

# Community Knowledge Sharing in Practice: The Eureka Story

*Daniel G. Bobrow and Jack Whalen*

An organization's most valuable knowledge—its essential intellectual capital—is not limited to the information in official document repositories and databases, such as scientific formulae, “hard” research data, computer codes, codified procedures, financial figures, customer records, and the like. It 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 informal (often tacit) knowledge is deeply rooted in individuals' experiences and the culture of their work communities. It commonly originates as practical solutions—through everyday inventions and discoveries—to the problems they must solve and thus serves as the critical resource for ordinary work practice (see, especially, Brown and Duguid, 1991, 2000). Much of this knowledge often remains embedded in practice. Small circles of colleagues and work groups commonly share crucial steps in a new practice and fresh solutions to recalcitrant problems through conversations and stories, with members filling in the background and gaps from their own experience. These groups and communities use the local vernacular to express these instructions and stories.

Organizations face the challenge of somehow converting this valuable but mainly local knowledge into forms that other members of the organization can understand and, perhaps most important, act on. Here we present a detailed account of one organization's effort to encourage inventiveness, capture new ideas, and use technology to then share the best of this knowledge beyond a local work group.

Our account is based on our experiences during seven years with the design, development, deployment, and evaluation of the Eureka system at Xerox Corporation. Xerox uses Eureka to support the customer service engineers (CSEs) who repair the copiers and printers installed at customer sites. In four iterations, the system went from an experiment that researchers at the Xerox Palo Alto Research Center (PARC) designed to measure the value of codified field experience to a system deployed to 20,000 CSEs worldwide. By focusing 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 involves pushing changes within a corporate culture; understanding the Eureka experience and the problems facing all knowledge systems to be deployed in the real world requires equal focus on these challenges.

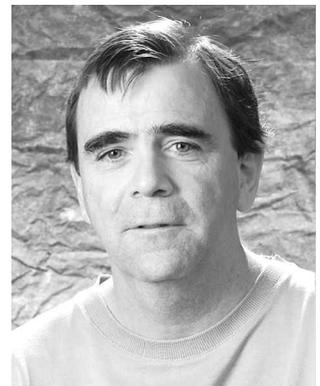
Our narrative covers the history of this project, carefully detailing the fundamental interrelationships between the social and the technical. We include a framework for building these kinds of community systems (see the sidebar) and our reflections on the barriers to organizational change that their proponents confront.

## Breaking the Frame

Xerox has more than 20,000 technicians worldwide who help to ensure that Xerox machines are performing as customers expect. As Orr (1996) pointed out, this is a triangular



Daniel G. Bobrow  
Research Fellow  
Systems and Practices Laboratory  
Palo Alto Research Center  
bobrow@parc.com



Jack Whalen  
Principal Scientist  
Systems and Practices Laboratory  
Palo Alto Research Center  
jwhalen@parc.com

## Building Community Knowledge Systems

How much of the Eureka story can be generalized to other organizations that want a similar, socio-technical system for knowledge creation and sharing? Answering these questions can help build such a system.

### Community: Who and Where

- **Who are the members of the work community?** Shared identity and practices define “community.” Because members share practices, communication between them can draw on background understanding or knowledge that doesn’t have to be explicitly stated. It is easier to build a knowledge-sharing system based in community life that stays within the community than 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, technicians write tips for other technicians, so the information is not only understandable in context but also trustworthy.
- **Do members work in close proximity to each other?** Working shoulder to shoulder supports continuous apprenticeship learning in which people can share knowledge that has not yet been articulated and documented. For people working primarily in separate locations, documents are especially important for sharing and learning. Moreover, when a community is large, documents help scale knowledge more rapidly across numbers, time, and distance. For example, Xerox service technicians spend most of their time alone in the field at customer sites. Extensive community knowledge sharing requires digital documents that they can read on a laptop.

### Knowledge: What and Why

- **What constitutes valuable knowledge for the community?** Observation of how people do their work will reveal what kind of information they most often share because they value it. For example, we saw that technicians valued not only diagnostic tips but also hints about making certain tasks easier and corrections or improvements to documentation.
- **Why do members share particular kinds of knowledge?** Understanding the motivations for sharing is important for grasping the natural incentives within the community. Successful knowledge-sharing systems should build on this structure. External rewards can encourage sharing,

relationship among the technician, the customer, and the machine. On many service calls, the technician needs to repair or adjust the machine; on some, the technician needs to help the customer adjust his or her expectations, procedures, or knowledge of the machine. In the early 1980s, because technicians trained by the armed services to debug complex equipment became increasingly unavailable, Xerox decided to use less skilled, less experienced service people. It moved away from the documentation and training that described the principles of product operation, which required skilled technicians to determine the appropriate repairs. It moved toward “directive” repair and adjustment procedures or documented instructions in a decision tree. Each decision step was in the form of “do the following setup and test; make the following measurement (or observation); if the result is A, do X; or else do Y.” The intuition embodied in this form of documentation is that technicians need only be trained how to use the documentation correctly to diagnose and repair any machine failure.

### *The Rapper Project*

Our group at Xerox PARC has a background in artificial intelligence and in modeling electromechanical systems. In particular, we have expertise in building programs that diagnose machine faults given an abnormal symptom and the ability to observe or get 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 guide a technician in diagnosing and repairing problems in this module. Our hypothesis was that if we were successful, a model-based expert system on a laptop that technicians carried in the field could replace the documentation and support a work process for isolating faults. In addition, this approach could improve Xerox’s speed in bringing supported products to market, because the need to create documentation inhibited deployment. Models could be created in parallel with design. Moreover, newer machines often used the same or similar subassemblies, making models reusable.

but there may be a danger in assuming that financial payoff is a naturally effective way to get quality information and participation. The service technicians felt that getting their job done more effectively and building a reputation for competence was a significant incentive.

### Sharing: How and When

- **How does sharing occur in the community every day?** 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 tell 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.
- **In what different work contexts does sharing commonly occur?** When a technician finds a particularly recalcitrant problem, he or she will tell the story at the next work-group meeting. This volunteering is often “just in time,” because when a problem crops up in one machine, it may come up in others. On the other hand, when people come to the group to help, they bring up old stories. Then they use the story to suggest possible unexpected linkages between symptom and cause.

### Implementation: What and How

- **What constitutes effective technological support for work practice?** Our experience strongly suggests the value of bringing a prototype to a pilot group in a community for participatory design and rapid turnaround in response to suggestions. The initial prototype provides something to which community members can react, which can indicate how the technology should change. Inventive community members will use the technology fruitfully in unexpected ways.
- **How can people learn the new system?** Learning to share knowledge involves learning what is valued, how to express it, how to find the knowledge, as well as learning about the technology per se. It is also involves having the incentive in the right context for learning. Learning should become a common, everyday activity in using the system, rather than an initial training activity separated from the work.

We succeeded in building Rapper (Bell et al., 1991), an expert system that used a model of the recirculating document handler to guide in isolating faults in that module. The model captured all the faults found when using the standard documentation. We asked technicians if a complete model for the machine would be useful. “Not really,” they said, “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 behind their negative response. First, to them, small improvements in the time required to isolate a 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. For many products (those produced by our Japanese partners), however, there were no full descriptions of operation. Additionally, the diagnostic documents were produced by inserting faults in the machines in a laboratory and then recording the symptoms. So the hardest problems were not those covered by the documentation; they were new problems.

We decided to spend more time observing what technicians actually did day to day. We started with US technicians, accompanying them on their service calls. Most of the time, they would look at the machine, talk to the customer, and know exactly what to do to put it in good working order. Occasionally, they ran into a problem that they hadn’t seen before and for which there was no documented answer. 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 two-way radios, or turn to the experts—former technicians now serving as field engineers—who were part of the escalation process. When they solved unusual problems, they would often tell stories about these successes at meetings with their coworkers. The stories, now part of the community, could then be used at similar gatherings and further modified or elaborated (see Orr, 1996; Brown and Duguid, 1991).

*A model-based expert system on a laptop that technicians carried in the field could replace the documentation and support a work process for isolating faults.*

*Would it be worth the time and effort to document the local knowledge just so it could travel beyond the confines of the local work group?*

This practice pointed to the importance of noncanonical 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; the work community itself could become the expert system, and ideas could flow up from the people engaged in work on the organization's frontlines (cf. Doubler [1994: 58]); quoted in Ambrose [1997: 67]).

*The Columbus Experiment*

A member of our group, a French national who worked at PARC, spent time with French technicians to see if their practices were similar to those in the US. At the time we started this research, Xerox France was competing for the Malcolm Baldrige National Quality Award. According to local doctrine, quality service meant uniform service. When first asked, the CSEs all said they followed the manual religiously, but when they found out that the PARC researcher was not from management, they shared their notes on their own clever solutions (see Bell et al., 1997). For example, many technicians carried cheat sheets of solutions their work group had invented to solve hard, undocumented problems. Technicians working on a new machine often asked more experienced technicians for copies of the cheat sheets.

At a series of workshops in France, when we asked technicians whether they had valuable knowledge to share beyond their work group, they were not sure, though they shared some stories about how they repaired difficult "problem" machines. Another CSE, hearing the story, commented, "If I had known that, I could have saved five hours last

week." We asked the technicians what the issues were if they shared hard-won knowledge. Some feared that they would then lose their performance advantage in benchmark comparisons to other groups. Others wondered if it would be worth the time and effort to document the local knowledge just so it could travel beyond the confines of the local work group.

But they (and we) believed that this knowledge could have significant value. The French service organization, including management and the "tigers" (the expert field en-

gineers who played a key role in the escalation process), gave us the backing to experiment. We required three things: an initial knowledge base of tips, a way to distribute this knowledge that would be easy for technicians to use, and an experimental design on which we could conduct a valid test.

We developed the initial case base by having the tigers edit and validate the stories that technicians had shared at the workshops, adding more tips that the tigers themselves used. The result was 100 to 200 tips, structured simply by *symptom*, *cause*, *test*, and *action*. We used a standard laptop running Columbus, a software package that our group wrote, to distribute the tips. A simple search using descriptive terms (such as copy quality or fault 10–200) would bring up, on an integrated "dashboard" interface, any material containing these terms, both from the tips database and from the standard documentation, which was also included on the laptop.

We worked interactively with the tigers in France to improve the software, often responding overnight; this transformed Columbus from our idea to their tool. This interaction became standard to our design methodology throughout the Eureka project: we codesigned everything with the user community, making necessary changes on a rapid, recurrent basis in response to suggestions and criticisms.

The experimental design for Columbus tried to account for the diverse technicians who serviced the target machine—whether they were dedicated to repairs of only this machine, whether they worked in rural or urban areas (city technicians drive less and take more calls per month), and how much experience they had in photocopier repair. We chose 40 technicians to participate in the experiment and gave them laptops and approximately three hours of training in the software. We chose another 40 as a control group, who were matched closely to the first group. We tracked all service calls made by both groups using the standard Xerox metrics, including cost of parts, service time, number of unscheduled maintenance calls, interrupted calls, and callbacks.

During the test, technicians responded positively. For example, those technicians not in the experimental group would borrow laptops to help them with difficult problems. Although this was encouraging, the metrics after two months were startling. The experimental group had an approximately 10% lower parts cost and 10% lower average service time than did the control group, without differing significantly in the other service metrics.

However, the test did not last long enough to convince Xerox's Worldwide Customer Services (WCS)—responsible for service strategy and technology throughout the corporation—of the value of field knowledge and the need to invest resources in a Columbus-like system for the entire service force. The results convinced our team at PARC, however, and Xerox France service management. One field engineer commented, "This is the first time people have truly paid attention to the field, to our knowledge." We decided to search for a way to extend the use of the knowledge base to all French technicians.



© Jonathan Liffgens

### *French Minitel Eureka*

To offer the technician-invented solutions to the entire French service force, we faced two problems: a method of distribution that would support technicians' work practice and a social process by which the database would have continuing value. We worked with the technicians to understand how to promote their ongoing participation and ensure continual updating of the knowledge base.

We could not continue to use laptops for distribution because funding was limited and, at the time, laptops were unavailable in France. In addition, in 1994, communication via phone lines or the Internet was too expensive. A printed booklet of tips was deemed ineffective: it would make existing information available but would not be an ongoing, growing resource.

We chose the French Minitel system for distribution, which Xerox France technicians already used for call management. Minitel, nationally deployed by the French telephone company, consisted of a small keyboard connected to the phone line and to a local display monitor (initially a television). Minitel was a general service with easy connectivity to private databases for commercial use.

We worked with the CSEs, the tigers, and the technical support hot-line specialists to figure out how to encourage contributions to the tips database, without seeming to threaten people's jobs. The hot-line specialists, who could have seen our effort as an attempt to cut positions in their organization, instead saw it as a way to potentially ease their workload. They could then spend more time thinking about common issues and generating their own tips. The tigers could have viewed it as "stealing" their knowledge, but they felt there were so many new problems that it would be advantageous to quickly disseminate new solutions. 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 problem:

1. **Quality.** To ensure quality, a validator known for expertise on the particular product line warrants each tip. At Xerox France, the validators are product specialists for each family of products in every district or "customer business unit" (CBU). The tigers oversee the process.
2. **Bottlenecks.** When a new tip is submitted on the Minitel, a message goes to the relevant group of validators, one of whom picks up the new tip within a few days. The validator converses with the submitter to ensure that the tip both captures the appropriate information and is written clearly. The CSE can edit and improve the tip, learning in the process.

3. **Incentives.** When we asked community members if they thought management should pay for each tip submitted, they said no. One tiger said, “This would make us focus on counting the number of tips created, rather than on improving the quality of the database.” The suggestion was to include the submitter’s name on each tip to act as a positive reinforcement for good tips and a negative one for badly flawed ideas.
4. **Integration with work practice.** Because the tip database was on Minitel, we added new information pages to the call handling to allow CSEs, when they faced a difficult call, to search the database for key symptoms taken from the call record. They could search the database from a customer’s site if they could have access to the local Minitel.

### Implementation and Deployment

Because implementing this system was not in WCS’s plans, and money for field service was limited, WCS declined to finance the countrywide experiment. A partnership of PARC and Xerox France paid for the system. The software was ready in about four months. A champion from the French tiger group and one of us talked with each group about service problems and how the CSEs could use the system. They met with more than 60 product leaders and helped train 1,300 French technicians. We carefully tracked participation by how many times the technicians referred to the database and how many new tips they entered. There were strong differences among workgroups. While one region might have high usage, another of the same size might have low usage rates. By revisiting the latter regions and training and reintroducing the purpose of the system, we encouraged broader participation. The strategy, then, can best be described as “hands-on, participatory implementation,” a marked contrast to a top-down, cascade model.

### Experience with Use

The Minitel system began with databases for only three products. By the end of the first year, CSEs had opened more than 40 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. Participation was extraordinary; more than 20% of the CSEs had submitted a validated tip, and CSEs were consulting the tip database an average of two or more times a week.

What did the technicians get from these tip documents? What did they consider important to share? The tips included some crucial diagnostic information, but also much more varied content. For example:

*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 white-out on the back wall near M.

Xerox France, compared to the rest of Europe, went from being an average or below average performer in service to a benchmark performer. The French service metrics were soon better than the European average by 5% to 20%, depending on the product. On a more qualitative basis, we have seen many different ways in which Eureka has affected the service process in France. In preparing for a call, technicians have found Eureka helpful in ensuring that they pick up a part likely to be causing failure before going to the customer site. On site, Eureka accelerates and improves diagnoses. It also reduces the number of calls that have to go to the next level, reducing the load on the technical support hot-line for recurrent calls about the same problems. It also significantly reduces the learning curve for new-product introduction.

### Spreading Eureka to Canada

In June 1996, we decided to bring Eureka to another community and to intersect directly with laptop introduction (only France had a system like Minitel to use as an alternative).

A senior manager who wanted to ensure the success of a new advanced color copier encouraged us to work in Canada. At that time, some 6,000 laptops had been deployed to Xerox CSEs, including all of the 1,200 in the Canadian service force—comparable in size to the French. We teamed up with a tiger from the Dorval Technical Support Center, near Montreal, who became a local champion for its development and deployment. The challenge was to adapt what we learned in France to the Canadian service environment. We initially confronted some critical, nontechnical issues:

- Those who had the laptops did not use them. Although Xerox was committed to using the computer to dispense technical information and manage work processes in the field, technicians depended on their traditional skills and practices and were skeptical about the new technology.
- CD-ROMs or floppies were used to distribute information to the laptop, so dissemination 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 leverage these independent applications.

We couldn't directly solve the laptop acceptance problem, but we hoped that Eureka would prove to be so valuable that technicians would want to use the computers. Our local champion from Dorval took existing technical information databases that had been distributed in paper form and converted them to the Eureka tip format. The technicians already valued this information, but it was hard to use or even track in paper form.

To address the distribution problem, we built a local client system that afforded rapid access to tips. This laptop client would be able to update the local knowledge base any time the technician was able to dial in to a central server. We still had to decide what kind of communication and server to use. 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 regularly to discuss problems and share ideas. This familiarity could work to our advantage as a platform for knowledge sharing in Eureka.

For accessing the knowledge base on the laptop, the technicians' work practices dictated that our search engine had to be extremely fast and 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 from a technical support group and now met all these requirements. Its integration with service applications on the laptop would have been both useful and technically possible. However, a central organization distributed and maintained the laptop software and documentation. Eureka was just an experiment operating on the periphery and had to remain a separate application. This peripheral status also meant that our Canadian champion had to perform his main job as a tiger, while simultaneously solving problems from technicians. Validators were also volunteers who were not relieved of their ordinary duties.

We had to adapt the 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 was essential.
- The French had rejected any financial incentives for authoring tips. In Canada, however, there was an existing financial incentive program 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 when they searched the knowledge base. The Canadian process would require technicians to explicitly download the latest information to update their database. How often would they want or need to do this to make the system effective?

Because the organizational structures were similar in Canada and France, with product specialists in each CBU, it was natural to make the Canadian product specialists the validators, just as in France. However, Canadian service management did not want to

give up the financial incentive program that they believed contributed 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 only for validated tips, rather than for all submitted tips.

Updating the laptop knowledge base proved to be a problem. Not all technicians used the BBS, and many found the process cumbersome. Moreover, the fact that Eureka was a separate application from call management created further complications and obstacles to frequent, easy use. As a result, when we checked after two months, many technicians—roughly 40%—rarely or never updated their knowledge base. To try to improve the situation, the Canadian champion visited each CBU to encourage updating and provide additional training.

Upgrading the software when we made changes was even more complex. We distributed floppy discs to everyone in the field and hoped that they were able to use them in a timely manner. This created so many problems that we eventually put a system in place for downloading the software components from the BBS.

Eureka was now an official, management-sponsored program, with certain expectations for improving service performance and with some financial support from a Xerox business division. Management had never dealt with a program in which the requirements emerged from experiments with pilot users, iterated until the users felt the program warranted large-scale deployment. Managers would try to set deadlines for us to get things done, independent of our process for rapid prototyping and debugging with extensive community involvement. The clash of these two different design and deployment methods had negative results. Some higher level managers lost some faith in the ability of the Eureka team to deliver.

Despite these conflicts, we successfully launched Eureka for 20 products in only six months, beginning in early 1997. The Canadian champion extensively trained product specialists, and the specialists then trained CSEs. We created a training video distributed on CD-ROM, reducing the need for more direct training. After six months, the Canadian Eureka really took hold and became the technicians' tool.

### *Eureka Moves to the US*

While Eureka had proved successful in less populous countries such as France and Canada, it was not clear how it would work in the US where 10,000 technicians are spread out over a huge area. More important, the dynamics are quite different in the US organization, which is much more bureaucratic and hierarchical, because of its size and complexity.

The US and Canadian technicians shared a common laptop/BBS infrastructure, so the only issue was to adapt the process to a differently shaped organization and to the scale of the US service force. US service management decided that validation would take place locally, with local groups selecting a validator for each product family. As in Canada, the validators would need to take on the task without reducing the rest of their workload.

Eureka was launched in the US in 1997 with pilot programs in several locations. The pilot took hold, however, only where there were local champions in the service force, as was the case in both France and Canada. Beginning in June 1998, Xerox Worldwide Customer Service (WCS) distributed Eureka CD-ROMs to the field managers, who were then expected to distribute them to technicians in their work groups. The CD included a computer-based training module; no hands-on training or direct engagement with technicians around the program was planned. This cascade strategy had been designed for mass distribution of software or documentation, but it was less effective with a socio-technical system like Eureka. In places where people became champions or where we engaged the local group, it was quickly adopted. In other places, it became just one of a dozen company-distributed programs that somehow had to be implemented over the next quarter, and adoption was correspondingly slow.

We had originally suggested to WCS management an alternative “participatory deployment” strategy in which the pilot champions, technicians, and managers most knowledgeable about Eureka would go to other locations in the US service community and talk about their experiences and ideas. Because these people were peers, the technicians would

trust them. This would have created more local champions and knowledgeable users, who could then have gone to still more locations to share information. During a relatively short time, Eureka would have spread across the entire country.

The up-front cost in time and travel for participatory deployment would have been greater than for the cascade distribution. But we believed that this cost would have been recouped because more technicians would have used the program quickly, resulting in a shorter learning curve and better performance. The results from France and Canada support this argument. But WCS management in the US did not understand the requirements of combining the social with the technical and did not approve this plan, so Eureka use spread slowly in the US.

Nevertheless, US technicians, once they learned about Eureka, were enthusiastic. One technician remarked, "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 US before moving to any other Xerox organizations in Europe, Latin America, or Asia, demand from technicians in these countries was so intense that the corporation had to begin distributing Eureka worldwide.

In 1999, US technicians authored approximately 2,000 tips. There were more than 9,000 "solves" using Eureka in the US and Canada in the fourth quarter of 1998 alone. The knowledge base for these problem resolutions included more than 30,000 records. 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 more than 300,000 records.

## Eureka in Practice

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?

After Eureka had been in the field in the US for six months, a member of our research group talked to technicians in San Francisco, concentrating on a particular work group. He and other members of our group also visited four CBUs around the country. We asked technicians if Eureka was worth using, and if so, how they used it and how we could make it better. When they learned that we had designed and launched Eureka, they made remarks such as:

Best reason for having a laptop. I use it on probably 50% of the calls where I don't walk in the door and immediately go, "Well, this sensor's broken," or something like that. Anytime something doesn't immediately jump out at me, it's the first thing I turn to. Most of the time before I get to a site, I look around in Eureka and see what's there so that I know what I'm gonna do.

When first designed, Eureka was conceived primarily as a tool to use when routine fixes fail to resolve a problem and past experience doesn't point to an answer or line of attack. Many technicians use Eureka only that way, whether with machines they are working on or for suggestions to give colleagues who ask for help. But some technicians use Eureka in other interesting ways. They use it as a tool of *first* rather than last resort. For example, one technician who works on high-volume copiers uses Eureka in combination with product documentation:

Before I go on a call, I like to look at some possible fixes in Eureka. If I feel that there isn't anything in Eureka that jogs my memory, then I go to the documentation. Keeping that footprint of some of the fixes and then just going through the repair procedures in the documentation accelerates things.

Thus, even before seeing the machine, this technician tries to develop several solid leads about the source of the problem, the likely repair procedures necessary, and needed parts. Another technician reported similar patterns in using Eureka and the documentation:

Eureka isn't so much an end, as a beginning. Someone will call over the radio with a fault code like, "I'm having 12-142s," and I can look it up in Eureka and scroll through common causes. It's faster to find it in Eureka than it is to go in and fire up the documentation CD for the repair procedures there.

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

Technicians also use Eureka as an informal *learning* tool. One who services mid-volume machines browses through the tips to see what has worked for others: "Whenever I download new Eureka data, I like to see what guys are doing. I look through the tips and bulletins. It teaches me a lot." By reading the tips and service bulletins somewhat casually, divorced from an actual repair situation, this technician uses Eureka as an instructor who offers a new set of lessons each week.

We also identified some barriers to effective use. Laptops have a long boot-up time, limited battery life, and an unstable operating system, and they complicate the updating of the database. In addition, many technicians simply mistrusted, were unfamiliar with, and resented computers, so didn't use the laptop except when absolutely necessary. Other technicians felt the laptop added time and work to their daily routine. One technician remarked:

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] and that's about it. The real problem is getting them to adopt the laptop generally, not Eureka.

Another barrier was that technicians had to do independent searches in the Eureka knowledge base and the documentation, which required them to enter information several times. The integration of all the tools and databases was the biggest request from our feedback meetings and was a primary design criterion for the next-generation laptop. Technicians wanted to move more easily between tasks.

Technicians not only used Eureka in creative ways, but regularly thought about making it more effective. This is exactly the kind of inventiveness that Eureka was meant to capture, and it stands as further evidence of the pervasive importance of working with users to make a system fit their needs—to artfully integrate the technology with their enhanced practice.

## Eureka II

The advent of cost-effective communication on the Internet allowed us to implement a new web-based Eureka—Eureka II—worldwide. To bring together everything that technicians need to do when connected, Xerox 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 downloads any updates to the knowledge base. This same mechanism updates the documentation 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 technicians have a problem, they can see where they may find helpful information in their "hit list" references to tips and to multiple places in the formal documentation. In addition, they can make annotations on already existing documentation, keeping such "post-its" in a private knowledge base or, if desired, submitting them as tips. 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.

Eureka II was so successful that it became a mainline program, and requirements poured in from many places. We constantly tried to balance our belief in simplicity with corporate managers' beliefs that if Eureka were the answer, they wanted to generate the question. For example, one manager felt that a big cost of the system was in training technicians. He wanted to simplify the training by embedding the Eureka application in a standard Internet browser (in this case, Internet Explorer). We thought this would complicate the implementation significantly because the software would then be dependent on each computer's version of the browser, operating system, and service packs. The manager felt this was less important than the simplicity of the training. Unfortunately, there were far more implementation and deployment complications than even we had imagined, and the delay in deploying Eureka II was significant. Our point here is not that the manager was wrong, but rather that decisions made in a standard software-development process contrasted with the bottom-up approach with which we had started.

In addition, because the project now had top management's attention, we sometimes had to set schedules based on managers' desires for certain goals, rather than on the necessary work to achieve a final state. Although we understood the pressures on the managers, their schedules often could not be realized, leading to internal battles and slipped schedules. We became very aware of the difference between singing in the spotlight and singing in the shower.

## Organizational Barriers to Change

After seven years, the Eureka story is a tale of how the development and deployment of a system for sharing knowledge from the front lines became a vehicle for organizational change. However, our story also reveals that this change was not without conflict and challenges. These messy details are rarely included in writings on knowledge management or organizational learning, so it is worthwhile to expand on their larger meaning for knowledge sharing.

In the initial stages of the project in France, few people in Xerox management believed that there was much value in what the technicians learned on their own in the field. In addition, they could not see how a tip system was much different from previous suggestion systems, all of which were highly centralized and controlled. And although technicians' tips quickly proved valuable, people in different parts of the company felt that it was more important to supply the technicians with centrally produced documentation than to support them in creating new knowledge. This different way of doing business made them nervous, for example, when a single flawed tip eventually slipped through.

Getting support for the project in the form of organizational resources 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 do the job. Moreover, in the initial stages, we sometimes had to operate like a guerrilla group because opposition was enough to kill the project if we openly challenged deeply entrenched convictions. We conducted our first experiment in France partly because it was out of sight of the central Xerox organization. After the French experiment, only by convincing one product manager in a business division to give us the funds for the Canadian experiment were we able to gather data that would convince the nonbelievers. (Later, WCS awarded one of our research group members, who had led the French effort, a plaque that read, "Despite the resistance of Worldwide Customer Service.")

We recognized at the start that the service organization would not accept informal responses to the collection of tips or the users' informal assessments. We knew we had to show hard bottom-line data. In some ways, resistance to making a wide-scale change like Eureka in just one step allowed us to gather better data. For example, we could put out a product with only minimal diagnostic documentation and then use the field force to help us understand where and how it needed improvement, that is, to construct the diagnostic documentation in the field.

Once Eureka became a major corporate program, the project ran into a different sort of problem, perhaps resulting from its success. Why was this a problem? In France and Canada, because we had been conducting guerrilla experiments, we could involve the users in the decisions on adapting Eureka to local needs and practices. Moving to the central WCS organization, however, engendered a change of philosophy, and WCS mandated a uniform, worldwide solution. Management had a policy of distributing corporate programs in a cascading fashion. Eureka works better when peers mentor each other on the uses of the system. Unfortunately, the rollout in the US was not done this way. While the US deployment was eventually successful, the problems that developed because technicians weren't closely involved in the process hindered the project's achievements.

WCS and Xerox Corporation emphasize cost savings in field service and view the service organization as a cost center. As a result, they discourage significant investment in service, unless it is matched with equivalent cost reductions. This had consequences for Eureka deployment and how the program now operates. Although the president of Xerox saw Eureka as 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. Xerox expected technicians to author tips and

validators to provide rapid turnaround and validation of submitted tips without any relief from their current workloads.

Did Xerox become a better learning organization as a result of the Eureka project? One answer would be the impact on field service performance and the degree to which most technicians use the program regularly, especially for learning new ideas and approaches to machine repair. The whole service organization has also been transformed to some degree by the bottom-up Eureka approach, which has had an impact on operational philosophy.

Despite these significant achievements, the corporation has not yet taken full advantage of the possibilities of a knowledge-sharing program like Eureka. For example, although many organizations use the Eureka knowledge bases informally, no formal process incorporates Eureka's information back into the documentation. Engineering could, but doesn't regularly, mine the Eureka knowledge base for ideas on continuous improvement. And manufacturing could use Eureka to augment the information flow needed to adjust rapidly at the initial launch of a new product.

The current Eureka process, which is dedicated to technicians authoring tips for fellow technicians, obviously cannot address all these areas or solve the interorganizational knowledge-transfer problem. It does point the way, however, to the need for collecting knowledge on the corporation's frontlines for use throughout the corporation.

At the same time, the "spirit" of Eureka has had some interesting effects on Xerox as a whole. We have many requests to help create a "Eureka-like" knowledge system for other operational units. Currently, we have deployed LinkLite, a simpler infrastructure, to support the Eureka process in a Xerox sales organization. Salespeople share knowledge about "customer solutions" (special configurations of machines and services that help customers solve important business problems), their successes, and other sales material. Perhaps the belief in and spread of this "spirit" has become the most important legacy of Eureka at Xerox.

## Acknowledgments

We would like to thank Olivier Raiman, the original Eureka champion in our group, and Bob Cheslow, the principal architect and implementer of Eureka. We are also grateful to members of the extended team: David Bell, Mark Shirley, Eric Delanchy, Michel Boucher, and the many technicians worldwide who graciously shared their experiences. We also appreciate the support of PARC management, particularly John Seely Brown and Johan de Kleer, and the people in Worldwide Customer Service, Bernie Colligan and Tom Ruddy. To others who have helped us along the way but whom we have not named explicitly, we thank you.

## References

- Ambrose, S.E. *Citizen Soldiers: The U.S. Army From the Normandy Beaches to the Bulge to the Surrender of Germany, June 7, 1944–May 7, 1945* (New York: Simon and Schuster, 1997).
- Bell, D.G., D.G. Bobrow, B. Falkenhainer, M. Fromherz, V. Saraswat, and M. Shirley. "RAPPER: The Copier Modeling Project" (International Logic Programming Symposium, San Diego, CA, October 1991).
- Bell, D.G., D.G. Bobrow, O. Raiman, and M. Shirley. "Dynamic Documents and Situated Processes: Building on Local Knowledge in Field Service," in *Information and Process Integration in Enterprises: Rethinking Documents*, eds. T. Wakayama, S. Kannapan, C.M. Khoong, S. Navathe, and J. Yates, eds. (Norwell, MA: Kluwer Academic Publishers, 1997): chapter 17.
- Brown, J.S. and P. Duguid. "Organizational Learning and Communities-of-Practice: Toward a Unified View of Working, Learning, and Innovation." *Organization Science* 2 (1991): 40–57.
- Brown, J.S. and P. Duguid. *The Social Life of Information* (Boston: Harvard Business School Press, 2000).
- Davenport, T.H. and L. Prusak. *Working Knowledge: How Organizations Manage What They Know* (Boston: Harvard Business School Press, 1997).
- de Kleer, J. and B.C. Williams. "Diagnosing Multiple Faults." *Artificial Intelligence* 32 (1987): 97–130.
- Doubler, M.D. *Closing with the Enemy: How GIs Fought the War in Europe, 1944–1945* (Lawrence, KS: University of Kansas Press, 1994).

- Nonaka, I. and H. Takeuchi. *The Knowledge-Creating Company: How Japanese Companies Create the Dynamics of Innovation* (New York: Oxford University Press, 1995).
- Orr, J.E. *Talking About Machines: An Ethnography of a Modern Job* (Ithaca, NY: Cornell University Press, 1996).
- Senge, P., A. Kleiner, C. Roberts, R. Ross, G. Roth, and B. Smith. *The Dance of Change: The Challenge of Sustaining Momentum in Learning Organizations* (New York: Doubleday, 1999).
- Stewart, T.A. *Intellectual Capital: The New Wealth of Organizations* (New York: Doubleday, 1997).

## Commentary

by Marleen Huysman

There probably is no living organizational learning researcher who hasn't read or at least heard of the article on communities of practice by John Seely Brown and Paul Duguid (1991). They refer to Julian Orr's study in which he analyzed the day-to-day learning, working, innovating activities of photocopier repairmen at Ranx Xerox. The story is so inspiring because it offers a refreshing look at normal daily activities: social learning in communities of practice.

Since this seminal article was published, I hadn't heard much about the Xerox repair people, which I had begun to interpret negatively. I was thus happily surprised to read Daniel Bobrow and Jack Whalen's article. The authors tell yet another story of the lives of Xerox copier and printer repair people. What was missing in Brown and Duguid's article, namely, the consequences of IT support, is now explicitly addressed. The authors give an excellent account of the life of a socio-technical system used by the reps to support their learning and working processes. Because of the detail and time frame covered, the story offers a multidimensional collection of do's and don'ts for introducing sociotechnical tools for knowledge sharing. For example, the authors illustrate the necessity of a bottom-up approach and the use of local champions. They show us how cultural differences can completely alter previous lessons learned and why top management approval has its downsides. They tell us about the process of appropriation and the new meanings attached to the tool the moment it is used in various local practices.

In addition to these and other requirements of sociotechnical systems, the story gives a nice example of the conditions for successful computer-supported knowledge sharing. Especially interesting is the possibility that Eureka offers in authoring, submitting, and validating tips and the absence of a need (at least in France) to financially reward tips. I agree (again) with Brown who sees this authoring and tipping process as contributing to the social capital of the organization: "The author attaches his or her name to the resulting story or tip, thus creating both intellectual and social capital, the latter because tech reps who create really great stories become local heroes and hence more central members of their community of practice" (Brown, 2000: 17).

The story provides insight into how to challenge the general lessons from the first generation of knowledge management, which often fail because the technology doesn't match the informal bottom-up collective knowledge-sharing needs and practices. The story also shows that there is no need to fall into the opposite trap of approaching all knowledge-sharing technologies as negative. In fact, when we are able to create a sociotechnical match between technology and the opportunities, abilities, and motivations (or degree of social capital) to share knowledge within communities, there are many possibilities. With this, the authors provide a valuable example of Boland and Tenkasi's (1995) claim: "Information systems aimed at knowledge management need to maintain the integrity of the social communities in which knowledge is embedded."

## References

- Boland, R.J. and R.V. Tenkasi. "Perspective Making and Perspective Taking in Communities of Knowledge." *Organization Science* 6-4 (1995): 350-372.
- Brown, J.S. "Growing up Digital, the Web and the Learning Ecology." *Change, The Magazine of Higher Learning* 32 (March/April 2000): 11-22.
- Brown, J.S. and P. Duguid. "Organizational Learning and Communities of Practice: Toward a Unified View of Working, Learning, and Innovation." *Organization Science* 2/1 (1991): 40-57.



Marleen Huysman  
Associate Professor  
Vrije University,  
Amsterdam, The Netherlands  
mhuysman@feweb.vu.nl