

Behind the Help Desk: Evolution of a Knowledge Management System in a Large Organization

Christine A. Halverson
IBM Research
650 Harry Rd
San Jose, CA. 95120, USA
krys@us.ibm.com

Thomas Erickson
IBM Research
19 Skyline Dr
Hawthorne NY, 10532, USA
snowfall@acm.org

Mark S. Ackerman
Univ. of Michigan, Schools of
Information Science and EECS
Ann Arbor MI 48109, USA
ackerm@umich.edu

ABSTRACT

This paper examines the way in which a knowledge management system (KMS)—by which we mean the people, processes and software—came into being and evolved in response to a variety of shifting social, technical and organizational pressures. We draw upon data from a two year ethnographic study of a sophisticated help desk to trace the KMS from its initial conception as a “Common Problems” database for help desk personnel, to its current instantiation as a set of Frequently Asked Questions published on an intranet for help desk clients. We note how shifts in management, organizational structure, incentives, software technologies, and other factors affected the development of the system. This study sheds light on some of the difficulties that accompany the implementation of CSCW systems, and provides an analysis of how such systems are often designed by bricolage.

Categories and Subject Descriptors

H.5.3 [Group and organizational interfaces]: Computer supported cooperative work

General Terms

Management, Design, Human Factors

Keywords

Frequently asked questions, design approaches, ethnography, distributed cognition, help desk

1. INTRODUCTION

The topic of knowledge management, in general, and help desks, in particular, has attracted considerable attention over the last decade. A particular concern has been with how people find and share the knowledge they need to carry out their work. In this paper we take a different stance. Rather than looking at how people *find* the knowledge they need, instead we examine how people *go about making knowledge available*. We look at a help desk, but rather than looking at clients, we look at the other side of the help desk: those who provide help. In this paper, we examine the processes through which the help desk codifies its knowledge so that it may be reused.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CSCW'04, Nov 6-10, 2004, Chicago, Illinois, USA.

Copyright 2004 ACM 1-58113-810-5/04/0011...\$5.00.

This study is of interest for two reasons. First, as others have found, capturing knowledge for reuse is not easy. Even setting aside difficulties having to do with what to capture, there are many problems associated with how to structure the captured knowledge so it can be used by others. In the case we examine, the help desk has developed a process for producing FAQs (Frequently Asked Question lists). Although this is a well-known genre for encapsulating and disseminating knowledge, there is a surprising dearth of empirical studies of the production or use of FAQs: our study appears to be the first.

Second, and more importantly, over the course of our fieldwork, we were able to observe the creation and evolution of the FAQ production process. Our observations raise a number of issues about how knowledge management processes are “designed” in organization contexts. While, at first glance, the development of this process seems a classic example of bad design, we argue that the evolution of the FAQ production process is, in fact, a highly logical response to shifting constraints in a dynamic organizational environment. We suggest that, far from being the exception, many knowledge management processes within organizations evolve as this one did. Thus, understanding how organizational processes come into being in such contexts is an important issue for CSCW.

In this paper, we proceed as follows. In section 2 we summarize previous work on knowledge management in general, and FAQs in particular. Section 3 briefly describes our approach: two years of ethnographically based fieldwork. Section 4 describes the field site, its organizational context, the structure of the help desk, and the processes by which it provides help and generates new FAQs. Section 5 focuses on how the FAQ production process came into being, tracing its evolution in response to a variety of complex, changing factors. In section 6 we discuss our findings, drawing on concepts from distributed cognition to make sense of the seemingly chaotic evolution of the FAQ production process.

2. BACKGROUND AND RELATED WORK

FAQs, in common with the home page, the search portal, and, most recently, the Blog, rank as one of the most common digital genres found on the Internet. They are also among the oldest, the first online FAQ being created on the ARPANet in 1982 (according to [21]). The FAQ list creation myth is that they arose in newsgroups as a response to repeated questions: news group regulars, irked by newcomers repeatedly posting questions which had been previously answered, compiled common questions and their answers into a list. The list was archived, and generally circulated to the news group on a regular basis, as an inoculant to newcomers’ questions.

We are most interested in FAQs as a form of knowledge management that focuses on collecting and packaging knowledge in a well-understood form that makes it readily available to its audience. Remarkably, there appears to be no empirical work on how FAQs are constructed by their maintainers, or used by their audiences. In spite of the seeming simplicity of FAQs as depicted in the creation myth, work on knowledge management provides grounds to suspect that the reality is not quite so simple. As many have pointed out, the use of knowledge management systems is not straightforward [4, 18, 20, 22], but is instead embedded in a complex weave of other resources, organizational norms, and practice, of the sort exemplified in Orr's work on Xerox repair technicians [16]. Researchers have also noted that packaging knowledge for dissemination—whether as an FAQ or in some other form—is a non-trivial process. For example, Randall et al. [18] challenge the oversimplified notion of a repository—a place where prepackaged solutions are made available.

A number of studies offer a CSCW perspective on KM in general. Schmidt and Bannon's [20] introduction of the notion of a Common Information Space (CIS) indicates how important organizational context is to knowledge. Others have detailed the specifics of that contextual importance. Orlikowski [13] shows how existing work culture acts as a disincentive for the adoption and use of Lotus Notes as essentially a knowledge management system. Ackerman and colleagues [1, 4, 2], have pointed to the need to understand the nuanced performances that mark knowledge work "in the wild," and have begun to apply that in system building. At a micro level Ackerman and Halverson [2, 3] point out that structuring knowledge for reuse requires first a decontextualization of the information, and then a subsequent recontextualization when it is applied to a new situation. Randall et al. [17] move up a level to categorize the kinds of things organizations need to remember, while [18] situates the knowledge management problem within the bounds of both legacy systems and daily work practice. These more situated approaches come together in a help-like setting, as Yamauchi et al [24] have shown in the example of Eureka—a system based on a deep understanding of how service technicians do their work, based on observations and insights starting with Orr [15] and developed by Whalen and colleagues [5, 24].

As well, there are several studies that examine the use of help desks and FAQs that are relevant. There are several studies of help desks that mainly look at them in terms of organizational or technical aspects (e.g. see [14]). The sparse literature that explicitly looks *behind* the help desk focuses on university settings—exploring the how and why of staffing, training and other issues. (For example, see proceedings of ACM SIGUCCS Conference on User Services.) Rarely is the issue of knowledge capture addressed, other than to point out that call records are extremely terse and unlikely to provide the fodder for automated systems to generate FAQs. For example, Takano [23] suggests a system using case based reasoning to record specific procedures as they occur and then save their traces. While this works for walking people through specific applications, such as the Adobe application they discuss, it does not address other less hands-on kinds of problems. Saunders [19] discusses FAQs that are automatically generated from the past 20 days of data, but is not specific about the content of those FAQs, or how (or whether) they are used. In contrast, Ng'ambi [12, 11] postulated that having

an FAQ didn't mean that it was a frequently referred to question, and thus of less value than one might expect.

In summary, while research has begun to explore some of the individual, social and organizational issues that make knowledge management a complex activity, there is still much to be learned. In this study, we examine the creation of an FAQ process as a way of managing and reusing knowledge in a large organization. We look both at how the current social and technical aspects of the system function, and also at how they evolve in response to organizational changes.

3. METHOD

The primary method for data collection was observation. The first author spent approximately 100 hours over a two year period directly observing and video taping across a number of roles and sites. The evolution of the FAQ production process began six months into the study, thus providing us with an 18 month window on it. Observations were supplemented with informal and semi-structured interviews conducted by phone and instant messaging. Data in all forms were transcribed, and reviewed against the original media for accuracy. We also had access to several of the corporate resources used to produce the FAQs. This included databases commonly available within the company, as well as the more restricted, group-specific databases. These were analyzed as described later in the paper.

4. FIELD SITE: THE TE HELP DESK

Our field site was a small part of the services arm of Gamma Corporation¹, specifically a sophisticated help desk that we will refer to as the Technical Engagement Help Desk, or TE Help. TE Help is not staffed by regular help desk personnel. As a consequence, TE Help differs from traditional help desks in a number of ways. Its staff (at least beyond the first level) has broad and deep expertise, and it must deal with complex, one-of-a-kind problems. Solving a consultant's problem can make the difference between the success or failure of a client engagement, and therefore finding a solution has an extremely high payoff. In this section we describe the organizational context in which TE Help is situated, and the processes and roles involved in providing technical engagement support.

4.1 The Context: Gamma Corp and NBS

Gamma is a large computer software and hardware company with a well established services arm supplying technical expertise. Our focus has been on a subset of these services, referred to here as Novel Business Services (NBS). NBS's aim is to create projects that support new business approaches, and it is an important one because it creates opportunities to sell Gamma's other products. One consequence of this is that Gamma's software often needs to be customized in order to link to legacy and third party software, and it is the consultants² who support this. These consultants, in turn, need support, and thus we come to TE Help's mission. TE Help was initiated by the training arm of NBS (N_Train), which has as part of its mission the education and mentoring of

¹ All names have been altered to protect the identity of companies and individuals.

² These are business consultants, ranging from those who propose an engagement to those who fulfill it.

consultants (something that further differentiates TE Help from traditional help desks). Furthermore, NBS is organized into a number of professional practices centered on particular areas. Our fieldwork focused on one of seven practices supported by TE Help: Web Business Application Development (WBAD), a practice that helps customers build web-based enterprises.

4.2 Structure and Organization of TE Help

Although TE Help only supports seven practices, they still get enough help requests that someone has to filter and distribute them. This filtering service is staffed by a more traditional, co-located group of first line (Level 1) help desk personnel from Gamma. While they are dedicated to TE Help, and have extensive experience on help desks, they have neither the technical nor the consulting expertise to provide substantive assistance to those who call TE Help. Thus, most help requests are passed on to the highly experienced experts who staff Level 2 of TE Help.

To understand the basic process of TE Help, we start with the consultant's entrée into the system (see figure 1). A consultant with a technical problem contacts the TE Help desk by calling a toll free number (1). This is routed to the Level 1 call center (2) already described. If possible, the Level 1 staff answers the question; however, it is more likely that they classify the call, identify the Level 2 practice under which it falls (e.g. WBAD), and then assign it to the portion of TE Help that supports that practice. The call is then referred to the team lead for the identified practice (3). The team lead does her own filtering (4) and then assigns it to a Subject Matter Expert (SME) within her team (5). The SME (6) initiates a dialog with the consultant (1) and begins to solve the problem. It is important to note that, in contrast to Level 1 help, almost none of the Level 2 staff of TE Help is co-located: most work out of their homes and communicate by email, instant messaging, and telephone.

4.3 How TE Help Processes Problems

Most of us are familiar with first line help desks where questions are straightforward and frequently repeated, such as a password reset. The problems TE Help gets are much more complex. They tend to be integration problems involving multiple products — often Gamma's own products combined with third party products. To solve such problems the SME draws on resources both inside and outside Gamma. (For details see

The SME compiles and packages the information, contextualizing it for the consultant. However, because one of the mandates of TE Help is to provide education and mentoring to consultants, the SME may also provide a broader answer than is strictly necessary. During the process of discovering and constructing the answer the SME records progress by saving a work history in a call tracking system and database used by TE Help as a whole.

At the beginning of our fieldwork, the processes by which TE Help provided support were those described above. However, six months after we began observing TE Help, we witnessed an effort to create a new process that would enable TE Help to package its solutions — as documented in the call tracking system and databases — in a reusable form. The next section describes the FAQ production process that resulted, as it existed eighteen months later at the close of our fieldwork.

4.4 The FAQ Process

An FAQ process was created to collect and disseminate common answers. The complex nature of the problems and solutions that TE Help deals with however means that this is not a simple matter of noticing a recurring question, and codifying its answer. Instead, producing FAQs requires much deliberation and consultation on the part of the already busy staff of TE Help, and thus requires a well defined process.

By the end of our study, TE Help's FAQ production process worked as follows: When the consultant agrees that their problem is solved, the SME "closes" the problem in the TE Help Database. If the SME believes that the problem has a common element that makes it a potential FAQ, he indicates that with a specific code, "Closed as FAQ," along with additional codes indicating the root cause of the problem. In addition, the problem and its solution are summarized in a Q and A format on the status line in the call tracking software.

Closing a problem as a potential FAQ triggers a partially automated process (see Figure 2) whose ultimate result is the display of the finished FAQs on a web page on Gamma's intranet. After a problem is closed as an FAQ (1), the problem is marked as a draft and transferred to a database application (2) that will eventually publish it to the web. Alice, a member of TE Help's administrative staff who oversees the FAQ production process, reviews the problem draft, and produces a draft FAQ based on the

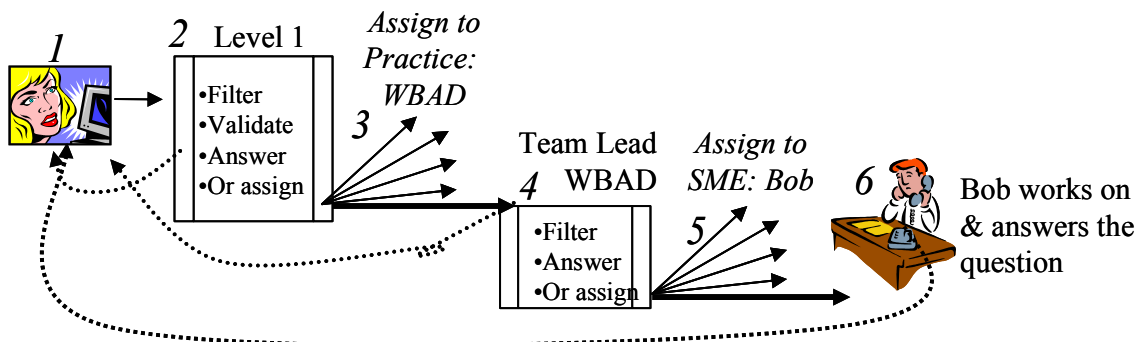


Figure 1. Process of a call to TE Help. The consultant (1) calls and is answered by Level 1 (2). They assign it (3) to a practice and a Team Lead (4) at Level 2. The Team Lead assigns the problem (5) to a Subject Matter Expert (6) who talks to the consultant (1) and works to solve the problem

summary in the problem’s status line, and the more detailed work history included in the problem record (3). Alice then notifies the SME that a draft of the FAQ is ready for their review (4). The SME may accept it, edit it, or redraft it completely (5). Alice does a final review of the draft FAQ (6), and if she is satisfied she marks it as “public” in the database. As the database application is connected to the intranet, marking the FAQ as public effectively publishes it, making it available for others within Gamma to find.

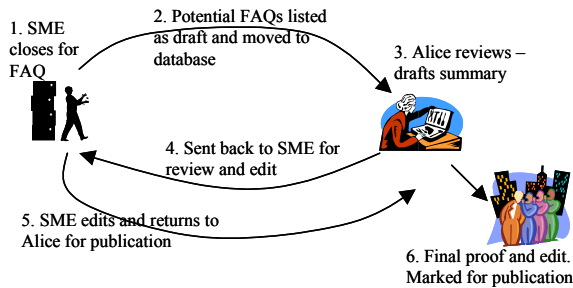


Figure 2. Outline of FAQ creation process.

4.5 The Resulting FAQs

TE Help considers their FAQ production to be successful: they have produced 477 FAQs in the 18 months that FAQ creation has been a goal, which averages to 26 FAQs made public each month (however, the actual rate of production is less consistent). Since there are no reports on FAQ production elsewhere, it is difficult to say much about the relative efficiency of TE Help’s FAQ production process.

To better understand the FAQs that TE Help produced, we coded the database across the dimensions in Table 1; the codes were based on the content of the FAQs, not on the categories used by TE Help for its internal metrics. The first author read through all of the FAQs. Coding categories were first derived from the field work. A pilot sample of the FAQs themselves required additional differentiations, which were derived from the content. Two more passes were made and the second author confirmed the categories.

Table 1. Coding Categories used for FAQ analysis

Application	Component	Process	How to	Who
Gamma Products	Operating System	upgrade	Basic Info (what is or how do I do)	People
3 rd Party Products	Programming Language	migration	Reference (where is)	
	Hardware	file conversion	Best Practice	
		installation		

Although our informants repeatedly told us that problems were complex and thus that FAQs were complex, we found that complexity did not translate into length. About 60% of the FAQs were less than ½ page; 20% were no longer than a single page; and the remaining 20% spread in decreasing frequency from one and a half pages to a maximum of 8 pages.

Instead, the complexity referred to by our informants comes out in the breadth of coverage in the content. That content, as illustrated in the coding categories in Table 1, is variable. The majority of FAQs fell into the *Application* category, but applications included 23 different *Gamma products*, and 27 *Third-party Products*. In

addition, a technical problem with an application could be complicated by different *Components* and/or *Processes*. For example, one FAQ dealt with the *Process* of migrating *Gamma product X*, which was *integrated* (another process) with *Third-party Product Y* from a Windows to a Sun platform (*hardware and operating system Components*). This FAQ is an example that crosses most of the categories on the left hand side of Table 1. Interaction between *Applications* and *Components* had a diversity of variations, covering 5 basic operating systems, 7 languages, and multiple hardware platforms (the left 3 columns of table 1).

Besides these technically focused FAQs, others are more focused on “how to”. This includes basic information about what something is (as in “has anyone heard of third party product Z”) or how to carry out a specific task, as well as where such information is located on the intranet and defining best practices. Finally a small subset of FAQs relate to locating people expertise with respect to a kind of problem

Taken as a whole, these FAQs illustrate the content that TE Help believes is useful to remember and reuse. While producing FAQs is not their primary job, nevertheless members of TE Help see it as a valuable activity in the context of their NBS mandate to educate and mentor consultants.

5. EVOLUTION OF THE FAQ PROCESS

Although the FAQ production process appears to be a clear and straightforward process, it is not the case that TE Help began, one day, to unproblematically execute the FAQ production process. It is not even the case that the process was really, in the normal sense of the word, designed ahead of time.

At the beginning of our fieldwork, the FAQ process did not exist. Furthermore, when TE Help began to make an effort to support the reuse of its knowledge, the initial process envisioned looked nothing like the final FAQ process: FAQs were not published to the web; problems were not closed as FAQs; nor, for that matter, was TE Help even thinking in terms of FAQs. There were a variety of other differences as well: different software was being used, a different manager was in charge, and, organizationally, the TE Help desk was in a different hierarchy. Over the course of two years, nearly every significant feature of what became the FAQ production process changed. The current form of FAQs and the FAQ production process have much to do with the changing needs surrounding the definitions of FAQs and the at times competing expectations of the various roles involved with their production. To show this, we describe the changing process and then go on to discuss the issues it raises.

5.1 How the FAQ Process Came into Being

To begin with, we need to re-emphasize the roots of TE Help. Although structured as a help desk, from the beginning its mission was viewed predominantly as helping to mentor and educate consultants – particularly those who were often by themselves at client sites. Its original funding and sponsorship came from the training arm of NBS (N_Train). One might argue that it should never have been cast in the mold of a traditional help desk. However, its first manager, Tom White, had experience administering help desks and formed it in that context. Thus, as TE Help was being set up, it drew on Gamma’s experience running help desks, and chose software necessary to support that endeavor—Tivoli Service Desk™ (TSD).

5.1.1 Alice and the “Common Problems” Feature

The original impetus for the FAQ process came, most directly, from Alice Liddle. Alice originally worked for the company that developed TSD. She was an expert in the software itself, training people to adapt and use the software for their particular needs, and in how the software’s knowledge base could help them over time. She identifies herself as a knowledge engineer. When Gamma hired her, she was still much more familiar with the capabilities of the software than with the needs of TE Help.

From her prior experience with other clients of TSD, it was evident to her that in a help desk there were core problems and issues that were constantly repeated. An initial connection she made was with the notion of Common Problems, a feature that TSD provided specific support for capturing in the software. Help desk software often has “boiler plate” associated with particular kinds of questions to speed up the process of creating a call record. For example, if a repeated question is how to reset a password, the Common Problems feature provides a way to save some text that can be put into the call tracking record in order to facilitate the speedy wrap-up of the call. As such, it also includes information about steps taken to solve the problem. Alice decided that this feature could be used to support TE Help. However she had a hard time convincing the SMEs that there was enough detailed repetition to fit this model. In addition she needed to lobby management that this was an important goal to spend time and resources on.

5.1.2 Common Problems become FAQs

Alice convinced Tom White of the importance developing the Common Problems aspect of TSD. Now she needed to work with the SMEs to identify the common problems, and develop the necessary text to elaborate them. She began to participate in the weekly operations calls in order to determine what the common problems were. As time went on this got transformed in the minds of the SMEs and Team Leads as FAQs. This is not surprising given that these are all highly technical people who would likely associate problem repetition with FAQs. However, this created a subtle shift from something that would help the help desk consultant do her job by providing common boilerplate, to something that potentially had a much broader reach.

As Alice Liddle pursued discovering the common problems she also began to develop a process to identify, create, and edit them and get them installed in the software. However there were two difficulties. Technically, the TSD software was unable to include attachments—a problem when answers require chunks of code. Second the SMEs were having difficulty understanding how to use the Common Problem feature within TSD. This came to a head when Tom White left with a sudden illness and his technical lead, Cooley Klein, took over.

5.1.3 Technical Lead Becomes Manager

Cooley had over 15 years with the company, and unlike Tom, had spent many years in NBS as a consultant. He had been the technical lead for TE Help, a natural outcome of his experience. When he took over from Tom White it was not surprising that one of his foci became how well TE Help’s software supported its activities. The software came up short in supporting FAQs.

Cooley’s experience as a consultant meant that his goals for FAQs differed from Alice’s initial effort associated with the Common Problems feature. Cooley wanted to expand TE Help’s

processes more in line with the education and mentoring charter that they had inherited from the training arm of NBS. As a consequence, he saw FAQs as a resource for consultants outside of TE Help. Coupled with his agenda for being more proactive at helping consultants solve problems, the shift in goals makes sense.

5.1.4 Supplementing TSD with a Notes Database

Cooley also initiated a shift in the software used for FAQs. While calls were still tracked in TSD, Cooley encouraged the use of a customized Lotus Notes™ database that he had on hand and was able to repurpose for FAQs.

Having a Notes database introduced a new set of constraints and opportunities. On the one hand, the underlying server infrastructure meant that the Notes database could be used both to work through the process and to publish to the web, making the FAQs more available. It also limited options because of the structure of the GUI. The Notes database was adapted from one customized for another use. Some of its inherited features had no value for the FAQ process, while others were exactly what were needed. Over all, the addition of the Notes database gave Alice and the SMEs something better to work with than TSD.

5.1.5 Linking TSD to Notes with “Closed for FAQ”

In parallel with the change in manager and applications, a number of process and incentive issues played out as well. When Alice Liddle first received the go-ahead to pursue Common Problems, she had to identify what they were. She began by sitting in on the operations calls, and requested the information from the SMEs. However, while Alice was initially focused on Common Problems, and the associated boilerplate to make the production of TSD call records easier and more complete, the SMEs had different goals. The SMEs wanted to provide aid to other consultants that would be longer term. As one told us, “I thought our job was to do ourselves out of a job.”

Over time the SMEs began to be more proactive in their FAQ identification, and began to develop team-specific processes. One team, upon deciding that a solution was a good candidate for an FAQ, began to summarize the one line description of the problem in a Q(uestion)/A(nswer) format. Another team met together each week in order to discuss potential FAQs. At the operations meeting they then presented Alice with a list, which might or might not be discussed further.

Both of these informal solutions ultimately led to changes in the official process and the applications. The Q/A format in the summary really helped to identify which problems were candidates. However, it still required that Alice browse the database to find them. This led to a request to add a new closed code in TSD, “Closed for FAQ”. Having a separate code meant that a report could be run, highlighting just those problems.

5.1.6 Fine Tuning with Manual Processes

There were some additional programming issues related to the incompatibility of the file formats. While the “Closed for FAQ” code allowed a report to be produced by TSD, there was no easy way to move the data from that report into the Notes database. This resulted in a number of manual processes to help Alice populate the FAQ database with the pertinent information.

The Q/A format also helped Alice summarize the more complicated problems without having to read through the entire

work history. However, there remained two problems. First, not every group used the Q/A format in the status line when they closed a problem. This led to a process change making it mandatory to do so. Second, not every problem could be successfully reduced to the Q/A format. However, by including a note on the status line to look for additional information in the work history, the SMEs were able to save Alice a lot of work.

Both of these improvements began to further alter the process to be more of a push from the SMEs to Alice, rather than her pulling from the SMEs, and bring us within sight of the final FAQ production process already described.

5.2 Organizational Pressures

It is time to revisit the FAQ production process outlined in Figure 2 from another perspective. We have explained how the process evolved as technology and management changed, and as new process components were invented and adopted. And we have done so primarily from the vantage point of TE Help. Now we are going to shift perspectives and look at the external pressures shaping the evolution of the process. Stepping outside of TE Help we look at how Gamma perceives help desks and how, within a context of changing organizational needs, the perception of an FAQ and its value changes.

5.2.1 *Costs, Values and Metrics*

At an organizational level, Gamma Corp. is answerable to their stockholders for profits. As a consequence, Gamma is interested in trying to assess and control costs relative to the values of what it produces. Commonly, help desks are a cost controlled service that is sold to clients—whether inside or outside the organization. Organizational control is often through measurements using standard metrics such as time spent per call, wait time per call, call volume, and so on. These measurements are converted to costs, subsequently used to evaluate the return on investment, and ultimately used to adjust costs or prices as desired.

Although TE Help is structured as a help desk, it was not originally envisioned as such. First, its services are not sold to clients. Instead it provides a service to Gamma's consultants that neither they nor their outside customers have to pay for. Second, the problems it deals with are exceptionally complex, and not amenable to quick answers. Third, it is primarily staffed (except for the first level) by SMEs selected for their expertise and experience, and thus has higher staff costs than most help desks.

The creation of TE Help by N_Train came about because of their goal of educating and supporting consultants by providing them with contextualized knowledge, rather than just answering questions. This goal of mentoring and training consultants isolated in the field justified the higher cost in time and staff. It also required different measurements to ascertain value and aid management and control. Reducing call durations is more cost effective in terms of labor; it is useless if the consultant's question is either unanswered or not understood. Thus, the unusual nature of TE Help means that the process of measuring its costs and outputs diverges from that used for traditional help desks.

TE Help collects a number of metrics to provide evidence of its value in relation to its special charter. These metrics focus on assessing the utility of the answer to the customer, identifying areas where N_Train could be more proactive, and where Gamma had software problems that needed to be addressed. One example is that TE Help randomly surveys its consultant customers to

obtain their evaluations of the success of the answers. Another example is that problem causes are collected and analyzed, as are the reasons for closure. These closed codes focus more on the ultimate user of the information³ than metrics normally used for a traditional help desk. In terms of FAQs, TE Help also tracks the overall numbers produced, the number produced per practice, how many web hits each FAQ receives, and so on.

At the same time, because TE Help uses a call tracking system (TSD) that is designed to support call centers, it produces a number of work traces that the organization can use to evaluate TE Help's performance along more traditional lines. The items tracked include the number and kinds of problems, how long the problem was open, and so on. Furthermore, since each step taken towards the answer has been saved in a work history in TSD, it provides the record that is used to generate FAQs. A side effect of the TSD based metrics is that they help both management and the SMEs make sure that they do not lose track of a problem.

Taken together, these metrics — both the TE Help specific ones and the traditional metrics provided by the TSD system — allow NBS to assess TE Help's performance in terms that are more appropriate to its charter than the standard help desk metric that the length (and therefore the cost of the call) is equal to its value. The bottom line is that, despite very long call lengths, and apparently small numbers of FAQs, TE Help itself, and the FAQ process, were both deemed successful, until the overarching organizational mandate changed.

5.2.2 *Organizational Shifts, Interpretive Shifts*

Toward the end of our study, as part of a corporate reorganization, several groups in NBS, including TE Help, were shifted to another area within the services arm called Managed Accounts (MA). Despite the other agendas driving TE Help, the fact that it was characterized as a help desk probably made it an obvious candidate for this transition. MA is a part of Gamma that primarily sells packaged services, like help desks and network operations centers, to Gamma's clients.

Unfortunately, this shift introduced a number of new expectations. MA is not chartered with mentoring consultants; instead, its experience is with traditional help desks, and thus it is much more concerned with managing the cost of providing a service. Under MA, the codes that TE Help used to signify that a problem had been resolved became a metric for the cost of answering a call.⁴ That is, while the metrics TE Help collected didn't change, the interpretive frame through which they were seen did. For example, the relatively long call times were not only interpreted as a high cost per call, but also a more valuable metric than the client's estimation of time saved. The mentoring role of TE Help was not part of MA's interpretive frame.

The shift to MA and the focus on managing costs also subtly shifted the incentives for SMEs to produce FAQs. Previously the

³ For example, a lack of documentation, or a bug, are things that parts of the product team care about. However, they are not things that TE Help or NBS have control over.

⁴ This happened around the time that the "Closed for FAQ" code was being instituted, and in fact that change—because it had the potential to adversely affect the metrics —was instituted with a great deal of concern.

SMEs' identity as a part of the larger consultant community meant they wanted to mentor less experienced consultants. This fueled their willingness to put together somewhat amorphous recurring core issues into an understandable package to be used in a broad range of contexts. As time pressures increased, however, this incentive got replaced with the need to focus on what kept them "in business"; that is, answering questions. If the SMEs' value was based on the number of questions they answered, and therefore how fast they could answer them, there was little or no incentive to spend the extra time to create FAQs.

6. DISCUSSION

The organizational environment of TE Help is complex and dynamic. Personnel, processes and metrics changed as the world changed. Of course, the importance of context and the continuously changing nature of real world practice have long been known. We want to move beyond just "recognizing the significant influence of these organizational elements [13]." We want to understand how the system evolves over time in response to its changing environment, and how as designers and builders of technology we might take advantage of this knowledge. We believe this to play out in three main areas discussed below.

6.1 Design and Adaptation

At first glance we face a common conundrum. Knowing how things are now — the FAQ process as shown in figure 2 — it seems obvious that that is how it *ought* to be. Surely it was planned that way from the start; if not, surely that is simply an indication of 'poor design.' We hope that our description of the evolution of the FAQ process indicates that neither is the case. While there was no overarching plan that produced the FAQ process, neither was its evolution entirely chaotic; rather, as the series of actions that lead to the FAQ process unfolded, each responded to particulars of the situation at that point in time. In this section, our aim is to examine this more closely.

Suchman's analysis of the conduct of work in part focused on how "plans are one of a range of resources which guide the moment-by-moment sequential organization of activity, rather than laying out a sequence of work which is then blindly interpreted." [6] Hutchins [8], goes beyond this to suggest that we can delve deeper and understand the specific changes taking place in an organization. That is, by looking at the organization as a cognitive system it is possible to analyze where and how changes are taking place. Paralleling Suchman, he says:

Common sense suggests that work is organized in accordance with plans that are created by designers who reflect on the work setting and manipulate representations of the work process in order to determine new and efficient organizational structures. Or, even if "outside" designers are not involved the reorganization of work is normally attributed to conscious reflection by members of the work group itself. A detailed examination of the response of a real-world group to a sudden and unexpected change in its informational environment shows that these common sense assumptions may be quite misleading. [8]

In *Cognition in the Wild* [9], Hutchins' micro-analysis of a navigation team's computations during a crisis shows how a solution evolves for the team. While it may appear to be the result of conscious reflection, is also due to adaptation to local task demands. "While the participants may have represented and thus

learned the solution after it came into being, the solution was clearly discovered by the organization itself before it was discovered by any of the participants."

Hutchins identifies two key ways in which organizational change happens: *adaptation by design* and *adaptation to local design*. Adaptation by design implies a conscious reflection on the whole system ("system" including both the social and the technological) — we will refer to this as *reflective design*. Adaptation to local design is much more of an immediate reaction to local constraints — we will refer to this as *local adaptation*.

6.1.1 Reflective Design and Local Adaptation

Some aspects of TE Help clearly come about through reflective design. In the path from vision to implementation the background of who is driving the reflective design has consequences. For example, N_Train was driven by a concern for the newly minted consultant in the field. Their vision centered on TE Help as a mentoring operation. However, TE Help's organization as a help desk resulted from Tom White's background and familiarity with them. His reflection on the system was cast in the light of his experience and, he formed it into something that resembled a help desk, even though N_Train's organizational goals didn't match a traditional help desk. Over time the mission of TE Help changed as being in the form of a help desk came to influence its function.

This is just one example of how individuals — despite their insignificance relative to the size of a corporation — can have great impact. That impact is not limited to reflective design however, as the evolution of TE Help shows a considerable amount of local adaptation. For example Alice Liddle's proposal to implement TSD's Common Problems feature — which was partly responsible for initiating the FAQ process — was started as reflective design based on her previous experience. However, once begun, local adaptation came into play in a number of ways. Experiences and expectations of the SMEs changed the common problems concept into something called FAQs. Cooley's reuse of an existing database to support the process added the shaping of software constraints. Finally, the move to MA further altered somewhat loose collections of information into tighter FAQ packages, with a process supported by the software.

6.2 Integration: Chaos or Bricolage?

Of course the reality of system design is not as simple as reflective design here and local adaptation there — rather, they are entwined. Our data suggest that the underlying mechanism is bricolage — putting together tools in a way to accomplish new needs, making unforeseen new use possible [10]. As we saw with repurposing the Notes database, reflection on the tools at hand and the needs of TE Help led to local adaptation of a customized database to handle the FAQ process. In fact, we clearly see two classes of bricolage in our data, depending on whether tools or processes predominate. We explore each of these below.

6.2.1 Software Bricolage- Tools predominate

When tools predominate we call it software bricolage — largely because the "tool" of choice tends to be software in some manner or form. We see this over and over in the FAQ process of TE Help—from the first adoption of TSD, reworking one of its features, to adoption of a Notes database and the concerns about software changes that might affect changing metrics.

In a technical setting such as Gamma, it might seem odd that system design is driven by bricolage. One would expect either the resources to build or purchase the needed software. However, organizational resources are often limited, development is expensive, and reuse is sensible. Reflection and experience help Cooley see that an existing database could be usefully reused rather than build or request a new one. However, reuse in a new situation is rarely a perfect fit, requiring additional resources for customization. These adaptations address the slight differences between the original and current situations and are more likely possible with limited funds. While they also take time, understanding that things will continue to change means adaptation will always be necessary and spending extra time and effort to create something new seems less reasonable.

From a personal point of view reuse makes sense as well. Given the choice between trying to learn a new tool that *might* do the job better and using an existing tool that will sort of do the job with some help, an individual often chooses the latter. This is also software bricolage because the subjects are constructing what they need from existing tools in which they have expertise. This is not a stupid choice. After all, learning a new tool is a big risk. What if it does not make things better but instead worse? It is often better to be a real expert in a tool known (by one's colleagues and organization) to be useful, than to have attained basic competence in using a tool no one has heard of.

6.2.2 Process Bricolage

The second form of bricolage is where process is predominant. We see a number of process changes throughout our observations, and our data suggest that the end result is a bricolage of both process implementations and their subsequent alteration. For example, adopting the Common Problems (CP) aspect of TSD was a reflective design choice that carried with it a specific process for identifying areas that needed repeated text in the call record. The shift from Common Problems (CP) to a vague form of FAQ was a local adaptation that carried its own process expectations based on online notions of FAQs. At the same time there was a shift from identifying repeating identical problems, for which the boilerplate of Common Problems could be used, to recognizing that their mentoring mission was better served by packaging a more general collection of problems that focused on a common issue as an FAQ. These two alterations shifted the process of creating the FAQ from one of unique subject lines to one where the problem solution was identified as an FAQ in the subject line by the Q/A formatting.

Once again, process bricolage makes sense for much the same reasons that software bricolage does. In addition, we have all had the experience of creating our own ways to work around a recalcitrant process. Over time some of these workarounds, such as the Q/A formatting, get adopted into the process itself. Thus, both individual and organizational perspectives reinforce the use of bricolage over new design. While most corporate situations rarely have the kind of crisis that sparks a rapid redesign (as Hutchins shows), we still observe a similar effect through local adaptation. At the same time, deep knowledge of the more general organizational needs and issues may result in reflective design. We argue that such bricolage is normal, even though it can appear chaotic. The case of TE Help may seem meandering and even dysfunctional, but the actions of the principals and their adaptations to local changes are actually quite operative.

6.3 Incentives and Organizational Change

In the previous two sections we have distinguished between reflective design and local adaptation and discussed how they were entwined in bricolage (software and process) over the course of the evolution of the FAQ process. In this section we turn to the issue of incentives, one of the mechanisms through which higher levels of an organization attempt to control (via reflective design) its lower levels. Not surprisingly changes in incentives produce responses from the lower levels of the organization that reflect both reflective design (as intended by those changing the incentives) and local adaptation. We see this most strongly in the incentive shifts caused by the move of TE Help from NBS's N_Train division to Managed Accounts (MA) and the consequent effect on the evaluation of TE Help's success.

The use of incentives as a management tool requires the existence of metrics to measure how well individuals and groups are fulfilling their commitments to the organization. Of course what are appropriate incentives or metrics vary. From the consulting perspective, one good metric for Gamma's field consultants has to do with the number of hours they bill; it is presumed that being able to continue to get and bill work means they are good workers. In contrast, a help desk counts things at a fine grain that are related to expense, such as time to call completion, number of calls resolved, and call waiting time. In turn these are used as incentives at the individual level, for example when a help desk responder is enjoined to get their calls below 15 minutes. It is also the case that not all incentives are organizationally produced; some are more personal. Consultants may derive personal satisfaction from doing a good job, or solving hard problems. In the context of TE Help, the SMEs – all of whom were also consultants – took their mentoring charter seriously, and got satisfaction from the feeling that by mentoring field consultants they were giving something back to their community.

How did the incentives shift as a consequence of TE Help's organizational move, and how did that shift interact with both reflective design and local adaptation? At the beginning, N_Train designed and used TE Help to mentor consultants in the field. In view of this charter, efforts to mentor—whether reactive (answering requests) or proactive (writing FAQs)—were valued and supported. Thus the cost of more experienced personnel, longer call times, and more complex problems, were all justified by the mentoring mission. Evaluation of success was based on customer reports (collected through surveys) of time saved and amount learned. At the same time, the SMEs also derived personal satisfaction from supporting the consultant community of which they were part – it was this that provided an incentive for doing the extra work to produce the FAQs. Another personal incentive was that working for TE Help allowed them to opt out of the consultant rat race for a bit: by working for TE Help they were able to reduce their travel, and to ensure a majority of their billable hours each month. Objectively, these incentives, while in line with N_Train's mission and the SMEs personal predilections, are also in line with the ways SMEs are evaluated from the consulting perspective. Thus local and organizational pressures overlap nicely.

However, over time, the fact that TE Help was cast as a help desk had its own effect. The mandate for the group subtly moved from mentoring to problem solving. The consultant in the field is generating money, and every minute lost to a problem is potential

revenue lost: the more quickly a problem is solved, the better. But this new attitude, of course, is at odds with the mentoring perspective; it also is in tension with the fact that the problems TE Help deals with often have at least one novel element requiring longer solution times. While under a mentoring mission it was considered acceptable to keep a problem open for a month while waiting for new code to be installed, as a help desk striving for a timely solution this seemed worrisome. Similarly, a long wait to produce a complicated FAQ was acceptable when FAQs were seen as a tool for mentoring, but was less so if it could not be demonstrated that FAQs reduced call volumes. Here we begin to see a conflict between local and organizational pressures. The SMEs were still being driven by the more personal incentives related to being part of Global Corp's consulting community, but the administrative staff was becoming more concerned about the information that the help desk style metrics were exposing to upper management.

When TE Help officially became a part of Managed Accounts (MA) these shifts became more institutionalized. Metrics for call completion, while longer than the standard 15 minutes in first line help desks, were set and expected to be met. Incentives shifted to meeting metrics rather than supporting one's colleagues. The SMEs are now operating under conflicting pressures. The organization wants them to hurry up, while their colleagues and their own sense of competency insists they must use care to ensure an appropriate answer to a complex question while providing the basis for other answers (via the FAQ process) later.

The transfer of TE Help from NBS to MA shifted the incentives and changed the scope of what was deemed appropriate work. These changes were, in part, the result of reflective design by upper management. However, the needs of the customers, the consultants in the field, had not changed. Thus local adaptations continue to take place that are not necessarily in line with MA's metrics and goals. This is because the individuals of TE Help—aware of the discrepancy between MAs incentives and the needs of their customers—continue to perform the local adaptations that make the larger socio-technical system competent.

7. CONCLUSION

Most houses are not designed by architects. Most rooms are not arranged by interior decorators. Nor is most clothing cut to fit by tailors. It is not the case, however, that design or designers are entirely absent. Design often has its effect at a distance. It acts through a variety of indirect means such as components designed for general use, and rules for configuring those components – whether regulations with legal force or socially propagated styles). It is up to individuals, in the course of their day to day activities, to configure the materials of their situation so as to best suit their own needs. The canonical idea of design—the careful assessment of the needs of users situated in a particular material, social, and organizational context followed by the crafting of an artifact or process that is made to measure—is not the norm. However desirable, however effective it is at producing results, this canonical approach to design is a luxury that many—whether individuals, groups, or organizations—do not pursue.

As it is with traditional design, so it is with the design of technology and systems. As much as we advocate the careful analysis of situations, and the design of systems that support the observed practices of their users, it is often the case that customized design does not take place. Rather, the design of

knowledge management systems bears a strong resemblance to what architects call vernacular design, design in the absence—although not beyond the influence—of designers. Our study of TE Help is important because it illustrates in some detail how this sort of design—at least in the context of large organizations—occurs more often than not.

Previous work has examined building technical systems and looked at how organizational knowledge works around those systems and in the wild. In contrast, this study demonstrates how a knowledge management system is constructed within a particular context. The TE Help Desk FAQ system began as a Common Problems database for use by help desk personnel. In response to shifting social technical and organizational pressures, the FAQs evolved into something quite different. This pattern of change looks from the outside like chaos or bad practice. In fact, it is normal practice.

What appear to be breakdowns or poor practices are adaptive responses that are organizationally functional. Some of the changes are the result of people moving things forward, contingent upon the specific organizational situation, and using bricolage as a design strategy. Some changes, such as the change of the status line content on closure, are adopted after being proven useful. Other changes, such as the switch to using the customized Notes database for FAQs, are more opportunistically driven—based on previous experience or availability. Although an existing solution or tool may not be the best for the problem at hand, the fact that it is familiar and ready to hand make it a sensible tactic for rapidly achieving an end.

And, in fact, TE Help did achieve its end. It produced a functional system that took knowledge captured by the help desk processes, and enabled a spatially distributed group of staff to decontextualize the knowledge and package it as FAQs. The FAQs were then made available, not just to NBS consultants, but to Gamma's much larger base of employees. This is knowledge management writ large, indeed. While we do know (because TE Help tracks it) that there are substantial numbers of hits on the FAQs, we do not know how, and to what extent, the intended audience of the FAQs is actually using the knowledge contained therein. This is an area for further investigation, both for our case, and the case of FAQs in general. And, if we can draw any conclusions from other research, it is that the usage practices that surround FAQs are likely to be as complex, in their own ways, as the production practices.

That said, are there general lessons that CSCW can draw from this work? A cynical response might be that we should give up. If most knowledge management processes are designed by bricolage, one might conclude that most of the nuanced knowledge that CSCW has developed is irrelevant to the larger part of the world. We should be sorry to provoke such a response. Rather, we wish to emphasize that design is not absent, but rather that it is a step or so removed. The local objects and processes that are the subjects of bricolage are the means through which design exerts its influence. It is the design of an artifact that makes it easy, or difficult, to turn to another end, and thus either suitable or inappropriate for bricolage. Thus, if we take seriously the claim that much of the design of knowledge management processes occurs via bricolage, it behooves us to ask how artifacts and processes might be designed to prosper under these conditions.

8. ACKNOWLEDGMENTS

We thank anonymous reviewers for helpful comments, as well as all the members of TE Help and those in Gamma who have supported this work. This work was supported in part by NSF grant IRI-0124878.

9. REFERENCES

- [1] M. S. Ackerman. Augmenting the Organizational Memory: A Field Study of Answer Garden. In *Proceedings of the Conference on Computer-Supported Cooperative Work (CSCW'94)*, 1994, 243-252.
- [2] M. S. Ackerman and C. Halverson. Re-Examining Organizational Memory. *Communications of the ACM*, 43, 1 58-63.
- [3] M. S. Ackerman and C. A. Halverson. Considering an organization's memory. In *Proceedings of the Computer Supported Cooperative Work*. (Seattle, WA, ACM,1998, 39-48.
- [4] M. S. Ackerman and D. W. McDonald. Answer Garden 2: Merging Organizational Memory with Collective Help. In *Proceedings of the ACM Conference on Computer-Supported Cooperative Work (CSCW'96)*. 1996, 97-105.
- [5] D. G. Bobrow and J. Whalen. The Eureka Story: Community Knowledge Sharing in Practice. *Reflections*, 4, 2 407-431.
- [6] P. Dourish and G. Button. On "technomethodology": Foundational relationships between ethnomethodology and system design. *Human Computer Interaction*, 13, 4 395-432.
- [7] C. A. Halverson. The Value of Persistence: A Study of the Creation, Ordering and Use of Conversation Archives by a Knowledge Worker. In *Proceedings of the HICSS 37*. (Waikaloa, HI, IEEE,2004, electronic.
- [8] E. Hutchins. Organizing Work by Adaptation. *Organization Science*, 2, 1 14-39.
- [9] E. L. Hutchins. *Cognition in the Wild*. MIT Press, Cambridge, MA., 1995.
- [10] C. Lévi-Strauss. *The Savage Mind*. The University of Chicago Press, Chicago, 1962.
- [11] D. Ng'ambi. Dynamic "Intelligent Handler" of Frequently Asked Questions. In *Proceedings of the Intelligent User Interfaces - IUI*. (San Francisco, CA, 2002, 210-211.
- [12] D. Ng'ambi. Pre-empting User Questions through Anticipation - Data mining FAQ lists. In *Proceedings of the SAICSIT-Annual research conference of the South African institute of computer scientists and information technologists*. (Port Elizabeth, South Africa, ACM Press,2002, 101-109.
- [13] W. Orlikowski. Learning from Notes: Organizational issues in groupware implementation. In *Proceedings of the CSCW92*. 1992, 362-369.
- [14] W. Orlikowski and D. Hoffman. An Imporvisational Model for Change Management: The Case of Groupware Technologies. In (eds.), 1997.
- [15] J. Orr. Narratives at Work: Story Telling as Cooperative Diagnostic Activity. In *Proceedings of the Computer Supported Cooperative Work*. (Austin, TX., ACM,1986, 62-72.
- [16] J. Orr. *Talking about Machines: An Ethnography of a Modern Job*. Cornell University Press, Ithaca NY, 1996.
- [17] D. Randall, J. Hughes, J. O'Brien, M. Rouncefield and P. Tolmie. Memories are Made of this: Explicating organizational Memory. In *Proceedings of the KMAC, The Knowledge Management Conference*. (Aston, Birmingham, UK, Operational Research Society.,2000, 255-264.
- [18] D. Randall, T. Rodden, M. Rouncefield and I. Sommerville. Remembrance of Designs Past: Legacy Data, Organisational Memory and Distributed Design. In P. Henderson (eds.), *Systems Engineering for Business Process Change. Vol. II*, London, Springer-Verlag, 2002.
- [19] A. Saunders. The 1, 2, 3's of Helping Yourself -- Search for answers, not for documentation. In *Proceedings of the ACM SIGUCCS Conference on User Services*. (Bloomington, IN., ACM Press,1998, 215-218.
- [20] K. Schmidt and L. Bannon. Taking CSCW Seriously: Supporting Articulation Work. *Journal of Computer Supported Cooperative Work*, 1, 1-2 7-40.
- [21] R. Shaw. *The FAQ Manual of Style: Creating Great Frequently Asked Questions Lists*. MIS Press, New York, 1996.
- [22] L. Suchman. *Plans and situated actions: the problem of human-machine communication*. Cambridge University Press, Cambridge, 1987.
- [23] A. Takano, Y. Yurugi and A. Kanaegami. Procedure Based Help Desk System. In *Proceedings of the Intelligent User Interfaces - IUI*. (New Orleans, LA, ACM Press,2000, 264-271.
- [24] Y. Yamauchi, J. Whalen and D. G. Bobrow. Information use of service technicians in difficult cases. In *Proceedings of the CHI 03*. (Ft. Lauderdale FL, ACM Press,2003, 81-88