage of knowledge management tools; largely adopt personal communication and collaboration tools; and finally, need to gather information about customers from a complex set of stakeholders. The system administrators also indicated in our survey that the most useful sources of information in handling complex incidents are: (i) the customer; (ii) the customer account team; and (iii) other employees who were experts both in the customer and in particular aspects of the delivery of services. This study indicates that knowledge management in IT service factories is very challenging and possibly should evolve from the often adopted passive model to a dynamic knowledge management style emphasizing both knowledge reusability through information technologies and knowledge sharing through informal 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

knowledge management, information technology management, IT service factories, system administrators, empirical study.

ACM CHIMIT 11, December 4, 2011, Boston, MA, USA.

1. INTRODUCTION

The last two decades witnessed a profound transformation in the ways computing and data processing are provided to organizations. While the in-house datacenter dominated the early decades of computing, many firms today outsource the management of their Information Technology (IT) infrastructure to other organizations which provide what is known as *Infrastructure as a Service (IaaS).*The delivery of IaaS is often made through gigantic IT service operations where hundreds or even thousands of support personnel take care of a network of thousands of servers, routers, and other IT equipment, often from multiple firms at the same time.

We refer to those IT service operations as IT service factories. Those organizations, often employing hundreds, sometimes thousands of specialized IT workers, are spread all over the world, but especially in India and other emerging countries like Brazil. The use of the term factory is purposively chosen to reflect the fact that those organizations borrow a great deal of structural resemblance to traditional manufacturing factories. Workers in IT service factories are very specialized, often focused on specific tasks repeated exactly many times during the day, and often with performance metrics that resemble the early days of Ford.

In IT service factories, a group of professionals often referred to as *system administrators*, a.k.a. *sysadmins* or *SAs* are responsible for some of the most critical functions necessary to maintain the customers' IT infrastructures running well and efficiently. System administrators are responsible for understanding the problems affecting IT systems, fixing them, and for preventively taking actions to keep the systems running without downtimes; for installing, updating, and upgrading the software installed in the machines; for protecting the system against attacks and other threats; and to manage thousands of user accounts.

System administrators in IT service factories differ from their counterparts in in-house datacenters because they usually work in large workgroups with specialized goals and competencies, sometimes supporting systems from multiple customers. As discussed by Gonzalez and colleagues [1], system administration in IT service factories are highly intensive knowledge jobs. Moreover, sysadmins have to find, understand, interpret, and apply not only technical knowledge, but also knowledge which is specific about each customer and its IT infrastructure.

Only recently researchers have started to pay attention to the work of system administrators in the context of service factories [2, 19]. Examples of previous work include the study by Haber and colleagues [3] and Botta and colleagues [4] that argue for the collaborative nature of sysadmins' work. Another example is the

work of Gonzalez and colleagues [1] that indicates how IT workers' activities are varied, fragmented, and overlapped.

While previous work has discussed important aspects of sysadmins' work, including their collaborative and knowledge-intensive nature, to the best of our knowledge, previous research has overlooked sysadmins' information needs. In other words, little is known about the type of information sysadmins seek, how often they seek information, and how much time sysadmins spend seeking information.

This paper reports on an empirical study conducted in a large IT services provider organization. This organization provides IT services for a variety of other large organizations through outsourcing contracts. 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 employing hundreds of system administrators around the clock to support and improve the information infrastructure of the customer organizations.

This paper examines the information and knowledge needs of system administrators using both qualitative and quantitative methods. The qualitative methods used include interviews and non-participant observation and were used in an exploratory manner to identify relevant problems in the work being conducted by sysadmins in an IT service factory. The quantitative method used was a survey where professional sysadmins from the same service factory were asked questions about their information needs while working with the maintenance of large complex IT systems. Questions in the survey were designed to gather details about the type of information needed by the sysadmins, the resources in which this information was (or was not) available (e.g., documentation, co-workers from the same or different departments, etc), the *effort* that the sysadmins spent in order to gather such information, and the frequency in which sysadmins needed to seek information.

Our study provides two major insights about the nature of the information needs from sysadmins in IT service factories. First, information and knowledge about the customer organization is absolutely critical for the successful delivery of IT services. As a matter of fact, our survey revealed that in complex situations, the most important sources of information for sysadmins 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. In other words, according to the sysadmins themselves, the most relevant information to help them deliver services is directly associated with the customer. Our second result is related to the collaborative nature of the work of IT sysadmins. As expected, our results corroborate previous work indicating that the work of sysadmins is highly collaborative [2, 3, 4, 20]. However, it suggests that there can be an imbalance in the work of sysadmins: while some professionals are overwhelmed spending several hours per week providing information for their colleagues, others are in a more comfortable situation. We looked at the expertise of these professionals, and that does not account for this difference between IT professionals.

The rest of this paper is organized as follows. Given that the IT service factory model of delivering IT services to customers is still unknown by many researchers, we start the paper by describing IT service factories and how work is carried out in these organizations (sections 2 and 3). In the following sections, we present the methodology used for both the qualitative and the quantitative study (section 4), and their results (section 5). Section 6 presents an analysis of our results. Finally, section 7 presents our conclusions and ideas for future work.

2. IT SERVICE FACTORIES

In the current highly competitive environment for IT outsourcing, IT service factories operate based on economies of scale. By employing hundreds of sysadmins and serving dozens of customers simultaneously, IT service factories can reduce costs by sharing sysadmins and infrastructure among their customers.

IT service factories are often structured in sub-organizations which handle different functions and needs of the IT infrastructure of the customer. A particular 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*." [6]

The incident management department of the BSF is responsible for some of the key tasks needed to maintain the customers' IT infrastructure, including restoring normal operations as quickly as possible with the least possible impact on the customers' business or user. Incidents are one of the smallest units of work in IT service factories, in such a way that each employee working in the delivery of services usually works on incidents for any of the different customers which have contracts with the service factory.

The quality of the service provided by IT service factories is regulated by the contracts between the IT provider and its customers. Those contracts often determine clearly how fast incidents have to be investigated and solved, to guarantee the least possible disruption in IT infrastructure. Often, the contracts establish formally Service *Level Agreements*, or *SLAs*, clearly determining the maximum time in which a certain type of incident must be dealt with and solved. In most situations, different types of incidents and IT components have different target solution times. Also, it is common to establish in those contracts a classification scheme for incidents, determining different levels of severity which correspond to the impact of the incident in the customer organization functions. 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 in the incident management database.

3. THE BIG SERVICE FACTORY (BSF)

This section describes the organizational structure [7] of the service factory we studied. As mentioned before, for confidentiality reasons, we call this factory the *Big Service Factory* (*BSF*). The anonymity of the IT service provider is necessary due to the nature of our study and its possible effect in a competitive landscape. BSF has first-class service quality levels which position the BSF among the top IT providers in the world.

BSF provides IT services for a variety of large organizations through outsourcing contracts in the U.S., Europe, Latin America, and Brazil. In this paper we call these organizations simply as the

customers of the BSF. There is a large difference among customers of the BSF, especially related to their size: while some of them have hundreds of IT components (servers, network equipments, etc), others only have dozens of them.

Following a common way to structure organizations [7], system administrators at BSF are assigned to work in *departments* where 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. Since the organizational structure is based on competence areas, in a given day, a sysadmin might work with incidents from completely different customers.

BSF implements almost all the principles proposed by the *ITIL* [6] 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. 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, BSF sysadmins are classified in three levels according to their expertise: *Level 1 (L1)*, *Level 2 (L2)*, and *Level 3 (L3)*. Level 1 sysadmins have limited technical knowledge on the competence area of their department and handle incidents that are very simple (from a technical point of view) and that often are 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 *customer experts (CEs)*, responsible for understanding in more details the IT environment of particular customers, and often knowledgeable of the IT and organizational structure 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 or selling additional services to current customers. 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, more importantly, monitoring and discussing the quality of the services delivered to the customer by the BSF. This team is composed of three different roles which are responsible for interacting with the customer and with the BSF internal service delivery teams at different organizational levels. 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. The third and last 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 is understood as representing a particular customer while interacting with the BSF competence teams, but representing the BSF as a whole while interacting with customers.

Finally, it is important to mention that BSF employees are spread in two different sites 100 kilometers apart from each other. Most of the sysadmins who work in the delivery of the services are located in one large site while some members of the customer account teams work in a smaller site, close to the sales organization. Also, some members of the customer account team work collocated with their customers. In other words, members of the service delivery, management, and customer account teams are spread through three different places making face-to-face communication often quite difficult.

4. METHODOLOGY

To study the knowledge and information 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 [24] characterized by collection and analysis of *qualitative* data followed by collection and analysis of *quantitative* data. The main activities of this first study are described in the following subsections.

4.1 Qualitative Data Collection and Analysis

The qualitative approach used in this study consisted of more than 20 semi-structured interviews [8] and more than 20 hours of non-participant observation [9] conducted over a period of about eight months. In this period, employees from different departments of the IT service organization and from different organizational positions were interviewed and observed. We were not able to record interviews, so we wrote down notes for later analysis.

As for the observations, we used non-participant observation [9] of employees who worked solely with incident management and 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 informant sysadmin, distant enough to not disturb the informant but close enough to be able to observe the contents of physical or digital objects which the informant was handling. Information collected during observation was recorded without distracting the informants. In some occasions, the observer asked clarification questions during the observations, not without first ensuring that such questions would not disturb the informant's activities. If it was not possible to ask clarification questions during the observations, those questions were asked later during interviews at the end of the day or during work breaks.

Observation was an important method to be used in this research because it allowed us to understand the overall context of work during the daily delivery of services, for instance, understanding which tools the informants used and how they used them. Moreover, it was a source of valuable insights to be explored during the survey. 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 [11] (see 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 Quantitative Data Collection

As mentioned before, a large scale survey was conducted at BSF. The goal of the survey was to collect quantitative data about the information needs of IT delivery personnel. We restricted the survey to BSF employees who worked in technical departments and, more specifically, in incident management, since we were interested in system administrators. 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.

The survey consisted of 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 additional effort from them. By effort we mean, the amount of mental and physical energy spent while working in the ticket. As part of the survey we collected demographics such as gender, age, education, employment status, etc. about the informants. We also collected information related to the work context of the respondents: work shift, job role, experience playing this role, and their experience with IT service delivery.

Informants were selected among BSF employees based on a stratified random sampling [16] 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.

The survey was reviewed by key stakeholders in the BSF 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.

In an enterprise it is important to create the perception among respondents that top-level management is interested in a survey, so employees feel more motivated and sanctioned to take time to answer the survey. Therefore, an initial e-mail message about the survey was sent to all service delivery personnel in the Incident Management organization by a high-rank manager. Two days later, invitations to fill the survey were sent again 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 answered the survey during a 14-day period during December of 2010. All e-mail messages clearly contained information stating 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*). Statistical tests were conducted 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 parts. Initially, we describe overall descriptive statistics of our data. Then, we describe sysadmins' information needs in complex tickets. Finally, we describe sysadmins' information needs in routine work (tickets).

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 do not work directly with incidents answered the survey, thus not conforming to our selection target criteria. Therefore, we eliminated these respondents from our 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 (see Section 3). However, about 19% of the respondents did not report their expertise at all. Respondents' experience working at the BSF and prior to it is presented in Table 1.

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%

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. These results are representative of the overall number of BSF sysadmins working in the delivery of services during different shifts. Finally, we also had good coverage regarding the highest level of education and age. In summary, we believe the survey covered well the different levels of expertise and experience of the sysadmins, as well as other factors which could potentially influence our results.

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

5.2 Information Needs in Complex Tickets

In the survey, sysadmins were first asked to answer a set of questions about incidents in which they worked and which demanded a resolution effort above normal (see section 4.2).

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 customers' IT environment. These incidents are managed as tickets in an Incident Management Tool (a.k.a, IPC tool) which stores information such as when the ticket was created, the observed problem, the IT component where the event happened, 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 find out to which extent the IPC tool provided enough information for the sysadmins to successfully work with an incident. 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

Another aspect we asked sysadmins was about which other sources of information they used to gather additional information to work with the tickets. Among the possible answers, we included people from the customer's organization, 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 high-severity incidents which require special attention from the organization. Results describing the most used additional sources of information are presented in Figure 1.

In the data presented in Figure 1, respondents could select more than one source of information. The median value was 2.5, while 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 several different sources of information).

The following question of the survey 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. Results show that the most helpful sources of information are people who have knowledge about the customer, i.e., either people from the customer's organization (21.33%), the customer account team (14.67%), or internal customer experts (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, artifacts, and other sources of customer information are considered the most helpful by only 9.33% of the respondents.

Respondents could also select "Other sources of information" in the question above, and did so in 10% of the answers. In this case, they could detail which source of information they considered the most 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). This last source of information seems to indicate the *means* used by the respondents to *find* the information instead of the *sources* of information themselves. Our results indicate – not surprisingly given the technical nature of the sysadmin 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. The figure shows related "categories" of information sources, clearly visualizing the main sources of information, namely: information sources related to customers, accounting for 54.99% of the answers[3]; artifacts and tools for technical information and knowledge management (18.96%); other personnel (17.48%); and finally, what we classified as "facilitators", i.e., BSF employees whose main work is to facilitate and manage the delivery of services.

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.

In general, when we aggregate several of the options, 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 differences between the fastest respondents 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, in a simple analysis, we did not find any sysadmins who worked in the Night shift being among the fastest ones. Since the most helpful sources of information are people, this result is not surprising, once people with customer knowledge are not readily available in the Night shift.

[2] This is a good example of how the qualitative part of the study influenced the quantitative part: several of these options were based on the observations we have conducted in the first part of this study.

[3] In this case we got a percentage of answers related with customer information that were in the "Other sources of information" answer. After reading these answers, we classified part of them as customer-related, technical information and so on. The customer-related percentage was added to the original 50% previously reported (from the customer, customer account team, and customer experts from the same department). The result of this addition is 54.99%.

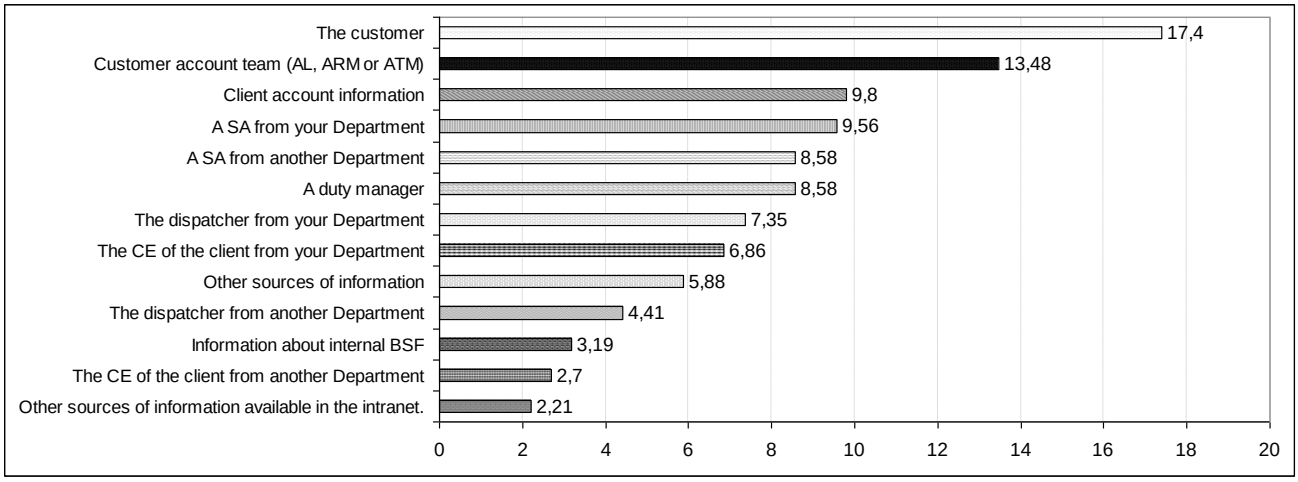

Figure 1 – Additional sources of information used by sysadmins.

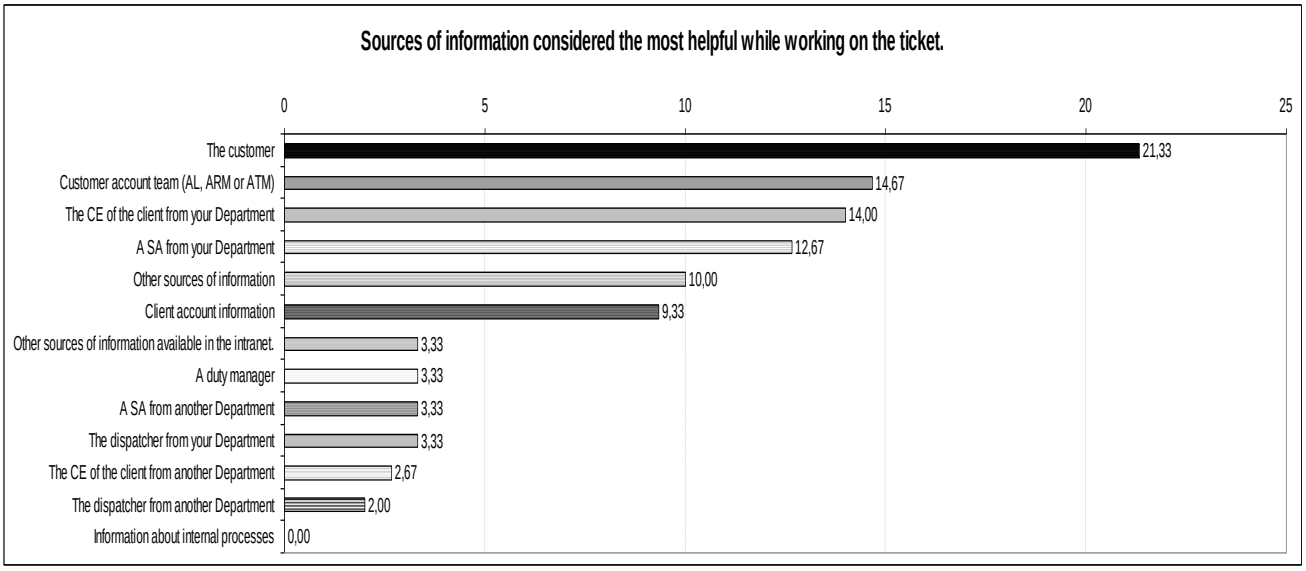

Figure 2 – The most helpful source of information used by sysadmins

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 classified them in different categories. These categories and their percentages are described in Figure 5.

The top most source of information required is in fact a better structure of information in current tools. In other words, respondents used the survey to provide negative feedback on the current BSF tools including the incident management tool.

The second most requested source is some sort of documentation "mapping" the IT components managed by BSF and the customer business systems and processes. For instance, when a particular server is down, what does it mean for the customer? Or 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 the 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 employees which use their IT environment. When there is an incident with one of those components, BSF sysadmins need to understand and, often, contact the user (among the hundreds of them) who are being, or might be, affected by the incident. Current IPC tools used at BSF do not always provide useful information to facilitate the identification of customer contacts.

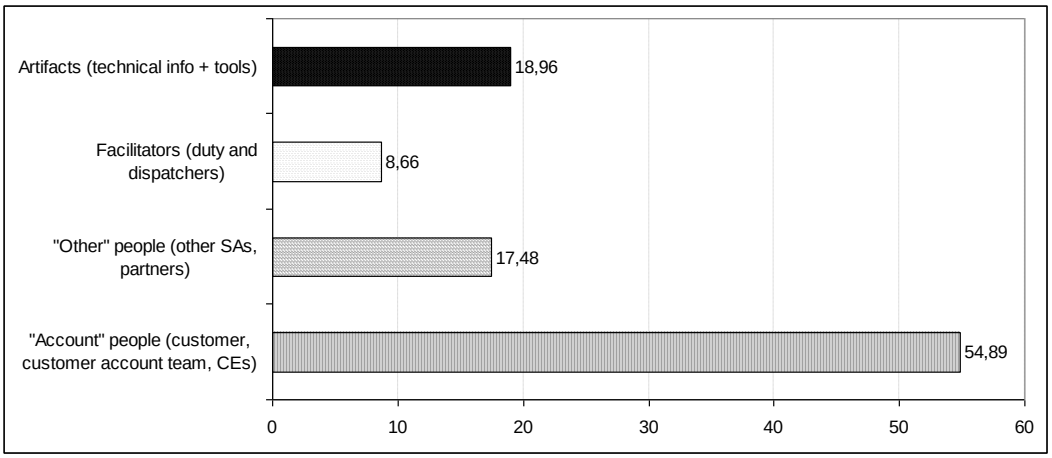

Figure 3 – Classification of the categories of the most helpful sources of information used by sysadmins.

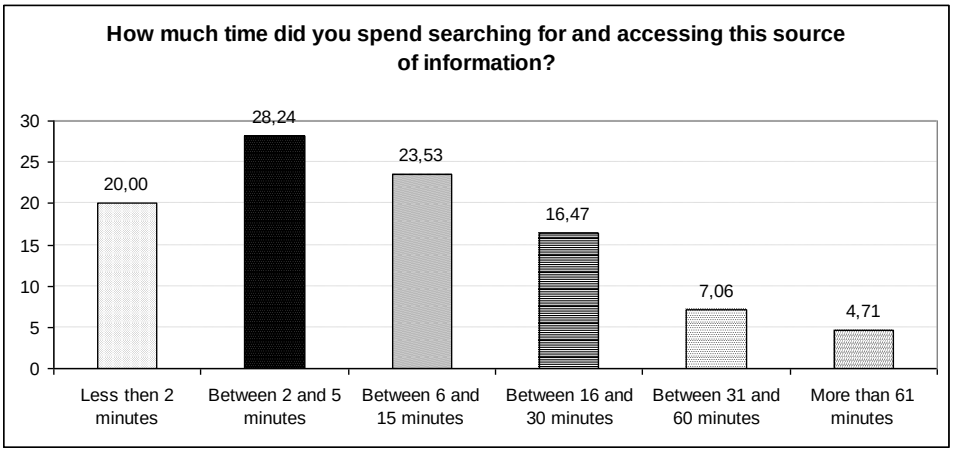

Figure 4 – Time spent finding and accessing the most helpful source of information

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 our results described before which suggest that people, as opposed to tools and artifacts, are the most used and helpful sources of information for sysadmins 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 easy to understand.

5.3 Information Needs in Routine Work

The survey had a second set of questions asking respondents their routine work with incidents, in contrast to complex tickets (see previous section). It asked them to rate the information sources used by the sysadmins in this case. 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. Questions were structured based on a 5-point Likert scale ranging from to "Strongly Agree" to "Strongly Disagree". Most respondents assessed positively those 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 neutral and negative 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 customer-related information from their colleagues (see 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 respondents. 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 the L2 respondents are not requested for help at all (see the far right column 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, others 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 [10], KM researchers have started to investigate styles of knowledge management [11] in order to find out whether the failures and/or successes were related to these styles. Some researchers (e.g., [17]) recognized that in order to be successful, knowledge management should also emphasize the human aspects —cognitive, social, cultural, and organizational— of knowledge management. Rather than focusing on the managerial level of an organization, expertise sharing focuses on the self-organized activities of the organization's members. Therefore, modern knowledge management approaches recognize the existence of two main KM 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."* [11]. Based on this view, Choi and Lee [11] identified three main 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*, and finally, the third approach, called *dynamic*, emphasizes both the "explicit" and "tacit" knowledge. Not surprisingly, Choi and Lee found out that companies which adopt this dynamic approach have better performance.

In the previous section, we found the following situation at BSF: (i) low usage of KM tools, (ii) high usage of communication and collaboration tools, and (iii) sysadmins' need of gathering customer-related information from a specific set of stakeholders. These findings suggest that the BSF currently adopts, implicitly or explicitly, a passive approach for knowledge management. Furthermore, according to our respondents, information about the customer is absolutely critical for the successful delivery of services. Again, this is based on the fact that the most useful sources of information for sysadmins 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.

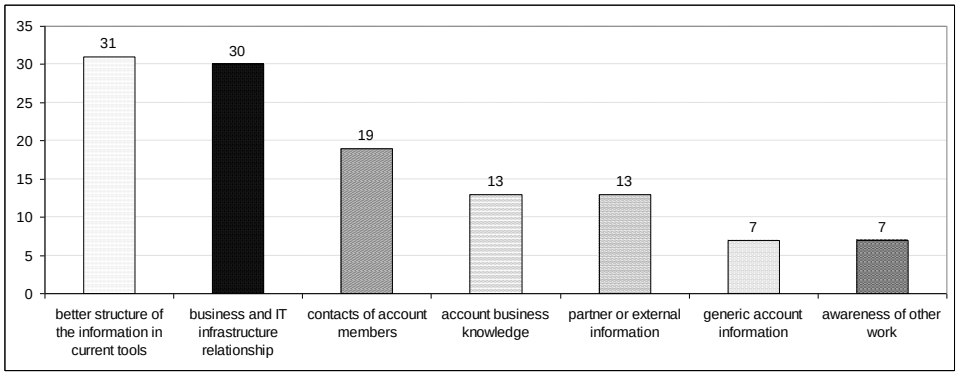

Figure 5 – Sources of information about the customer desired by the respondents.

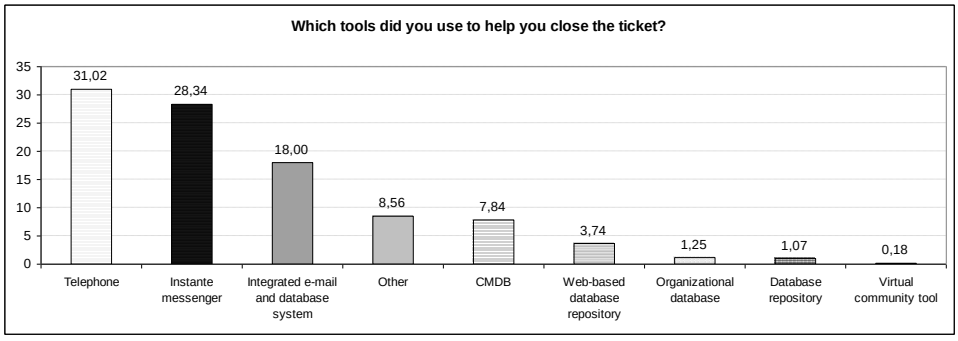

Figure 6 – Usage of tools by sysadmins working in complex tickets at BSF.

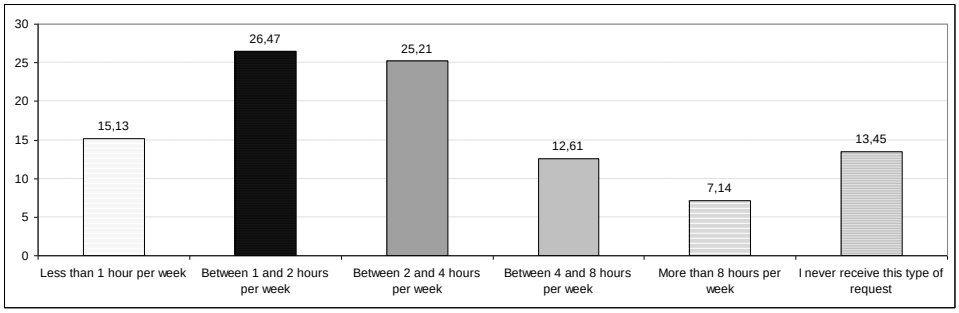

Figure 7 – Number of hours spent per week by sysadmins dealing with information requests about customers.

However, as pointed out by other researchers [10, 11] 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. Our work focused solely on BSF, but we believe that other IT service factories might have similar results.

If we take a point of view from Service Science, the need to obtain information from the customer during the delivery of the service is not a surprising result, since value is co-created [12], 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 [19]. 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 so on, from each customer, do impact the way by which the 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 [13] like crowdsourcing [14, 15] 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 [18], some of the tools proposed to do so might still be relevant.

Finally, our results suggest that while information management work is indeed collaborative [2, 3, 4, 20, 21], 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 [22, 23] 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. doi: 10.1145/1477973.1477976.

[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. doi: 10.1145/1866739.1866755.

[4] Botta, B., Muldner, K., Hawkey, K., Beznosov, K.: Toward understanding distributed cognition in IT security management: the role of cues and norms. Cognition, Technology & Work 13(2): 121-134 (2011)

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

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

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

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

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

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

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

[12] 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.

[13] 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. Available at SSRN: http://ssrn.com/abstract=1381502.

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

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

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

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

[18] 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.

[19] 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* (CHiMiT '09). ACM, New York, NY, USA, 2009.

[20] 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.

[21] 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.

[22] 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 (CSCW'06), Banff, Alberta, Canada.

[23] 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 (CHI'11), Vancouver, Canada

[24] Creswell, J. W.. Research Design: Qualitative, Quantitative, and Mixed Methods Approaches. Sage Publications Inc., 2003..