

Community Knowledge Sharing in Practice: The Eureka Story

**Daniel G. Bobrow
Jack Whalen**

**Systems and Practices Laboratory
Palo Alto Research Center**

3333 Coyote Hill Road
Palo Alto, CA 94304

bobrow@parc.com, jwhalen@parc.com

Abstract

This paper presents a detailed account of a knowledge-sharing system in practice. The Eureka system leverages the practical know-how and inventions of frontline employees, with the goal of helping the entire organization learn from their experience. This system and its enabling technology were developed through several iterations using a distinctive socio-technical methodology. Eureka is now deployed worldwide in a large corporation's service organization. We provide a précis of the methodology, and summarize the challenges faced in bringing such a change into the corporation.

This text is a prepublication version of a paper that is to appear and can be cited as:
Bobrow, Daniel G., Whalen, Jack, "Community Knowledge Sharing in Practice: The Eureka Story",
Reflections, Published by the Journal of the Society for Organizational Learning and MIT Press,
Volume 4 Issue 2, Winter 2002

Community Knowledge Sharing in Practice: The Eureka Story

Introduction

It is widely recognized that an organization's most valuable knowledge, its essential intellectual capital, is not limited to information contained in official document repositories and databases – scientific formulae, “hard” research data, computer code, codified procedures, financial figures, customer records, and the like – but also includes the largely undocumented ideas, insights, and know-how of its members (see, for example, Nonaka and Takeuchi 1995; Stewart 1997; Davenport and Prusak 1997; Senge et al 1999).

This more tacit or informal knowledge is deeply rooted in the experiences of individuals and the culture of their work communities. It commonly originates as practical solutions – through everyday inventions and discoveries – to the problems that must be solved to be successful at work, and thus serves as the critical resource for ordinary work practice (see especially Brown and Duguid 1991, 2000). Much of this knowledge often remains embodied in the practice. Crucial steps in a new practice and fresh solutions to recalcitrant problems are commonly shared through conversations and stories among small circles of colleagues and work groups, with members filling in the background and gaps from their own experience. These instructions and stories are naturally expressed in the local vernacular of those groups and communities.

It is also recognized that it is a significant challenge for organizations to somehow convert this valuable but mainly local knowledge into forms that other members of the organization can understand and, perhaps most important, act upon. This paper presents a detailed account of an effort in one organization at encouraging inventiveness and capturing new ideas generated by people at work, and at using technology to then share the best of this knowledge beyond the local work group.

This account is based on our experiences over seven years with the design, development, deployment and evaluation of the Eureka system in Xerox Corporation. Eureka is being used in Xerox to support the customer service engineers (CSEs) who repair the copiers and printers installed at our customer sites. Four iterations of the system took it from an experiment initiated by researchers at the Xerox Palo Alto Research Center (PARC) designed to measure the value of codified field experience, to a system now deployed to 20,000 CSEs worldwide. Through a set of glasses that focus on communities and how they share knowledge in ordinary practice, we developed a set of questions and a methodology that we hope will enable others to build similar community knowledge sharing systems. However, deploying any knowledge system equally involves pushing changes within a corporate culture; understanding the Eureka experience, as well as the problems facing all knowledge systems that have to be deployed in the real world, requires equal focus on these challenges.

Our analysis follows a narrative format, recounting the story of Eureka in Xerox. But it has a generalizable theme: Large organizations face certain knowledge problems when established doctrine – well-known policies, procedures, beliefs – prove inadequate or even erroneous as guides to effective action, and traditional organizational methods for generating new knowledge and managing its distribution won't suffice. In this kind of not terribly uncommon situation, supporting the creativity and inventiveness of those members of the organization closest to the problem, warranting their insights, and then sharing those new solutions widely, can make all the

difference between inspired success and abject failure. This is just the sort of inventiveness and sharing that Eureka was designed to enable.

The narrative covers the complete history of this project – carefully detailing the fundamental interrelationships between the social and the technical. It concludes with a précis of a framework for building these kinds of community systems and our reflections on the barriers to organizational change that their proponents confront.

Breaking the Frame

Helping the Xerox field service force

Xerox has over 20,000 technicians worldwide who help to ensure that the machines we sell to our customers are performing as expected by them. As Orr (1997) pointed out, this is a triangular relationship – among the technician, customer and the machine. On many service calls, the machine needs to be repaired and/or adjusted; on some, it is the customer who needs adjustment – in their expectations, procedures, and/or knowledge of the machine. In the early 1980's, with the decreased availability of technicians trained by the armed services to debug complex equipment, Xerox decided to make it easier to use less skilled and experienced service people. They did this by moving away from documentation and training that simply described the principles of operation of a product, leaving it to the technicians to determine the appropriate repairs. They moved toward what was called “directive” repair and adjustment procedures. These are documented instructions in the form of a decision tree. Each decision step was of the form “do the following setup and test; make the following measurement (or observation); if the result is A, do X; else do Y.” The intuition embodied in this form of documentation is that technicians need only be trained how to correctly use the documentation to be able to diagnose and repair any failure of the machine.

Our group at Xerox PARC has a background in Artificial Intelligence, and in modeling electromechanical systems. In particular, we had expertise in building programs that perform efficient diagnoses of machine faults given an abnormal symptom and the ability to obtain observations/measurements from the machine (de Kleer and Williams 1987). As a test of our technology, we decided to build a model of one complex module of a particular photocopier, and demonstrate how a program could provide guidance for a technician in diagnosis and repair of problems in this module. Our hypothesis was that if we were successful, the model-based expert system running on a laptop carried by technicians in the field could replace the documentation and support an optimized work process for isolating faults. In addition, this model-based expert systems approach could improve Xerox's speed in bringing supported products to market, because creating documentation was a gating factor in deployment, and models could be created in parallel with design; moreover, newer machines often used the same or similar subassemblies, making models reusable.

We succeeded in building RAPPER (Bell et al 1991), a model-based expert system that used a model of the recirculating document handler to provide optimized guidance in isolation of faults in that module. The model was constructed with the help of Xerox engineers and technicians; it captured the behavior of all the faults that could be found using the standard documentation. We showed it to some technicians. “That's amazing,” they said. “Would it really be useful if we had a complete model for the machine?” we asked proudly. “Not really – though it is amazing – rather like a bear dancing. It is surprising to see it do it at all.”

We probed further for the issues that underlay their negative response. First, small optimizations of the time required to isolate any known fault were not worth much. Only a relatively small portion of their average two-hour call was actually devoted to diagnosis. Second, they usually knew the procedures for the common faults, and so required no guidance. The thorniest issues, though, were much more surprising to us. For many products (those produced by our Japanese partners) there were no full descriptions of their operation. Additionally, the diagnostic documents were produced from experience in inserting faults into the machines in a laboratory, and then recording the symptoms. But the hardest problems in the field for all machines were not those covered by the documentation – they were new problems.

There are a number of reasons why new, unexpected problems occur in the field that are not predicted from laboratory tests. Xerographic machines work differently in different environments – extremes of temperature, humidity and dirt can all cause surprises. As machines age, components develop new failure modes; and when a new machine comes on the market, there may be new components with unexpected interactions. With vibrations and other slow acting processes, faults develop that are intermittent, and these are the hardest to track down. These difficult cases caused most of the long calls, and were the reason the average call time went up.

Accordingly, we decided to spend more time observing what technicians actually did in their day-to-day practice. We started with technicians in the United States, accompanying them on their service calls. Most of the time, they would approach a machine, talk to the customer, and know exactly what to do to make the machine available again in good working order. Once in awhile, they ran into a problem that they hadn't seen before and for which there was no answer in the documentation. They would try to solve these problems based on their knowledge of the machine. This often worked, but sometimes they were stuck. They might call on a buddy for ideas, using their radios, or turn to the experts – former technicians now serving as field engineers – who were part of the escalation process. When unusual problems were solved, they would often tell the stories about these successes at meetings with their co-workers. The stories, now in the possession of the community, could then be used in similar gatherings and further modified or elaborated (see Orr 1997, Brown and Duguid 1991).

This practice pointed to the importance of non-canonical knowledge generated and shared within the service community. It suggested to us that we could stand the Artificial Intelligence approach on its head, so to speak, with the work community itself becoming the expert system, and with ideas flowing upward from the people actually engaged in work on the organization's frontlines (cf. Doubler 1994, 58; quoted in Ambrose 1997, 67).

The COLOMBUS Experiment

A member of our group, a French national who worked at PARC, then spent time riding with French technicians to understand if their practices were similar to those in the United States. At the time we started this research, Xerox France was competing for the Malcolm Baldrige Award. According to local doctrine, quality service meant uniform service. When asked, the CSEs all said that they followed the manual religiously. However, when it became known that the PARC researcher was not from management, and interested in their real practices, they shared with him their notes on clever solutions they had discovered (see Bell et al 1997). For example, many technicians, carried cheat sheets of solutions their workgroup had invented to undocumented hard problems. Technicians who were starting on a new machine often asked more experienced technicians for copies of these cheat sheets.

In a series of workshops in France, we asked technicians whether they thought they had valuable knowledge to share beyond their workgroup. They were not sure, though they shared some stories about difficult “problem machines” and how they repaired them. Others in the workshop made comments like, “Knowing that, I could have saved five hours last week.” We asked the groups whether they would be willing to share these solutions more widely. This was not a simple question. Would a group that shared its hard-won knowledge lose their performance advantage in benchmark comparisons to other groups? And would it be worth the time and effort to document this local knowledge so it could become portable, could travel beyond the confines of the local work group?

We believed that this knowledge could have significant value, and received the backing of the French service organization – including management and the “tigers,” (the expert field engineers who played a key role in the escalation process) – to try an experiment. This required three things: an initial knowledge base of tips; a vehicle for distributing this knowledge that would be easy enough to use that technicians would turn to it; and an experimental design that would let us conduct a valid test.

We developed the initial case base by having the tigers edit and validate as tips the stories that had been volunteered at the workshops, adding more tips that the tigers themselves used. This resulted in 100 to 200 tips, structured simply in terms of Symptom, Cause, Test, and Action. The vehicle for tip distribution was a standard laptop running COLOMBUS (the French spelling of Columbus), a software package written by our group, and designed for ease of use by the technicians. A simple search using descriptive terms (like copy quality, or fault 10-200) would bring up, on an integrated “dashboard” interface, access to any material concerning this symptom, both from the tips database and from the standard documentation for the focus machine, also included on the laptop.

Interacting iteratively with the tigers in France around a prototype of this software led to rapid improvement, often overnight. This not only improved the functionality; it transformed COLOMBUS from our idea to their tool. This experience became a standard part of our design methodology throughout the Eureka project: co-design everything with the user community, making changes whenever necessary on a rapid and recurrent basis in response to suggestions and criticisms.

The experimental design for COLOMBUS tried to account for the diversity in technicians who service the target machine: how dedicated they were to repairs of just this machine, whether they their areas of work were rural or in cities (city technicians drive less and take more calls per month), and how much experience they had in photocopier repair. Forty technicians were chosen to participate in the experiment. They were given laptops (no others were available in France) and approximately three hours of training in using the software. Another forty were chosen as a control group, matched closely for dedication rate, experience and locale (often from the same workgroup). All service calls made by both groups were tracked using the standard Xerox metrics for service. These included: parts cost, service time, number of unscheduled maintenance calls, broken calls (leaving a job with the machine not yet repaired – say for a needed part), and number of callbacks (when the machine failed again quickly).

During the test, the informal response was extremely positive. For example, technicians who were not in the experimental group would ask to borrow the laptop over lunch to help them with a difficult problem they were encountering. Although this was encouraging, the metrics after two months were startling. For the same, two-month time period, the experimental group had about

10% lower parts cost and 10% lower average service time than did the control group, without differing significantly in the other service metrics.

However, this was not a long enough test to convince Xerox's Worldwide Customer Services (WCS) – the organization responsible for service strategy and technology throughout the corporation – that field knowledge was valuable and was worth investing resources to develop a COLOMBUS-like system that could be used by the entire service force. The results were convincing enough to our team at PARC, though, and – more important – to Xerox France service management. One field engineer put it this way: “This is the first time people have truly paid attention to the field, to our knowledge.” This led us to search for a way to extend the use of the tip knowledge base to all French technicians.

French Minitel Eureka

In extending the use of technician-invented solutions to the entire French service force, we faced two problems: the technical means for distribution and access that would fit into and support technicians' work practice, and the social process to make the database have continuing value. The latter required us to work with the technicians to understand what would promote ongoing participation of a significant part of the work force and ensure continual updating of the knowledge base from experience.

We could not use laptops to solve the distribution and access problem. Money for service investment was limited, and rollout in France of the laptop was at least one year away. In addition, in 1994 communication through phone lines or even the Internet was too expensive. A suggested stopgap was a printed booklet with the tips. This was not deemed an effective alternative: it would make existing information available, but not help at all to make this an ongoing and growing resource.

The French Minitel system was then chosen as the key to the distribution problem. Minitel is a nationally deployed system of the French telephone company (PTT). It consists of a small keyboard that connects to the phone line and to a local display monitor (initially a television). The PTT distributed these keyboards free of charge after realizing they could quickly recoup their cost by printing fewer telephone books and have fewer calls to information. Minitel grew into a general service that provided easy connectivity to private databases for commercial use. Xerox France technicians already used a Minitel-based service for call management – picking up new calls entered by the work support center when customers phoned in, and closing out their calls when finished.

Now we had the medium. How were the messages to be generated? We worked with a number of different communities that were part of the service force to create a design for the process: the CSEs, the tigers, and the technical support hot-line specialists who both answered customer complaints and were first in the escalation process when a CSE had a tough problem. The hot-line specialists could have viewed this as an attempt to reduce the number of jobs in their organization; but job cuts had already taken place, and they looked upon this as a way to potentially ease their workload. If it did, they could spend more time thinking about the common issues and generating their own tips. The tigers could have viewed this as “stealing” their knowledge, but they felt there were more than enough new problems. They also felt it was advantageous to get quickly out to the field any new solutions to tough problems. The CSEs liked the idea that their hard-won knowledge could travel beyond their own work group. They worried about four things, however. If they submitted a tip, would it disappear into a black hole? Would they get credit? How would they know they could trust all the tips? And how would they

get the right tips at the right time? In workshops and meetings with all the different community members, people came up with solutions to each of these problems.

Quality: To ensure the quality of each tip, each is warranted by a validator known for their expertise on the particular product line. In the Xerox France organization, there were already people designated as product specialists for each family of products in every district or “Customer Business Unit” (CBU). The tigers oversee the process to provide for a wider view.

Bottlenecks: When a new tip is submitted on the Minitel, a message is sent to the relevant group of validators, thus ensuring that there isn’t a single point of failure. One of these should pick up the new tip within a few days of its submission. Rather than just pass/fail, the validator converses with the submitter to ensure that the tip both captures the appropriate information, and is written clearly enough. The CSE gets a chance to edit the tip, and improve it, learning in the process.

Incentives: We asked the community if they thought we ought to push management into paying for each tip submitted. Not at all! As one French tiger put it, “This would make us focus on counting the number of tips created, rather than on improving the quality of the data base.” Their suggestion was to include on each tip the name of the submitter. This acts as positive reinforcement for good tips, and a negative one for badly flawed ideas that waste people’s time. The validator’s name, submission date and acceptance date also appear on each tip, providing open feedback for how well the process is working.

Integration with work practice: Since the tip database was going to be on Minitel, new screens (information pages) were added to the call handling. These allowed CSEs, when they decided they might be faced with a difficult call, to search the database based on key symptoms taken from the call record. They could also get directly in to the database from a customer’s site if they could get access to the local Minitel. The organization of the tips was simplified to Problem, Cause, Solution. Entering new tips, and commenting and voting on the success of existing tips would usually be done from the office or at home in the evening.

Implementation and deployment: Since implementing this system was a change from the plans that WCS had for the ultimate solution, and money for field service was limited, they declined to finance this countrywide experiment. A partnership of PARC and Xerox France paid for the implementation of this Minitel based system. The software was ready in about four months, but the key issue was how should the field be encouraged to use it. A champion from the French tiger group went around the country with one of our group. They talked with each group about problems with service, and how the CSEs could use this system to help themselves and others improve. They met with over 60 product leaders, and helped train all 1300 French technicians. Participation was carefully tracked, both in terms of the number of times that the database was referred to, and in terms of the number of new tips entered. There were strong differences among workgroups. While one region may have been high, another of the same size might have quite low usage rates. Revisiting the latter regions, providing some training through examples of use, and reintroducing the purpose of the system helped encourage broader participation. The strategy, then, can best be described as “hands-on participatory implementation,” which is a marked contrast to the top-down, cascade model.

Experience with use: The Minitel system opened with databases for only three products. By the end of the first year, CSEs had opened over forty databases encompassing products from convenience copiers to high-end printers. Also by the end of the first year, more than one new tip

was being added to the database each day. Over 77% of the tips were being validated within five working days, and 80-90% of the submitted tips were accepted. Participation was extraordinary. Over 20% of the CSEs had submitted a validated tip, and CSEs were consulting the tip database on the average two or more times a week.

So what were the technicians getting out of these tip documents? What did they think it was important to share? The tip content included diagnostic information, since that was obviously crucial but there was also much more. Some examples:

- Diagnosing unusual, costly failures
Bimetallic corrosion builds up on A and causes intermittent failures that seem to be B. Replacing B makes the problem seem to go away because A is moved in installation. First clean A, and later replace by new gold-plated AA, available as Part #1234.
- Workarounds
Paper curl in a dry environment causes excessive jams on baffle Q. Putting Mylar tape from tool kit on edge will ease problem.
- Easing the job
To make it easier to adjust M, paint Whiteout on the back wall near M.
- Comments on documentation
Before replacing the sensor P as the manual says in Repair Procedure R, inspect reflector Y because it might be out of alignment, causing this symptom.

Does this quite varied content make a difference to the company? The bottom line is the most telling. Comparing France to the rest of Europe, it went from being just an average or below average performer in service to being a benchmark performer. The French service metrics were soon better than the European average by 5-20%, depending on the product. This unexpected “experiment” has been made possible by the fact that laptops were not available for several years in Europe, and no other European country had the communication infrastructure in place to adopt a wide-scale Eureka process.

On a more qualitative basis, we have seen a number of different ways Eureka has affected the service process in France. In preparation for a call, using Eureka has been helpful in ensuring that a part likely to be causing failure can be picked up before going to the customer site. This reduces broken calls where a technician has to go back to the parts depot to replace the offending part. On site, Eureka accelerates and improves the diagnoses that CSEs make for the hard problems. It also reduces the number of time a call has to be escalated, thus reducing the load on the technical support hot line for recurrent calls about the same problems that start cropping up in the field. It also significantly reduced the learning curve for technicians on new products that are introduced into the field.

Spreading Eureka

In June of 1996, we decided to try bringing Eureka to another community, and to intersect more directly with the laptop rollout (only France had a system like Minitel to use as an alternative). At that time, some 6000 laptops had been deployed to Xerox CSEs, including all 1200 in the Canadian service force –comparable in size to the French. We received encouragement to work in Canada from one senior manager who wanted to help ensure success

of a new advanced color copier being introduced there. In addition, we were able to team up with a tiger from the Dorval Technical Support Center, near Montreal, who was keen about the Eureka concept and would become a local champion for its development and deployment. The challenge was to adapt what was learned from France to the Canadian service environment. We confronted an initial set of critical technical issues in this regard.

- The laptop was not widely used by those who had it. While the corporation was committed to the computer as a platform to dispense technical information and manage work process in the field, technicians had long depended on their traditional skills and practices and most were skeptical about the need for this technology.
- CD-ROMs or floppies were the primary means of distributing information to the laptop, which meant that distribution was sporadic and slow.
- Separate applications were used for call management (dispatching and tracking all customer service requests), for the now electronically presented documentation, and for parts inventory and ordering management. Thus, there was no easy way to make a connection between these independent applications and searching the Eureka tip knowledge base.

A review of how we decided to deal with these issues will provide further evidence for how technical problems necessarily interacted with the social and community context.

We couldn't directly solve the laptop acceptance problem. But we hoped that Eureka would prove to be valuable enough to technicians that it would give them an important additional reason to use the computer in their work. To achieve this, our local champion from Dorval hit upon the idea of taking existing technical information databases that had been distributed in paper form to the field and converting them to Eureka tip format. This information was already valued by the field, but hard to access in paper form -- or even keep track of, since it was difficult to keep it all in one place.

To address the distribution problem, we decided to build a server-client system, with a thick client on the laptop, and a local database that could be easily reconciled with the community database on the server. This local client/database afforded rapid access to tips while the laptop was disconnected as it usually was in the field. At the same time, it supported continual updating of the knowledge base, as frequently as a technician was able to dial in to the server.

We still had to resolve the issue of what kind of communication and server to use. Once again, we decided to build off common community practice. The common communication infrastructure for technicians at that time was a dial-in telephone connection to a bulletin board service (BBS). Some technicians used the BBS on a regular basis for discussing problems and sharing ideas, and we saw that this familiarity could work to our advantage as a platform for knowledge sharing with Eureka. There already was support for phone lines to the BBS and technicians knew the protocol for gaining access. Still, this platform had no direct support for a Eureka process workflow (submission, validation, and updating of multiple databases). This meant we had to build programs to support these in a rather baroque programming environment.

For accessing the knowledge base on the laptop by content, technician work practice dictated that our search engine had to be extremely fast, as well as easy to learn and use. A software engineer in Xerox's Printing Systems Group (PSG) had designed SearchLite, a program that had evolved through community feedback with a technical support group, and now met all these requirements; this software engineer has now become an integral part of the Eureka team. . The integration with service applications on the laptop would have been both useful and

technically possible. However, the laptop software and documentation was distributed through a central organization and maintained by them. The Eureka experiment was just that – an experiment still operating on the periphery of this organization. At this point, then, Eureka would have to remain a separate application. This peripheral status also meant that our Eureka champion for Canada had to play this role while simultaneously performing his main job as a tiger, handling escalations from and solving problems called in by technicians in the field. Validators were also volunteers who took the job on without being relieved of any of their ordinary duties.

We also had to address community process issues with regard to adapting the Eureka tip authoring, submission, and validation methods developed in France to the Canadian context.

- Product specialists did all validation in France, with field engineering overseeing the process. Would this same division of responsibilities work well in Canada? Validation had turned out to be such an important aspect of the system's success and value in France that managing this process for each different community that wanted a Eureka system was plainly an essential task.
- The French had rejected any financial incentives for tip authoring. In Canada, however, there was a financial incentive program already in place for submitting service suggestions. Should this same system be applied to Eureka?
- Because French technicians were using Minitel, they always had the most recent information at the time that they searched the knowledge base. The Canadian process would require technicians to explicitly download the latest information to update their client knowledge base. How often would they want or need to do this to make the system effective?

Because the organizational structure was very similar in Canada and France, with product specialists in each CBU, it was natural to make those Canadian product specialists the validators, just as in France. Incentive structures were another matter. Canadian service management did not want to change or give up the financial incentive program that they believed was contributing to significant improvement in service performance. Consequently, technicians received the same small financial reward for tips as for any other service suggestion. Later, the reward procedure was changed to compensate technicians for only validated tips, rather than for all submitted tips.

Updating the laptop knowledge base proved to be problematic. Not all technicians had experience using the BBS and a significant number found the process cumbersome. Moreover, the fact that Eureka was a separate application from call management, requiring dialing in to a completely different server (with a different phone number), created further complications and obstacles to frequent, easy use. As a result, when we checked after two months, a large number of technicians – roughly, 40% – rarely or never updated their knowledge base. A social process was used to try to improve the situation, with the Canadian champion making visits to each CBU to encourage updating and provide additional training.

If updating the knowledge base was a problem, upgrading the software when we made changes was even more complex. It required distributing floppy discs to everyone in the field, and hoping that they were able to use them in a timely manner. This created so many problems that a system for upgrading by downloading the software components from the BBS was eventually put in place.

Finally, Eureka was no longer only an experiment initiated by PARC researchers, with few management expectations. Even the Minitel implementation, although it involved all French

technicians, was still experimental in character and “feel.” The Canadian implementation was quite different. Eureka was now an official, management-sponsored program, with certain expectations for service performance improvement and with some financial support from one of Xerox’s major business divisions. But management had never dealt with a program where the norm was that requirements emerged from experiments with pilot users, leading to rapid incorporation of improvements, iterating till the pilot users felt it was good enough for large-scale deployment. Standard programs start with requirements that are first “locked in,” using traditional methods for interviewing stakeholders to collect them. Implementation hews to those requirements, and large-scale deployment follows quickly on the heels of limited pilot testing, and widely spaced, scheduled releases are used for any changes.

Given program expectation, managers would try to set deadlines for when they wanted things done, independent of our process for rapid prototyping and debugging with extensive community involvement. This clash of two very different design and deployment methods – one expecting people to adapt to new tools as given to them, the other expecting a co-evolution of people’s work practice and the tools they were using to support and enhance that practice – had negative results. When manager’s announced schedules could not be kept, the field was disappointed and higher-level managers lost some faith in the ability of the Eureka team to deliver. The team was discouraged.

Despite these conflicts, Eureka was successfully launched for twenty products in only six months (the launch began in early 1997). Extensive training of product specialists by the Canadian champion, with the specialists then training CSEs, took place for another four months. A training video was also created and distributed on CD-ROM, reducing the need for some of the more direct training. After six months, Canadian Eureka really took hold, and became the technicians tool.

Confronting the organizational challenge: Eureka moves to the United States

While Eureka had proved successful in less populous countries such as France and Canada, it was not clear how it would scale. The United States has 10,000 technicians spread out over a huge area. More importantly, the organizational dynamics were quite different in the U.S; it is much more bureaucratic and hierarchical, due largely to its size and thus complexity.

The U.S. and Canada shared a common laptop/BBS infrastructure; so there appeared to be no new technical issues. All the issues revolved around adapting the process to a differently shaped organization, and to the scale of the United States service force. Canada has around 100 validators for 1200 technicians. It is easy to create a community of 100, where people can get to know each other and the skills they bring to a task. The same is not true when scaled by a factor of eight. Accordingly, U.S. service management decided that validation would take place locally, with local groups selecting a validator for each product family. As with Canada, the validators would need to take on the task without any reduction in the rest of their workload.

Eureka was launched in the United States in 1997 with a pilot program in several locations. The pilot only took hold, however, where there were local champions in the service force, as was the case with both France and Canada. For the general distribution, beginning in June 1998, Eureka CD-ROMs were distributed via the mail to field managers, who were then expected to distribute them to technicians in their work groups. The CD included a Computer-Based Training module, but no hands-on training or direct engagement with technicians around the program was planned. This cascade strategy was designed for mass distribution of software or documentation, but was less effective with a socio-technical systems like Eureka. In places where people knew

about the system and became champions, or where we engaged the local group, it quickly gained traction and was widely used. In other places, it became just one of twelve programs that had to be somehow implemented over the next quarter and adoption was correspondingly slow.

We had suggested to WCS management an alternative “participatory deployment” strategy to scale for an organization with 10,000 technicians. This would have started with the pilot champions, the technicians and managers who were most knowledgeable about Eureka, going to a few other locations in the U.S. service community and talking about their experiences and ideas. Because these people were peers of technicians at those locations, they would be trusted. This would have created more local champions and knowledgeable users, who could then have gone to still more locations and followed a similar sharing method. Over a relatively short period of time, a direct, hands-on engagement around Eureka would have spread across the entire country.

The up-front cost in time and travel for participatory deployment is of course greater than for the cascade distribution, “spray and pray” strategy. But we believed that this cost would have been recouped fairly quickly because much larger numbers of technicians would have started sooner to use the program effectively, resulting in a shorter learning curve and better performance. The results from France and Canada support this argument. The main difficulty with getting management in the United States to embrace this idea was the lack of understanding of what combining the social with the technical actually required.

Nevertheless, as has been the case with each service community, United States technicians, once they finally learned about Eureka, were enthusiastic in their acceptance. One technician’s remarks are typical: “In all my years in Xerox, the two best things ever given to us are the radios and Eureka.” In fact, although the original plan was to complete rollout in the United States before moving to any Xerox organizations in Europe, Latin America, or Asia, demand from technicians in these countries for Eureka was so intense that the corporation had to begin distributing the system worldwide. This was so even in places where technicians could not update their knowledge base through connecting to the BBS because of cost and technology limitations. Here, technicians came up with the idea of using floppy discs to share and update their knowledge base.

One story that nicely captures the value of sharing the knowledge worldwide is featured in the 1999 Xerox Annual Report. A technician from Montreal traced a chronic problem with a customer’s high-speed color copier to a fifty-cent fuse holder and authored a Eureka tip. A technician in Brazil working on a similar problem with the same copier – a problem so severe the customer wanted the machine replaced – discovered this tip during a test run of Eureka, saving Xerox the \$40,000 replacement cost.

In 1999, U.S. technicians authored approximately 2000 tips. Results were promising. There were over 9,000 “solves” using Eureka in the United States and Canada in the fourth quarter of 1998 alone. The knowledge base for these problem resolutions included over 30,000 records in all: tips, Service Bulletins, and other technical information documents (the ease of distribution through Eureka made it become the standard digital “portal” for all service information). All the records were authored in English, but there was a language translation function that could be used on selected phrases and terms. Tips from the U.S., Canada, and several European countries made up about 20% of the total and were the most rapidly growing portion of the knowledge base.

By the first quarter of 2001, the size of the database had grown considerably as the number of countries using Eureka increased; with close to 50,000 technician-authored tips and over 300,000 records in all. Moreover, the number of problems solved using Eureka had increased to nearly 200,000 annually. In 2001, it is estimated that there will be over to a quarter million solves. Since a solve can save a few to many hours of frustrating (to both technician and customer) down time, possible escalations to experts, and sometimes the replacement of a machine, the use of Eureka provides many millions of dollars in savings annually for the company. As “side” benefits, both customer and employee satisfaction is increased.

Having reviewed the story of how Eureka was designed and deployed in three different service communities, focusing on the adaptation to their different organizational and technical environments, we will now try to answer three questions based on studies in the field. How have users responded to their experiences with Eureka? How have they adapted it to their work practice? What barriers to more effective use have they noticed?

Field Studies of Eureka in Practice

After Eureka had been in the field in the United States for six months, a member of our research group spent some twenty days over the course of three months traveling with and talking to technicians from in San Francisco, concentrating on one work group in particular. He and other members of our group also visited four different CBUs around the country to try to understand the user’s experience of this new tool. We did this through group and individual interviews, intending to frame the discussion by asking technicians if Eureka was a tool worth using, and if so, how they used it and how it could be made better. However, as soon as they discovered that we were from the team that had designed and helped launch Eureka, they volunteered remarks like:

Best reason we have for having a laptop. I use it on probably 50% of the calls that you go to where I don’t walk in the door and immediately go “Well, this sensor’s broken” or something like [that]. Anything where something doesn’t immediately jump out at you, it’s the first thing I turn to.

We also had specialists tell us that they tried to encourage everyone to use it when they have a problem they can’t solve, as the following remarks indicate.

Most of the time before I get to a site, I try and figure out some directions to go. I look around in Eureka and see what’s out there so that when I get there I know what I’m gonna do.

If we’re going to an escalation, we pull it up before we get there, if it’s not something that we’re familiar with, something that may be new. So it’s something we’ve gotten in the habit of doing, so we don’t even think about it. It’s very key and very important that way. A lot of times you’ll read something and you’ll think, “Man, I never would of thought of that.” Now maybe it doesn’t actually work that time, but now you’re thinking about a lot of things that will work that wouldn’t have ever occurred to you. You know there’s another train of thought, another set of eyes.

This was not surprising, since when first designed, Eureka was conceived primarily as a “tool of (near) last resort,” when routine fixes fail to resolve a problem – particularly the intermittent ones that are so difficult to diagnose – and past experience don’t point to an answer or line of attack. And many technicians only use Eureka that way, whether with machines they

are working on or as a source for suggestions to be given over the radio when colleagues are asking for help. But we have discovered that technicians use Eureka in several other interesting ways.

Most striking is that many use it as a tool of first rather than last resort. For example, one technician who works on high volume copiers told us how he used Eureka in combination with the product documentation before he went out on a call.

Before I go on a call I just like to see what picture I'm looking at as far as some [possible] fixes [in Eureka]. Then I drop into the documentation, if I felt that there wasn't anything [in Eureka] that jogged my memory. And then I'm on my way. Because keeping that footprint of what some of the fixes were and then just going down through [the repair procedures in the documentation], it just accelerates things.

Thus, even before seeing the machine, this technician tries to develop several solid leads regarding the source of the problem, the likely repair procedures necessary, what parts were needed and the like.

Many other technicians report similar patterns of Eureka and documentation use, as in the following descriptions:

I check Eureka before I go to a site, that way when I get there I already know what parts to bring in. The customer is always impressed when you show up and already have an idea on how to fix the machine, and have the part with you. They just love it. It also saves me a lot of time. Whether it's a whole trip back to get a part, and then go back to the customer, or it's just a trip back to the car, I can use any spare time I can get.

[Eureka isn't] so much an end as it's a beginning. Someone will call up over the radio with a fault code, [like] "I'm having 12-142s," and I can look it up in Eureka and can start scrolling through common causes. Actually it's faster for me to find it in Eureka than it is to go in and fire up the documentation CD and go for the repair procedures [in there]. And not only that, getting into the book is just the book, you know, where here [in Eureka], if I bring it up and get a clue, not only am I gonna get a clue what that jam is, but also the potential area of fixes.

This technician also reported that he felt Eureka was useful even on those occasions where the tip didn't provide the precise solution because it allowed the team to rule out certain sources of trouble, thus narrowing the search for the correct one.

Finally, technicians also use Eureka as an informal learning tool. One technician who services mid-volume machines told us that browsing through the tips allowed him to see what has worked for others.

And that's basically what I do. Whenever I download [new Eureka data] I like to look through it. See what guys are doing. I go looking through the tips and Bulletins. It just sort of teaches me a lot.

Essentially, by reading the tips and Service Bulletins somewhat casually, divorced from any actual repair situation, this technician is able to use Eureka as an instructor who offers a new set of lessons each week. This technician went on to note that attending to how others found workarounds has helped him develop his own skills. That is, beyond the actual content of the tips, he found reading how others solved problems useful in broadening his own way of working, his own knowledge.

But we also identified some barriers to effective use. Laptops have a long boot-up time, and technicians hate to wait for it. Leaving the Laptop on all the time is not a solution because of limited battery life, and it is often difficult to find an available electrical outlet at the customer's site. Other difficulties include the instability of the Windows operating system (crashes are frequent), and difficulties when in the field of finding a free phone line to dial in to the Xerox server and update the database.

In addition to these technical barriers, however, many technicians simply mistrust, are unfamiliar with, and resent computers, and so don't use the laptop except when absolutely necessary. A representative remark from a technician about this issue:

Half my team is basically uncomfortable on a computer no matter what's on it. They use the laptop as little as they can. They clear calls [using the call management function on the laptop] and that's about it. The real problem is getting them to adopt the laptop generally, not Eureka.

Other technicians told us they do not use the laptop because they feel it adds time and work to their daily routine. These technicians explain that when laptops were first introduced they actually added workload by forcing technicians to do call handling on the computer, a task they previously accomplished by phoning in to a Xerox work support center. That is, it was done manually, through a relationship to human beings and using a dependable device, the telephone. Two years passed, they claim, before time saving software was introduced to make the call handling application on the computer more efficient and useable. Thus, while the laptops certainly have the potential to make technicians more effective in the field, at least some CSEs feel laptops have made their lives more difficult. The first question CSEs who feel this way often raised with us (when we were introducing Eureka to them) makes this point clear: "Is this another call-handling situation, where they are using the computer to try to force us to do more work?"

Our regular field observations confirmed these reports. CSE's have both Laptops and radios. We found we could travel with CSEs for several days and see the laptop used only once or twice a day, and usually only for call handling, with some work team members never using it. Many teams have developed workarounds to the computerized call-handling system, with one member of the group logging on to the system, usually by connecting from a Xerox parts drop and almost never from a customer site, checking the queue and taking care of all the dispatch assignments for her co-workers. The radios allow workers to keep each other updated on the status of the call queue for the group, and to communicate about dispatch assignments. In general, then, many technicians have not easily integrated the laptop into their everyday work practice. Radios on the other hand have become an indispensable tool for communication and collaboration and thus an integral part of the technician's work life. In part this is because radios are always on (no booting up needed), provide a persistent connection to their work community, have long lasting batteries, a simple interface (no operating system) and rarely crash. Finally radios serve ongoing social and technical purposes in the work practice of the group.

Another barrier technicians identified was that they had to do independent searches in the Eureka knowledge base, and the e-doc. This required them to enter information several times,. Too often the division between the tools that technicians had to use reflected their separate institutional development rather than the needs of those for whom they were designed. For example, e-doc had been completely separate from Eureka because it was developed separately. The integration of all the tools and databases was the biggest "ask" received from the field that

came out of our feedback meetings, and was a primary design criterion for the next generation Laptop. Finally they wanted it to be easier to move between tasks. Putting Eureka on the BBS separated the connecting and updating from the more routine connecting to the call management system. As one CSE put it, “Is there any way to move Eureka out of BBS and into a system we routinely use on a daily basis?”

We see, then, that technicians not only use Eureka in creative ways, they are regularly thinking about making it more effective for their work. This is exactly the kind of inventiveness from the field that Eureka was meant to capture. It stands as further evidence of the pervasive importance of working with the users to make the system fit their needs – to artfully integrate the technology with their enhanced practice.

Eureka II

Even before we did our field studies, we recognized many of the issues that were facing the technicians in using the laptop, and the opportunities new technologies were presenting, in particular the advent of cost effective communication on the Internet. Although we will not go into detail about the new web-based implementation of Eureka – named “Eureka II” – that has since been rolled out worldwide, we will mention a number of features in terms of the problems they solve for the work practice of technicians. Eureka II and the accompanying laptop suite of tools for technicians were developed by a multi-organizational team that included Xerox personnel from PARC and Worldwide Customer Services, with additional programming skills obtained through contracts with software houses.

To bring together everything that technicians needed to do when connected, Xerox has deployed a global service network with multiple servers. As technicians log on to the call management system to get their next service call, the Eureka web server is able to download in the background any updates to the knowledge base. This same mechanism updates the e-doc, and if necessary updates the software on the laptop.

All the information sources are accessible through a single search mechanism based on SearchLite. So when a technician has a problem, they can see in their “hit-list” references to tips, and to multiple places in the formal documentation where they may find helpful information. In addition, annotations can be made on already existing documentation. Such “post-its” can be kept in a private knowledge base, or if desired submitted as a tip. When validated and shared, an annotation appears not only in the hit list directly, but as a link on the page where the annotation was made.

The success of Eureka as a knowledge-sharing framework has inspired other communities to adopt this practice. For example, software analysts have started using Eureka. But this required some adaptation. The Eureka tips had a minimal structure with explicit fields for only Problem, Cause, Solution. For some tips, software analysts wanted to include a field for Software Patch. Tips dealing with interactions of different software systems could have a separate field for Operating system, etc. In Eureka I, these changes had to be added by the software team. Eureka II includes a tool that allows user communities to define and evolve their own fields (through XML) and their own ways of viewing these records (through HTML frames). This kind of extended customizability is pervasive in Eureka II.

We abandoned the bulletin board system. It had too many features, both administrative and functional, that were unnecessary for Eureka. We created a much simpler web server that would hold the knowledge base, and support the simple workflow from authoring to validating to

subscription management to downloading. Simplification and stripping out the unnecessary were our goals.

A pressure in the other direction came from the fact that Eureka was so successful that it became a mainline program. This meant that requirements poured in from many places; we were constantly trying to balance our belief that simpler was better with corporate managers' beliefs that if Eureka was the answer, they wanted to be the one to generate the question.

For example, one manager's intuition was that one of the biggest costs for the system would be in training the technicians to use the system. He felt that the best way to simplify the training was to make the Eureka application be embedded in the standard Internet browser (in this case Internet Explorer). We explained that this would complicate the implementation significantly because the software would then be dependent on exactly which version of the browser, which operating system, and which service packs were loaded on each machine. His response was that this was not as important as simplicity of training. We explained that knowing how to browse had only minimal overlap with knowing how to use the functionality of Eureka for searching, authoring, etc. But as manager of this major program, with long experience working with the field, he insisted that the tradeoff would be worth it. The program subsequently made that choice; unfortunately, we found that there were far more implementation and deployment complications than even we imagined, and the delay introduced in deploying Eureka II was significant. In hindsight, we all agree that the tradeoff would have been better made the other way. The issue here is not that a manager was wrong, but rather how decisions are made in a standard software development process, in contrast with the bottom up approach with which we started.

In addition, because the project now had the attention of top management, schedules were sometimes set based on managers' desires to achieve certain goals by certain times, rather than on the work that was necessary to achieve the desired state as estimated by the implementers. We understood the pressures that these managers felt, but none-the-less this often led to internal battles, and then to the appearance of slipped schedules as managers' pronouncements could not be realized. It made us very aware of the difference between singing in the spotlight and singing in the shower.

Building Community Knowledge Systems

How much of the Eureka story can be generalized for other organizations who wish to build a similar, socio-technical system for knowledge creation and sharing? In this section, we try to capture methodological issues for system design. These are presented as key questions to address in building a system that puts the work community at the center – where this community becomes the expert system.

In reflecting on the design process we followed in building the first version of Eureka on Minitel, and then the later client-server version for Canada and the U.S., we realized that we had been continually trying to address a set of closely interrelated problems. These had to do with community membership and relationships, the kinds of knowledge community members regularly shared with each other and their motivations for doing so, the existing practices and contexts through which this sharing took place, and the process for implementing a socio-technical system to enhance these practices. Using this framework, we can then extrapolate the following questions as the basis for our design methodology.

- **Community: who and where**

- § **Who are the members of this work community?** Shared identity and practices define “community.” Because members share practices, communication between them can draw on a great deal of background understanding, knowledge that doesn’t have to be explicitly stated. The shared practice and common environment for the work support all terse descriptions, indexical terms, and elision of “obvious” steps in a process description. For these reasons, it is much easier to build a knowledge sharing system that is based in community life, and stays within community boundaries, than it is for one that crosses distinct boundaries. Moreover, community membership is the basis for trust, and effective knowledge sharing depends on trusted information. In the case of Eureka and the Xerox service community, only the key information needs to be presented in a tip, not all the steps involved in fixing a problem. And since tips are written by technicians for other technicians, the information is not only understandable in context but also trustworthy.

- § **Do members work in close proximity to each other or at a distance?** Knowledge sharing can take many different forms, using different media if people are working together in a single location or if they are primarily alone. Working shoulder-to-shoulder supports continuous apprenticeship learning where knowledge can be shared that has not yet been articulated and documented. For people working primarily in separate locations, documents – since they can travel across time and space with ease – will be especially important for sharing and learning. Moreover, when the community is large, even if some groups work together, documents help in scaling knowledge more rapidly across numbers, time and distance. For example, Xerox service technicians spend most of their time in the field at customer sites, usually at some distance from each other. Although they gather together at parts drops and for meals and occasional meetings, and may accompany one another to some customer locations to collaborate on especially obstinate problems, they are commonly alone and have only their radios for communication. These features of service work mean that extensive community knowledge sharing requires digital documents that can be read on a portable device like a laptop.

- **Knowledge: what and why**

- § **What constitutes especially valuable knowledge for this community?** The answer to this question is not always obvious to people who simply look at a formal description of the work process for a job. Observation of how people do their work will reveal what kind of information is most often shared, because it is highly valued, in actual practice. For example, in the service world, we saw that it was not only diagnostic tips that were valued, but also hints about how to make certain tasks easier, corrections or improvements to documentation, and workarounds for particular engineering shortcomings. The reality of servicing machines in different environments, with different machine histories and customer profiles, makes necessity the mother of invention in this work. On the other hand, one would expect that consultants might share different kinds of information such as proposals, approaches to take with clients, etc. Sales representatives do not want to share their leads, but are happy to share stories of their successes.

- § **Why do members share particular kinds of knowledge?** Understanding the motivations for sharing is important for grasping the natural incentive structure of the community. If a knowledge sharing system is to be successful, it should build upon this structure. External rewards can also be used to encourage sharing, but there may be

danger here in assuming that financial payoff is a naturally effective way to get quality information and participation. With service technicians, we found that getting their job done more effectively and building a reputation for competence was a significant incentive.

- **Sharing: how and when**

- § **How is sharing taking place in the community on an everyday basis?** This is perhaps the most critical question to answer, because an effective knowledge sharing system should honor natural sharing practices, and the style people follow to exchange information, seek and give advice, and otherwise support each other. Service technicians use stories of particular machines and their problems to share their learning and experience. The style of the tips, although they are written documents, tends to follow this narrative structure. These stories are naturally shared under two interestingly different circumstances.

- § **What are the different work contexts in which the sharing commonly occurs?** The issue here is the context of sharing during work. When a technician has just found a particularly recalcitrant problem, or found an ingenious workaround, they will often tell the story the next time their workgroup gets together. This volunteering is often “just-in-time” because when a problem crops up on one machine, it may mean that it is going to come up in other members of that machine population. On the other hand, old stories are brought up when people come to the group to help. Then the story is used to suggest possible unexpected linkages between symptom and cause. This suggests that there is utility to supporting both notification of what is new in a knowledge base, and a way of retrieving “similar” stories, even if a tip does not directly solve a problem.

- **Implementation: how to**

- § **How to determine what constitutes effective technological support for work practice?** Our experience has strongly suggested the value of bringing a prototype to a pilot group in a community, and using rapid prototyping, and participatory design. The initial prototype provides something for community members to react to. Their reactions and the use they make of the prototype – both give indications of the ways that the technology needs to change. However, even when technology cannot change (for technical or economic reasons), the experience may lead to a better process, for community members are often good at inventing workarounds. Inventive community members also will use the technology fruitfully in unexpected ways.

- § **How can the new system be learned?** Learning for knowledge sharing is as much about what is valued and how it should be expressed, and how to find the knowledge that you want in a system, than it is about learning about technology *per se*. It is also about having the incentive in the right context for doing the learning. When people need knowledge, it is an excellent time to teach them how to get it. Learning should thus become a common, everyday activity in the use of the system rather than be seen as merely an initial training activity separated from the work. So for example with service technicians, when they call the tigers for help on a problem, the tigers may look in Eureka to see if there is an appropriate tip. If there is, the tiger can use this as an opportunity to teach both the value of the Eureka system and how to use it.

Organizational Barriers to Change

At the end of the day (really seven years), the story of Eureka is a tale of how the development and deployment of a system for sharing knowledge created on the front lines became a vehicle for organizational change. However, our story also reveals that this change was not without conflict and challenges. These are precisely the kinds of messy details that are rarely included in the writings on knowledge management or organizational learning. It is worthwhile to expand a bit on these details and their larger meaning for knowledge sharing in practice.

In the initial stages of the project, during the experiment in France, few people in Xerox management believed that there was much value in what the technicians learned on their own in the field. Additionally, they could not see how a tip system from the field was much different than prior suggestion systems, all of which were highly centralized and controlled. And although technician tips quickly proved valuable, people in different organizations in the company felt that it was more important to get centrally produced documentation in front of the technicians than support them in creating new knowledge. This new way of doing business made them nervous; for example, when a single flawed tip eventually slipped through – this happened after the U.S. deployment began, with a tip that was appropriate for Europe but not the United States due to the different electrical requirements – the doubters went into an “I told you so” mode, even though this was an extremely rare occurrence and was soon caught by the oversight process that was put in place.

With these obstacles, getting support in the form of organizational resources for the project naturally proved difficult. Time and again, as the Eureka story makes clear, we ended up relying on local champions who somehow managed to cobble enough resources together to get the job done. Moreover, in the initial stages, we had to sometimes operate like a guerilla group because opposition to the idea of a bottom-up knowledge system based on the value of field know-how was enough to kill the project if we openly challenged such deeply-entrenched convictions; our first experiment in France was conducted there in part because it was out of sight of the central Xerox organization. After France, it was only by convincing one product manager in a business division to give us the funds to try the Canadian experiment that we were able to finally get data that would convince the non-believers. (Interestingly, one of our research group members, the person who led the French effort, was later given an award from the Worldwide Customer Service organization for getting Eureka launched, with the award plaque reading “despite the resistance of WCS.”)

And speaking of convincing data, we recognized at the start that informal response to the collection of tips and the users informal assessment of their value would not be accepted by the service organization. We knew we had to show hard, “bottom line” data from believable experiments. But comparative studies are more difficult if people are clamoring to get a feature. Then everyone improves, but there are no comparative metrics. In some ways, resistances in the organization to making a wide scale change like Eureka worldwide in just one step allowed us the opportunity to gather better measures of the kind they would accept. We believe that some of the more important measures are the responses of the field to being listened to and the potential for changing how we do business. For example, we could put out a product with only minimal diagnostic documentation (getting the new product to market more rapidly) and then use the field force to help understand where and how it needs improvement – that is, to construct the diagnostic documentation in the field.

Once Eureka did begin to receive official recognition and support from Worldwide Customer Service (they assigned one of their best people as the WCS champion), and as the performance data and, more than anything, the enthusiasm of the technicians themselves overcame most resistance, deployment in the U.S. was approved. Then, however, the project ran into a very different sort of problem, one that could be described, perhaps, as resulting from success. Eureka became a major corporate program. But why was this a problem? In France and Canada, since we were conducting guerilla experiments, we could involve the users more in the decisions of how to adapt Eureka to local needs and practices. Moving to the central WCS organization, however, engendered a change of philosophy. As Worldwide Customer Service, they felt they had to design a uniform, worldwide solution. The power was for the most part wrested from the users' hands.

Related to this, management had a policy of distributing corporate programs in a hierarchical ("cascading") fashion. Eureka works much better when peers mentor each other on the uses of the system. The French and Canadian experiences demonstrated the value of participatory deployment. Unfortunately, this was not how the rollout in the United States was done. While the U.S. deployment was eventually successful, the problems that developed because technicians weren't closely involved in the process (especially those technicians who had the most experience with Eureka and could have been an invaluable resource for their co-workers in other parts of the country) slowed the project's achievements over the first year and more.

This can be explained in large part by recognizing that WCS, and Xerox Corporation more generally, has emphasized cost savings in field service, and has viewed the service organization as a cost center. As a result, significant investment in service is discouraged, unless it is matched at the time of investment with equivalent cost reductions (instead of seeing investment as a front-end expense that will have a back-end, long-term payoff). This had consequences for Eureka deployment, and for how the program now operates. When the service organization was asked to change how it viewed the importance of field knowledge, the field technicians who were charged with contributing to the new knowledge base and maintaining its quality were not given enough time to learn and engage in the change activities, because this would have required a significant investment in time (and thus money). As a result, there was not enough time to reflect on and practice the new way.

Although the president of Xerox said that Eureka was a key program, we could not get sufficient support from operating organizations for building the kind of process infrastructure – such as training resources and time – to make it more successful. They had serious cost concerns, and did not want (or were not able) to make the necessary investment. More important, with respect to program operation, technicians were expected to author tips and validators were asked to provide rapid turnaround and validation of submitted tips – a critical role in the process – without giving any of them relief from their current workloads.

Finally, we have to ask; did Xerox become a better Learning Organization as a result of the Eureka project? One way to answer the question would be to point to the impact on field service performance, and to the degree to which most technicians use the program on a regular basis and in many different ways, including especially learning new ideas and approaches to machine repair. The service organization, taken as a whole, has also been transformed to some degree by the bottom-up Eureka approach, which has had an impact on this particular organization's operational philosophy.

Although these are certainly significant achievements, the corporation – the corporate “organization of organizations” – has not yet taken full advantage of what a knowledge-sharing program like Eureka makes possible. For example, although there is informal use by many organizations of the Eureka knowledge bases, no formal process is in place, for example, to incorporate some of the information in Eureka back into the documentation. Engineering could, but doesn’t regularly, mine the Eureka knowledge base for ideas for continuous improvement. And manufacturing could use Eureka to augment the information flow it needs to make rapid adjustment at the initial launch of a new product. ∴

The current Eureka process, which is dedicated to technicians authoring tips for fellow technicians – and not, then, for documentation authors, engineers, quality control experts in manufacturing, or trainers – obviously cannot address all these areas or solve the inter-organizational knowledge transfer problem. It does point the way, however, toward what a solution could be, and to the need for making use of knowledge collected on the corporation’s frontlines (no one in the corporation spends as much time with Xerox customers or knows as much about Xerox machines as technicians) throughout the corporation.

At the same time, what we could perhaps describe as the “spirit of Eureka,” which is the belief in the value of bottom-up knowledge sharing, of the importance of drawing directly on the experience and creativity of frontline employees, has had some interesting effects on the corporation as a whole. The most interesting of these is the number of requests to our group at PARC from different parts of the corporation for helping to create a “Eureka-like” knowledge system for their organization or operational unit. Presently, we have deployed LinkLite, a new simpler infrastructure to support the Eureka process, in a Xerox sales organization. Sales people share knowledge about “customer solutions” (special configurations of machines and services that help customers solve important business problems), what their successes have been, and material that supported the sale. It may well be that the belief in and spread of this “spirit” becomes the most important legacy of Eureka in Xerox.

Acknowledgements

We would like to thank two people whose work has been crucial to the success of Eureka: Olivier Raiman, the original PARC champion for Eureka, and French national who led the effort in France and Canada; and Bob Cheslow, the software engineer behind SearchLite, and later the major architect and implementer of all the versions of Eureka, including the current LinkLite system. We are also grateful to others who have been members of the extended team over time: David Bell and Mark Shirley from our group at PARC; Eric Delanchy, a French tiger; Michel Boucher, a Canadian tiger; and the many technicians worldwide who graciously shared with us their experiences. We also appreciate the support of PARC management, particularly John Seely Brown and Johan de Kleer, without whom we could never have made the project successful; and the people in world wide customer service, Bernie Colligan and Tom Ruddy, who have helped to make Eureka real and continue to live and pay off in a community of 20,000 people. To others who have helped us along the way but who we have not named explicitly, we thank you.

References

- Ambrose, S. E. 1997. Citizen Soldiers: The U.S. Army From the Normandy Beaches to the Bulge to the Surrender of Germany, June 7, 1944-May 7, 1945. Simon and Schuster, New York.
- Bell, D. G., D. G. Bobrow, B. Falkenhainer, M. Fromherz, V. Saraswat, M. Shirley. 1991. RAPPER: the copier modeling project. International Logic Programming Symposium, San Diego, CA, October.
- Bell, D. G., D. G. Bobrow, O. Raiman, M. Shirley. 1997. Dynamic documents and situated processes: building on local knowledge in field service. In Information and Process Integration in Enterprises: Rethinking Documents, edited by T. Wakayama, S. Kannapan, C. M. Khoong, S. Navathe, J. Yates. Kluwer Academic Publishers, Norwell, Ma.
- Brown, J. S., P. Duguid. 1991. Organizational learning and communities-of-practice: toward a unified view of working, learning, and innovation. Organization Science 2 40-57.
- _____. 2000. The Social Life of Information. Cambridge: Harvard Business School Press.
- Davenport, T. H., L. Prusak. 1997. Working Knowledge: How Organizations Manage What They Know, Boston: Harvard Business School Press.
- de Kleer, J., B. C. Williams. 1987. Diagnosing multiple faults. Artificial Intelligence 32 97-130.
- Doubler, M. D. 1994. Closing With the Enemy: How GIs Fought the War in Europe, 1944-1945. University of Kansas Press, Lawrence KS.
- Nonaka, I., H. Takeuchi. 1995. The Knowledge-Creating Company: How Japanese Companies Create the Dynamics of Innovation. Oxford University Press, New York.
- Orr, J. 1997. Talking About Machines: An Ethnography of a Modern Job.
- Senge, P., A. Kleiner, C. Roberts, R. Ross, G. Roth, B. Smith. 1999. The Dance of Change: The Challenge of Sustaining Momentum in Learning Organizations. Doubleday, New York.

Stewart, T. A. 1997. Intellectual Capital: The New Wealth of Organizations. Doubleday, New York.