AI Magazine Volume 21 Number 4 (2000) (© AAAI)

The Road Ahead for Knowledge Management

An AI Perspective

Reid G. Smith and Adam Farquhar

■ Enabling organizations to capture, share, and apply the collective experience and know-how of their people is seen as fundamental to competing in the knowledge economy. As a result, there has been a wave of enthusiasm and activity centered on knowledge management. To make progress in this area, issues of technology, process, people, and content must be addressed.

In this article, we develop a road map for knowledge management. It begins with an assessment of the current state of the practice, using examples drawn from our experience at Schlumberger. It then sketches the possible evolution of technology and practice over a 10-year period. Along the way, we highlight ways in which AI technology, present and future, can be applied in knowledge management systems.

Although there are many different definitions of knowledge management, we can take the following as a representative statement of the primary goal.

> Improve organizational performance by enabling individuals to capture, share, and *apply* their collective knowledge to make optimal decisions … in real time.

By *real time*, we mean the time available to make a decision—to take action that will materially affect the outcome.

In 1997, The Delphi Group, Inc., reported that although only 28 percent of companies surveyed were investing in knowledge management at that time, 50 percent expected to be investing within the next year, 77 percent within 1 to 2 years, 93 percent within 2 to 4 years, and 98 percent after more than 4 years. Three years later, according to studies by KPMG and the Conference Board, 80 percent of the world's biggest companies have knowledge management efforts under way. The predictions of the earlier survey have been borne out by action.

To date, most companies have embarked on knowledge management work in search of near-term efficiency, productivity, and service quality improvements through knowledge reuse. However, they also envisage longer-term benefits, including personnel motivation, faster learning, and increased innovation. Substantial bottom-line impact has been reported by various companies, including Texas Instruments ($1.5 billion over 3 years), Chevron ($2 billion annually), and BP ($30 million in the first year) (O'Dell et al. 2000; Payne and Elliott 1997).

Knowledge Management Issues

Practitioners and business managers alike agree that issues of technology, process, people, and content must be addressed to achieve success.

An organization must have "good enough" technology to make progress, especially in the transnational business environment of today. However, to achieve bottom-line success, even more attention must be paid to the other issues. Experience in many organizations has shown that no more than one third of the knowledge management budget should be devoted to technology (O'Dell, Elliott, and Hubert 2000).

The basic organizational unit of knowledge management is the *community of practice* (Brown and Gray 1995), which is a group of people who share a common area of expertise and/or who search for solutions to common problems. Nurturing vital communities is hard enough when the members are in a single location with good connectivity. When the members are spread around the world, in remote

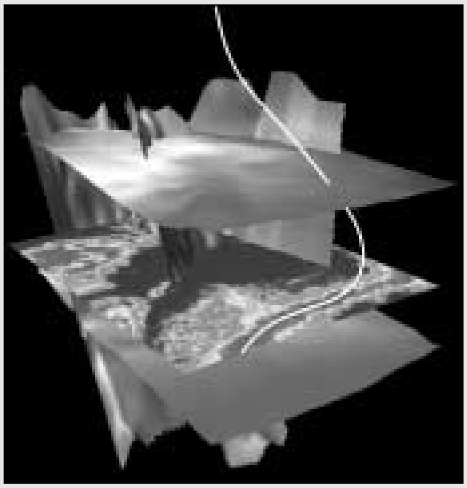

Efficiency and Productivity

Vignette 1:

Efficiency and Productivity

Success Statement: The organization knows what it knows and uses it and knows what it needs to know and learns it.

Conversation between two research and development managers:

André: In Paris, we just finished development of the new service for imaging fluid flow in horizontal wells. The insight of that new engineer who works for you in Houston was invaluable.

Martha: Glad to hear it. He saw on the intranet that your team was in need of assistance with computational fluid dynamics—and he has a Ph.D. in that area.

…or… He told me your team found out about his expertise via the intranet.

If this sort of conversation were to be an everyday event, it would clearly indicate a company that is able to connect the people with the knowledge to the people who need it and bring that knowledge to bear on important problems. It would also indicate a company that is able to make its accumulated experience available to new staff, thus helping them become effective contributors soon after joining.

areas as well as population centers, the challenges are enormous. Therefore, significant attention must be paid to issues of process and people.

Communities must have processes in place that enable them to capture, share, and apply what they know in a coherent fashion across the organization.

According to a 1997 Ernst & Young Center for Business Innovation 1997 survey entitled "Executive Perspectives on Knowledge in the Organization," the biggest impediment to knowledge transfer is corporate culture (54 percent), and the biggest difficulty in managing knowledge is changing people's behavior (56 percent). (See also Bock [1998].) Organizations must find ways to motivate individual community members to share what they know and to apply the knowledge of others.

Finally, validated and trusted content is vital. Organizations must be concerned with putting in place a work environment that encourages and supports both maintaining the current content and replenishing the stocks over time.

From this 1997 Ernst & Young survey, business managers indicated that the most important types of knowledge to have include knowledge about customers (97 percent), knowledge of best practices–effective processes (87 percent), and knowledge about the competencies and capabilities of their company (86 percent).

For the purposes of this article, we can define *information* as "interpreted data" and *knowledge* as "information in action" or "information transformed into capability for effective action." However, we dwell no further on distinctions between the two. Rather, we take the point of view that such distinctions do not offer much useful guidance, at least through the first few stages of knowledge management implementation in an organization. See Davenport (1996) and O'Dell, Elliott, and Hubert (2000).

What Constitutes Successful Knowledge Management?

By way of making the goals of knowledge management more concrete, we present four examples of what it might mean to succeed. Each is drawn from our experience in the oil and gas business at Schlumberger. (We learned this technique from Stephen Denning [1999], who has used storytelling to advantage at the World Bank.)

The Knowledge-Powered Enterprise

If we can make the vignettes "normal" in an organization—make them such every day occurrences that they would not merit special attention—we argue that a new era in organizational

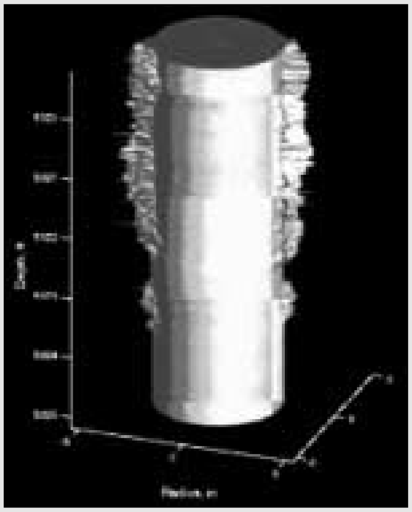

Vignette 2: Rapid Dissemination and Implementation of Innovative Ideas

Apply Everywhere What Is Discovered Anywhere

Success Statement: For any project, for any customer, the project team delivers the knowledge of the overall organization.

Conversation with a customer:

> *Well Engineer:* Yesterday, one of our teams was working with your company on well N21, offshore Nigeria. They came up with a new approach to solving wellbore stability problems. Our scientists in Cambridge have run some tests and recommend that we use it on well G56 that you're drilling in the Gulf of Mexico.

A wellbore stability problem encountered in an oil or gas well while it is being drilled can result in a "stuck pipe incident." This is analogous to a hand drill becoming stuck in a piece of wood, except that with oil field drill pipe, it can have much more serious and expensive consequences.

This vignette indicates a company whose people are aware of what their colleagues are doing, even on the other side of the world. Furthermore, it indicates a company whose people have access to the latest developments—they are operating in real time. They are also able to draw on the research and development resources to understand how to bring to bear the knowledge and deliver it to a customer in a timely fashion.

To put this vignette in perspective, consider an alternate "failure" version: The customer in the Gulf of Mexico learned about a wellbore stability problem encountered in Nigeria yesterday. He calls the well engineer to determine how to prevent the problem from happening to him. Unfortunately, the well engineer is not aware of what happened in Nigeria. Such an outcome reduces the confidence of the customer in the well engineer and the service company. It also affects the well engineer—no one wants to be embarrassed in front of a customer.

performance will have been reached. We call it the era of the "knowledge-powered enterprise."

In this new era, knowledge management happens in the background—in real time. It is done by everyone as part of the day-to-day job, embedded in the workflow. People are easily able to obtain the data, information, and knowledge they need to do their jobs. They interact effectively with their colleagues—anywhere, any time, and by any means. They are supported by a new generation of business and technical applications that contain integrated decision support, data-mining, and simulation technology.

In addition, the enterprise fosters knowledge creation and innovation by continuous learning to replenish and renew its stocks of knowledge. Not only has the enterprise achieved substantial improvements in efficiency, productivity, and service quality, it has reinvented itself as a provider of products and services that are only possible because it is able to leverage the collective knowledge of its people.

Why Now? What's New?

Why has business interest in knowledge management picked up dramatically over the past few years? After all, organizations have been sharing information and knowledge for many, many years.

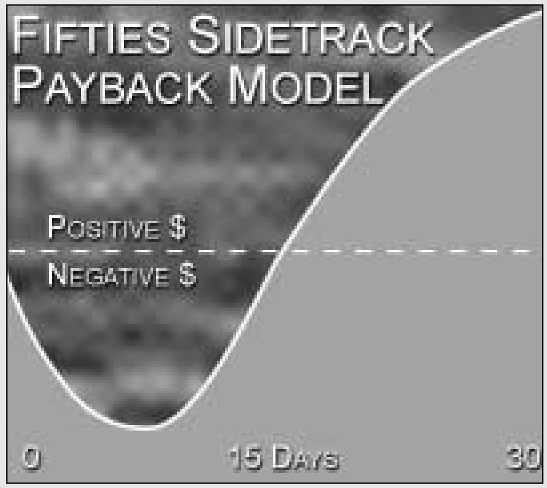

Integrated Information, Simulation, and Decision Support (from Baird [1997]).

Vignette 3:
The Right Information, to the Right People, at the Right Time

Success Statement: The organization delivers the right information, to the right people, at the right time—with the tools they need to use it!

Conversation with a potential investor:

OilCo Well Engineer: We're planning to sidetrack well B23 in the Fifties field and are hoping you'll buy in. For $10 million you can become a full 1-percent partner. A 96-hour shut-in test ended a few minutes ago. It indicates additional reserves of 750,000 barrels of oil and an initial sand-free production rate of 4,300 barrels per day.... We could do this tomorrow.

Investor (consulting his knowledge hub): I see that a sidetrack can double my rate of return, with a payback time of 30 days for 1-percent interest.... Count me in.

It's not enough to have the right information delivered to the right people at the right time. They must also have the tools to apply that information. In this vignette, the investor has the simulation and decision support tools necessary to integrate the geoscience, production, and economic data delivered by the well engineer. He/she is able to determine the potential payback and risk and make a decision—in real time.

Consider the following points:

First is the collaboration space; it is now more virtual than physical. Companies are increasingly distributed worldwide. As a result, it is more difficult to collaborate with one's peers. It's hard enough when the team works in a single location. It is much harder when it is spread out across the world.

Second is intellectual capital. In the 1970s, the ability to set up a world-class manufacturing environment was a key differentiator. In the 1980s, it was quality; in the 1990s, it was customer service and support. Today, the ability to learn quickly and continuously and to operate inside the learning circle of the competition is a key differentiator. Knowledge is seen as supplanting other physical assets as the most important competitive resource. It is argued that most of the valuation of a company is based on its intangible assets, including its intellectual capital. In addition, for many companies, knowledge is their product as opposed to physical assets or actions. See Stewart (1997).

Third is information technology. This is perhaps the main driver. As a result of information technology, it is now possible to do something

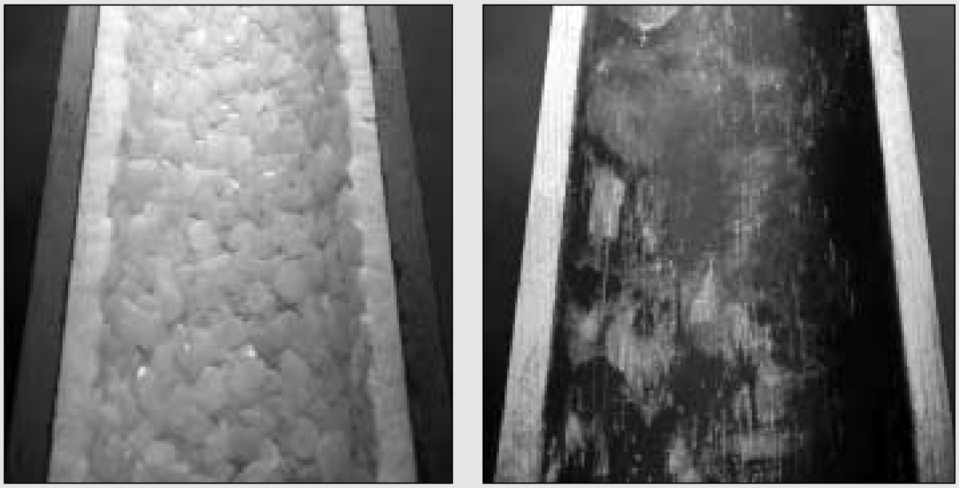

On the Same Page as the Customer.

Vignette 4:

On the Same Page as the Customer

Success Statement: The perspective of the employees is aligned with that of the customers.

Conversation with a customer:

Customer: My well B23 in the Fifties field has a serious scale problem. I'm looking at the ScaleBlaster information on your web site. From the video and case histories, it looks like this might be the answer.

From the "Service Delivery" section, I also see that you can deliver this service tomorrow.

Sales Engineer: Yes, I'm looking at the same page.

By the way, our research and development staff has been working on new applications for ScaleBlaster. I think the results might be relevant to some of your other wells. Let's talk about it when we meet tomorrow.

Knowledge management is an essential ingredient for e-business. Vignette 4 demonstrates how the world of e-business naturally leads to increased transparency between customers and suppliers.

At Schlumberger, we use the same information structure to present information to employees that we use to present information to customers—and we use the language of the customer. As a result, the employees are literally on the same page as the customers. The only difference is that the employees can typically see more information than the customers.

In this vignette, the sales engineer can see what the customer sees but in addition can see unreleased information on the current state of the research and development program, enabling the sales engineer to suggest a conversation about new applications for ScaleBlaster.

about knowledge management. Today's internet, intranet, and web technology permits practical capture, sharing, and leveraging of information and knowledge throughout organizations.

Knowledge Management and AI

Knowledge management shares a goal with the AI community—development and use of computational technology to improve the performance of individuals and communities.

Knowledge management needs AI. The ideas and technology that have been developed by the AI community are required for successful knowledge management; for example:

Knowledge base. Different words can be used, corporate memory, knowledge repository, best-practices database, and so on, but it is clear that a high-quality knowledge base is fundamental for successful knowledge management. What has been learned in the AI community about knowledge acquisition, representation, and inference can all be brought to bear for knowledge management. Indeed, this experience *must* be brought to bear to achieve the success stories presented thus far.

Search. Community members must be able to find and bring to bear the relevant knowledge in the repository. In today's knowledge management systems, the search engine is the main workhorse. Search engines apply rudimentary natural language–understanding techniques. Over time, we expect what has been learned about indexing and retrieval in case-based reasoning systems to be applied. Furthermore, agents and distributed problem-solving technology will play an increasingly important role.

Knowledge management challenges AI. The issues are largely people oriented, distributed, and messy.

Everyone is a contributor. To be successfully applied for knowledge management, AI technology must work for a broad population of people—the knowledge workers inside companies. It is important to remember that they are not computer specialists.

Knowledge is not neatly packaged. It is difficult to extract, it is ephemeral and may only be approximately correct. One of the chief barriers to the construction of expert systems is the difficulty of knowledge acquisition—and this barrier must also be overcome to achieve success in knowledge management. However, there is some good news. The problem may be more tractable in the new context. Systems that support knowledge management typically do not try to solve a problem alone. Rather, they try to find the knowledge (best practices, lessons learned, tips, solutions to related problems, and so on) that assist people in developing their own solutions. Stated another way, the goal today is to help people solve problems. Contrast this with the original goal of expert systems—to solve problems themselves at an expert level. This is not to say that the AI community should give up on the expert systems goal, but rather, that achieving the goal of giving powerful assistance to people as they solve problems is of great interest to the knowledge management community.

And now for the bad news.

Knowledge management shares the problems of hype and inflated expectations that have plagued AI. For good or ill, a comparable number of column inches have been devoted to the two subjects. Perhaps as a result, negative conclusions have already been drawn by some about the lack of early success in knowledge management. This is, of course, familiar to the AI community.

However, despite the hype, knowledge management is regarded as important by many chief executive officers, and substantial progress is being made in a number of organizations (O'Dell, Elliott, and Hubert 2000; O'Dell et al. 2000). Indeed, knowledge management has passed the "peak of inflated expectations" and the "trough of disillusionment" and is moving up the "slope of enlightenment" on the GartnerGroup maturity curve. As a result, if the AI community is able to develop something of value in this area—the "*killer app*" for knowledge management—there is an audience waiting to use it.

Communities of Practice

It is clear from our earlier goal statement that a primary focus of knowledge management work is finding effective ways to connect groups of people. Not surprising, most practitioners agree that the CoP is the fundamental organizational unit in knowledge management. Brown and Gray (1995) offer the following explanation, referring to CoPs at Xerox and National Semiconductor.

> What are CoPs?… At the simplest level, they are a small group of people (in this case, about 20) who've worked together over a period of time. Not a team, not a task force, not necessarily an authorized or identified group. People in CoPs can perform the same job (tech reps) or collaborate on a shared task (software developers) or work together on a product (engineers, marketers, and manufacturing specialists). They are peers in the execution of "real work." What holds them together is a common sense of purpose and a real need to know what each other knows. There are many communities of practice within a single company, and most people belong to more than one of them.

CoPs are known by many names in different organizations: communities of interest, knowledge communities, technical communities, knowledge ecologies, professional networks, best-practice networks, and so on.

There is by now a wealth of documented industry experience on approaches to organizing, operating, and nurturing CoPs (Dixon 2000). Consulting companies have been among the leaders, including Hewlett-Packard Consulting, Arthur Andersen, Andersen Consulting, and Ernst & Young. Oil and gas companies have been active as well, including BP, Caltex, Chevron, Conoco, Marathon, Mobil, PDVSA, Shell, Statoil, and TOTALFINAELF. Companies such as Intel, Lucent, Siemens, and

Xerox have also worked to advantage in this area.

Some communities have formal objectives and full-time support staff (for example, Mobil's downstream best-practice networks [Hauswald 1999], the Schlumberger INTOUCH program). Others are self-governing (for example, the Schlumberger Eureka technical communities).

Example

In this subsection, we illustrate CoP operation by drawing from our experience at Schlumberger. Unlike the vignettes given earlier, everything in this example is in use today.

Figure 1 depicts the Schlumberger intranet and knowledge hub support in place today for each community of practice, composed of people in field operations, research and development centers, marketing, personnel, finance, and so on.

The story begins in the upper left corner of the figure. Community members are engaged in what we call *field activities*—working directly with customers—often on remote seismic vessels or drilling rigs to find and develop oil and gas reservoirs.

While engaged in these activities, they carry on a dialog with their colleagues—the members of their CoP. There are many CoPs in the company, with foci such as geophysics, reservoir characterization, well engineering, and well stimulation.

The engineers ask each other questions, pose problems and request help, exchange ideas, and debate solutions to problems. The culture of these communities is such that the members feel a responsibility to respond. It is a normal part of the way they work.

The dialog is carried out over the company intranet, depicted by the large oval. We cannot understate the importance of the investment that the company has made over the years in this technology. The Schlumberger SINET is today one of the largest intranets in the world and is essential to the business. The rest of what we have to say about CoP operation must be seen in this context. We assume worldwide connectivity and a culture that encourages its use.

E-mail: This is the tool of choice to support CoP interaction. There are three reasons why this is true. First, many of the engineers are "sometimes connected" to the intranet. On drilling rigs scattered across the world, often offshore, they do not always have the continuous high-speed connections seen in offices and research and development centers. When they do get the chance to connect to the net, they are able to download mail messages in the background, while they catch up with their local colleagues, make phone calls, and the like before rushing off again. Second, the communities are spread out across the world. The engineers know that they can pose problems and ask questions while their colleagues are asleep, then go to sleep themselves and expect to see answers when they awake. Third, although Schlumberger uses English as a common language, it is the mother tongue of no more than half of its people. E-mail is a permissive means of communication with respect to incomplete mastery of a foreign language. It masks unequal command of the language by allowing someone who reads and responds to a message to take the time to understand it and compose a reply, asking colleagues or consulting a dictionary as needed.

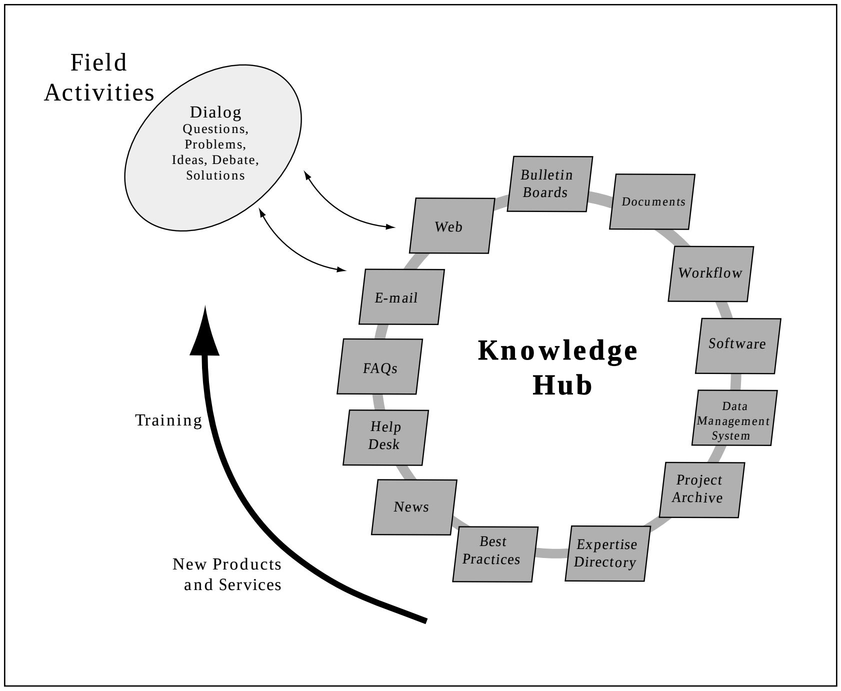

Figure 1. Communities of Practice.

Web: Although e-mail is the tool of choice to support community interaction, in the years since 1993, the web has become the primary means for finding information. As a result, one of the first things we did in knowledge management was to construct a one-stop shop to bring together information that previously had been difficult, if not impossible, to find on the intranet. We call it "the hub." It is an example of what is now called a *portal.*

Proceeding clockwise around the oval, the remaining boxes illustrate the kinds of information that CoP members are able to find with the hub.

Bulletin boards: These are analogous to usenet news groups. The e-mail messages are captured in web-searchable archives. There are several hundred of these in the company. The technology is not leading edge—it's 25 years old—but it has the advantage that it is actually used by the community members. A lesson we have learned is to focus on the technology that we know the community members actually use rather than attempt to push more exotic technology as we also try to change process and behavior.

Documents: These are the training materials, product information, technical manuals, policies, and so on, that the community members need to do their jobs.

Workflow: This is the series of steps to be followed to perform particular tasks.

Software: For engineers working in the oil and gas industry, the primary software tools are for job planning, simulation (for example, fluid dynamics, wave propagation), and diagnosis.

Data management system: With technical communities, access to the measured data—the hard data—is essential, for example, seismic data, well-log data, production data, economic data, and business data.

Project archive: The communities also have access to current and past project information.

The project archive plays three roles: First, it is an electronic work space for project teams—the clearinghouse where they capture and store the project customers, vision, road map, plans, milestones, progress reports, presentation materials, software, experimental data, successes, failures, lessons learned, and so on. Because team members can see advantages for themselves, they are motivated to contribute to it. In addition, the software is simple enough so that the effort required of members is minimal, not much more complicated than storing the information on their personal computers, which of course they are already doing. The archive improves the efficiency of project teams.

It is interesting to note that the technology was primitive and hard to use when we started in our Cambridge, England, laboratory in 1997, yet the project teams stuck with it. The need was there, and using the project archive became part of their normal workflow. Happily, the technology is much improved today.

The project archive helps to bring new team members up to speed more quickly by giving them access to the historical record, which also minimizes the number of questions that they need to ask of their colleagues. This is positive for the new members (because they are not embarrassed by the need to constantly ask "simple" questions). It is also positive for the more experienced team members (because they do not have to answer the same questions repetitively—they are only asked the more difficult and interesting questions).

The second role played by the project archive is that of real-time activity reporting. It enables others around the company to keep abreast of what is happening, what the problems are, and so on. The project archive is one of the vital elements needed to make possible the success of vignette 1.

Finally, the project archive reduces the "knowledge drain" as people move around or leave the company. It fosters reuse of information and helps avoid duplication of previous efforts.

Today the approach is applied primarily in the research laboratories. However, it is clear that it can have widespread application in the field as well.

Expertise directory: The community members can also find each other by searching the corporate directory. "Know who" is often as important as "know what" and "know how."

In Schlumberger, there is a single LDAP (lightweight directory access protocol) directory (LDAP Central 2000) for the entire company. The directory contains the basic "coordinates" of each employee—telephone number, e-mail address, physical address, business unit, the manager(s) to whom they report, photographs, pointers to their personal web pages, and so on.

It is important to note that employees are responsible for keeping their information up to date. It is not the responsibility of an Information Technology or Personnel department.

From the directory, one can generate organization charts of any part of the company. The LDAP directory is also used by many applications to provide authentication and authorization information to identify people and give them access to information (for example, their health and pension benefits).

The directory is where people record their interest in various community discussions and subscribe to them.

Finally, people can record their expertise in the directory, which helps community members find others who might be able to offer assistance. For example, they can search for people with expertise in "computational fluid dynamics." This is the sort of search that was done to make possible the success of vignette 1.

Today, Schlumberger employees are able to search for others in the company with the needed expertise in three ways: First, they can use keyword expertise search in the corporate directory. Second, they can use full-text search through the community network profile (CNP) database of community members. The CNP is a kind of resume for each community member. An example is shown in figure 2. Third, they can search the intranet for documents that mention the needed expertise.

Best practices (and lessons learned): These represent the ready-at-hand knowledge of the community—its shared common sense. They include recipes and examples that detail the community's understanding of the best way to accomplish a task or solve a problem. In the next two subsections, we discuss best practices in more detail and present an example from the Schlumberger directional drilling community.

News: This is the "buzz" of the community—the issues in which the members are interested today. News includes upcoming community events, recent successes and failures, and newly published best practices and lessons learned. For example, imagine that we learned something interesting in Nigeria today. We need to ensure that the community is alerted so that we can reuse the knowledge gained elsewhere in the world tomorrow. This kind of information dissemination underlies the success of vignette 2.

Contributors are also highlighted in the news for each community. It is one way to recognize their efforts.

Help desk: Many companies have been significantly "flattened" in recent years. In our own company, there used to be a complete hierarchy of "technique" people in the field. Among other things, they were responsible for helping engineers solve day-to-day problems. Today, that hierarchy is gone. In its place is a small distributed help desk, called INTOUCH, staffed by experienced field engineers based in the technology centers of the company. The INTOUCH engineers respond to requests for help from the field. Sometimes they know the answers themselves, and sometimes they must find the people who know the answers.

It is important to note that being an INTOUCH engineer is a sought-after position. These are not the most junior people on the staff. Rather, they are experienced field engineers with extensive knowledge and a web of contacts in the technology centers and in the field organization.

Frequently asked questions (FAQs): The Schlumberger community culture is such that if an engineer has a problem, he/she is expected to do the homework first; that is, search the knowledge hub for solutions, best practices, and FAQs. The next step is to contact the help desk. The INTOUCH engineers have the responsibility to refresh the FAQs—to ensure that field engineers can find the answers to the common questions themselves by browsing the knowledge hub. This has the additional benefit of reducing the repetitive load on the help desk, which allows the company to dedicate a smaller staff to the task than would be required without following a knowledge management process.

Figure 2. Career Networking Profile.

As shown in figure 1, the company must close the loop on CoP operation. The community activity is a source of valuable intelligence on opportunities to develop new products and services. In addition, the best practices and FAQs must be used to continuously refresh the training courses and manuals. Finally, what is being learned must be applied in the day-to-day job of each community member—serving the customer.

As noted at the beginning of this section, everything presented in the example is in use today in the business. It is not a vision of the future. However, we must note that what we have shown is not used in all businesses, with all customers, every day, anywhere around the world. Our primary goal is to make what we have shown the norm throughout the organization.

The Best-Practice Knowledge Management Process

In this subsection, we look more closely at the basic process associated with best practices in a CoP. The basic cycle is shown in figure 3.

In the course of their normal work, community members—practitioners—apply the current best practices. As they encounter new situations, new customers, or interact with their

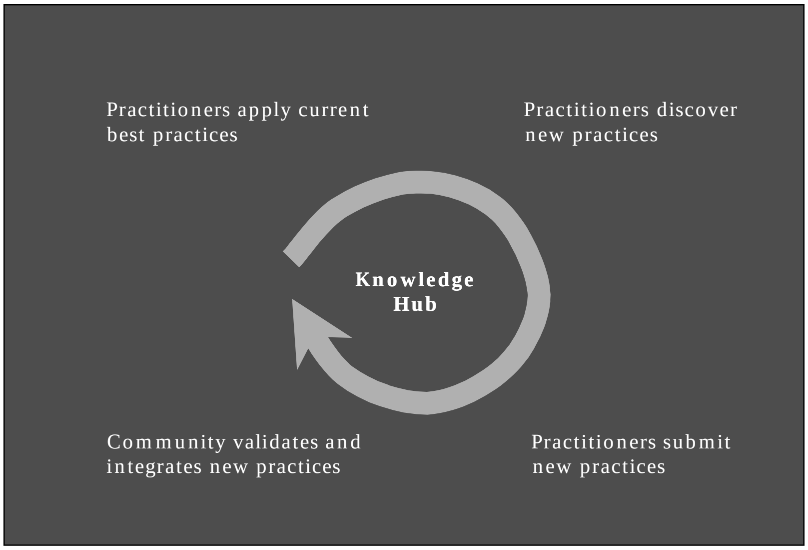

Figure 3. The Best-Practice Knowledge Management Process.

colleagues, they discover new practices. They submit them as best-practice proposals. The community validates and integrates the new practices into the overall set stored in the knowledge hub, and the cycle continues.

In the days before the intranet web, the community knowledge hub might have been a person with a filing cabinet and a telephone. Today, it is most likely a web-enabled database. Submission can take many forms—a telephone call, an e-mail message, filling out a best-practice web form.

Different companies adopt different approaches to validation. Some take the point of view that every community member should be trusted. One is laying his/her reputation on the line when posting a best practice, and that is seen as sufficient evidence of its validity. Other companies have formal validation processes. They might involve the entire community or only a few domain experts. Some companies involve a variety of people outside the community in the process (for example, lawyers, marketing staff members).

It is important to balance the tension between the quality of the knowledge and the time it takes to validate it. The risks of disseminating and applying invalid best practices are obvious. However, another risk is that the time to validate can be so long that the published knowledge will only be good for archival purposes; that is, it will be published too late to be of any practical day-to-day use in the community.

At Schlumberger, best practices represent the validated knowledge assets of each community. A *knowledge champion* is responsible for animating the community, encouraging the members to participate, highlighting successes, recognizing the contributions of individual members, and so on. The knowledge champion coordinates the processes of validating the practices submitted by the members and integrating them into the overall community knowledge repository. Because the knowledge champion is well placed to understand the issues in which the community members are interested today, he/she edits the community news.

It is important to note that the community must address an ongoing maintenance task. It is the *content maintenance task*—keeping the community knowledge hub fresh and up to date. A person who knows the community well—a knowledge champion—must perform this task. The champion addresses questions such as, What are the problems being encountered by the community today? What best practices should be retired? What best practices can be integrated? Where is the repository weak? In what areas do best practices need to be solicited from the community?

Finally, note that all community members are involved in the task of knowledge acquisition, which is a marked change from the past. In the expert system days, this was a task for a group of dedicated *knowledge engineers,* AI specialists charged with capturing and codifying the knowledge of the community. It involved painstaking interviews and coding sessions, a slow process exacerbated by the fact that the knowledge engineers had to take additional time to achieve an adequate level of understanding of the domain. Today it is assumed that community members are able to capture the knowledge largely by themselves. At best, a knowledge champion, perhaps coached by a small group operating across the organization, helps them. What knowledge engineers remain are likely to be part of that small group.

Example: Drilling Best Practices

Figure 4 shows part of the first page of the Drilling Best Practices section of the Schlumberger knowledge hub. It was built and is maintained and used by the company's directional drilling community, approximately 1400 engineers.

On the left, note that the best-practice contributors are highlighted in the news. In the center, most of the Yahoo!-like categories are linked to the best practices for various specialized types of drilling (for example, extended-reach drilling [ERD], horizontal wells). Two important categories are the following:

Experiences mined from bulletin boards: To get the best-practice process started in the directional drilling community, the knowledge champion examined the past year's messages on the community bulletin board. He found messages that summarized best practices. He reformatted them, categorized them, and seeded the knowledge hub with them. He did this so that on the day the knowledge hub was announced to the community, there was a real possibility that the members would be able to find something of value and hence be encouraged to return to the site in future and to contribute to its evolution.

He was aided in this task by an informal process that had been followed by the community. When a question was posed to the bulletin board, a number of responses received, a back-and-forth debate carried out, and the question eventually resolved to the satisfaction of the members, one of them would write a summary message, detailing what had been decided. Placing these summaries on the knowledge hub in the appropriate categories improved the ability of the members to find them quickly and therefore increased the likelihood that they would use them.

Best-practices study: The first task undertaken by the knowledge champion was a study of what had been done in other organizations. The set included several customers and other companies, such as Xerox, that had gained a reputation for knowledge management. Examples of what was learned are the importance of a simple process (from Xerox) and the utility of a tiered approach to validation (from Chevron). The three tiers that were eventually adopted are (1) good idea, (2) local good practice, and (3) Schlumberger best practice. A *good idea* is as yet unproven, but the presence of this tier offers the possibility for all members to contribute. A *local good practice* has been implemented and improved results in a particular location. A *Schlumberger best practice* has been determined to be the best approach across the organization.

Two points are important: First, the drillers built their knowledge hub themselves; the central knowledge management group coached them on the process and provided them with simple portal technology so that they could take charge of their own content. Second, they had an effective knowledge champion (figure 2). It is difficult to overstate the importance of real people behind every successful knowledge management project.

Figure 4. Drilling Best Practices.

Lessons Learned

To complete this review of today's knowledge management practice, we summarize lessons learned to date in our work at Schlumberger, a report from the front lines after three years of action.

The Biggest Challenge Is to Nurture a Knowledge-Sharing Culture

A *knowledge-sharing culture* is one in which people share their knowledge and learn from others as a matter of course; they see it as the right thing to do.

First, a culture change requires the support of line management. Hence, knowledge management must be visible on the business radar screen; targets must be chosen to address important business problems that capture the attention of line managers.

Second, incentives and recognition must be provided to encourage people to exhibit knowledge-sharing behavior—to contribute their own best practices and to reuse those of others. The line managers are best placed to encourage this behavior—by setting knowledge management objectives and assessing performance based on them.

Third, changes must be made to the business processes to embed knowledge management in the workflow. An example of such a change is a mandatory check for best practices before every job. Another is to contribute lessons learned after the job has been completed.

The line managers are also in the best position to implement organizational changes that result from a strong knowledge-sharing culture. An example is the introduction of the Schlumberger INTOUCH help-desk organization discussed earlier.

Finally, line managers can provide important personal support. This can be as simple as asking the right questions during every management review or field visit. For example, "What solutions have you developed recently?" is a standard question, but following it up with "With whom have you shared them?" "What have their results been?" and "What solutions have you learned from others?" drives change if the questions are consistently asked.

Vital Virtual Communities Are Essential

The community of practice is the central organizational unit in knowledge management. Communities are supported by technology, but relationships among members and a shared sense of purpose bring them to life. The knowledge champion plays a critical role: to be the focal point for questions, report on progress, push the process (for example, validation), edit and publish the news, highlight successes, and so on. The champion also "kick starts" the process by building the initial knowledge repository by mining the existing information sources (for example, bulletin boards) for content. Without a person behind it, any web site will be static and die. It must be dynamic to live.

In the example of the directional drilling community, there were several ways in which the knowledge champion kept knowledge management visible. He made presentations to classes at the learning center and at management meetings. He made field trips and discussed knowledge management with visitors from the field. He built a team of local knowledge champions around the world and wrote a monthly letter to them. He maintained an up-to-date news section of the community hub. He sent updates of knowledge management activities to the community bulletin board. He worked with management to put in place a program to recognize and reward knowledge management successes. He also armed managers with questions to ask on field visits.

While Technology Isn't Everything,… Little Progress Will Be Made without It

In companies that operate across the globe, if not in all companies, technology can be likened to the entry fee. With technology alone, you won't win the prize of achieving bottom-line success, but without technology, you won't even get into the game. Referring back to the previous lesson, it is important to get started as soon as possible, rather than wait for the next generation of technology. Begin to build the community with the available technology and upgrade it later (Hagel and Armstrong 1997). Again, a useful rule of thumb is to spend no more than one third of the knowledge management budget on technology (O'Dell, Elliott, and Hubert 2000).

Build It and They Won't Come

Just because an intranet portal has been built filled with world-class technology, it is not a given that community members will flock to it. Do not overwhelm them with all the features that computer scientists can think of that "clearly" would be beneficial. Instead, be cautious. Determine first what technology the community members actually use. In our case, we had solid evidence that they all used e-mail every day and that about one-third used the web daily. As a result, we started the communities off with a small evolution to the technology they already used. Now we are working with the community leaders to continue to refresh the technology to meet their needs over time.

It's about the Content

An up-front investment is required to seed the initial knowledge repository. It is difficult, if not impossible, to convince community members to contribute to an empty shell. (How many content-free web sites do you revisit?) Not only must there be content from the launch date, but it must be quality content as well. The lesson here is to use technology that enables those with the knowledge to be responsible for publishing it, upgrading it, and retiring it as needed. Moreover, the credibility of a best practice is directly related to the reputation of the person who wrote it. It is therefore doubly important to recognize and publish the names of contributors.

Everyone Is a Potential Contributor

If simple publishing interfaces are put in place—what we call *light publishing*—the possibility exists to have a large number of contributors. Contrast this situation with that of the past, where only a small number of people had the opportunity to publish, typically those working in the research and technology centers. They published outside the company in technical journals. The published documents represented substantial bodies of work, often extending over several months or years. However, the intranet portal offers a means to publish inside the company to a broader selection of people. The publishing is on a lighter scale than full-scale research projects. It includes small, useful ideas and best practices.

When we say everyone is a potential contributor, we mean it. For example, some of the strong contributors to our program have been retirees.

WIIFM: What's in It for Me?…Now!

We have learned that this question must be answered in the present tense, not the future tense. It is not good enough to explain what will be in it for the various stakeholders someday. The community members must solve their problems today, not someday. This again emphasizes the importance of seeding the initial knowledge hub and addressing the issues of incentives and recognition early.

Stakeholders Are Community Members and Business Managers

The stakeholders are the community members and the business managers. The question "What's in it for me?" must be answered from the point of view of the person who seeks information, the information publisher and the business manager (where the answers include cost savings, faster training, improved decision making, and improved asset management).

Walk the Talk

A good pair of interview questions for a knowledge management salesperson, especially one selling software is, "How do you manage your own team's knowledge? What tools do you use?" The evangelists who exhort the organization to get serious about knowledge management must be seen to practice what they preach.

The Road Ahead

In this section, we look forward at the road ahead for knowledge management—how the technology and processes that underlie knowledge management might evolve over the next decade or so.

Our approach is to present a "technology road map" that highlights computational technology on the critical path to reach a new era in organizational performance. We hope to stimulate the AI community and provide signposts for the research and development needed to achieve the successes hypothesized in the vignettes.

The road-map method has been popularized by Motorola (Willyard and McClees 1987) as an orderly process for developing a picture of future technology together with a projection of its evolution over time. It was originally intended as a practical tool to encourage business managers to give proper attention to their technological future. It also provides a means of communicating to engineers and marketing personnel which technologies will be requiring development and application for future products.

The process starts by stating a clear overall target, based on product market, competitive or technology trends. Armed with such a target, developing the road map is an iterative brainstorming process, often involving people from a variety of functions (for example, research and development, manufacturing, marketing, business). The fact that road-map time frames tend to extend well beyond conventional business and product planning horizons adds to the challenge. Another benefit of the process is that it forces the participants to be explicit about their assumptions—to make clear the problems that must be solved to reach the target, the order in which they will be solved, and the expected interim results.

The process has been adopted by the U.S. semiconductor industry overall. It has been

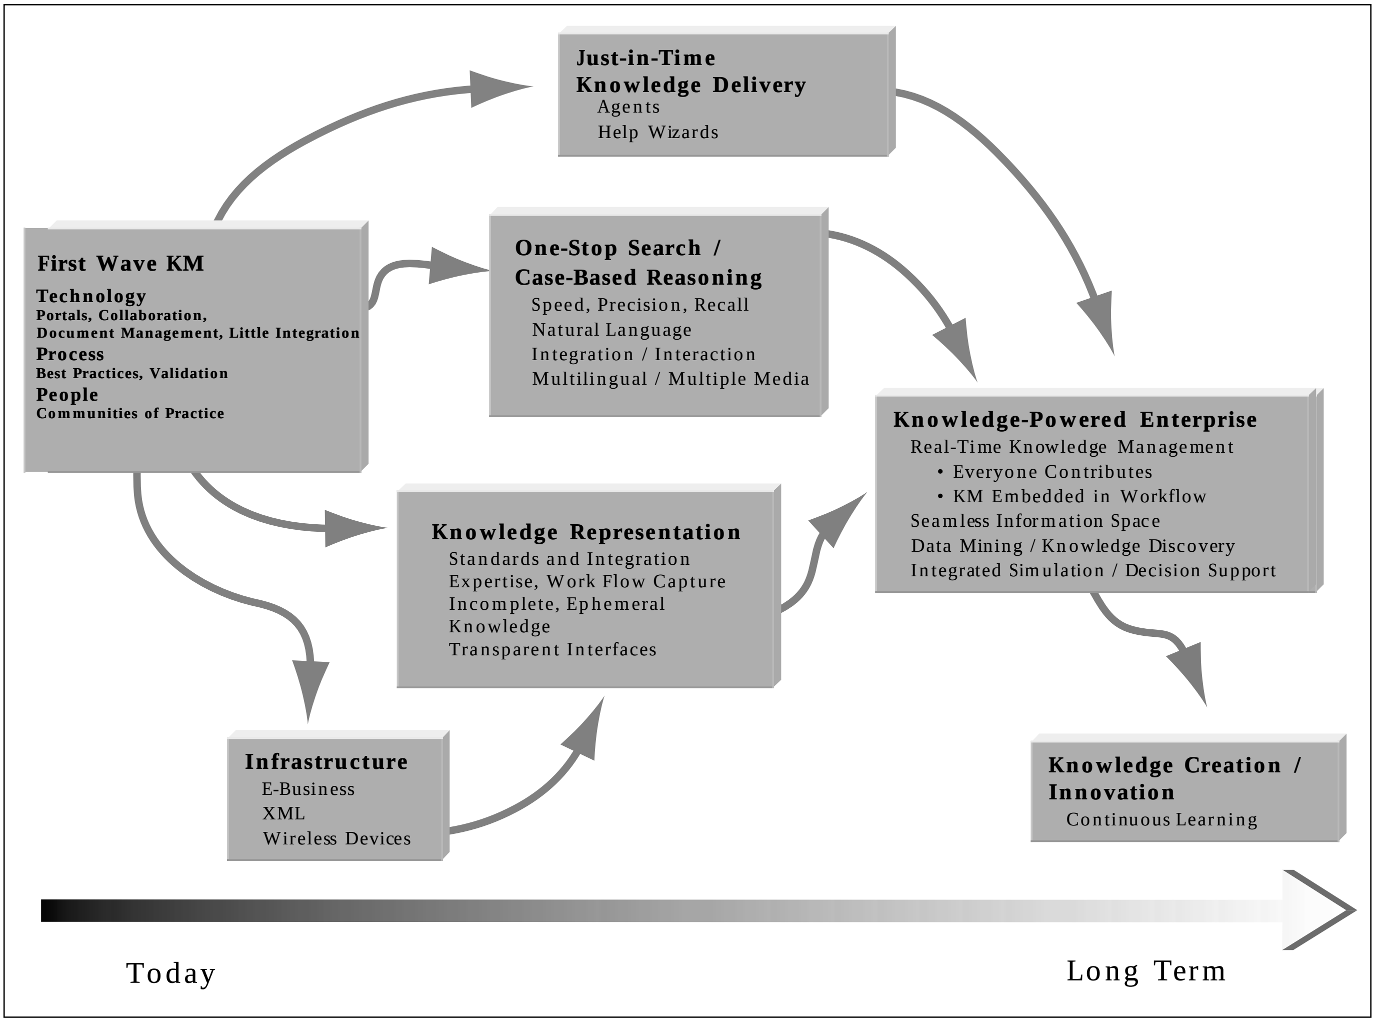

Figure 5. Knowledge Management Technology Road Map.

used to provide a common vision, a framework to guide research and development for all sectors of the U.S. semiconductor technology base—industry, universities, and government organizations (SIA 1994). The overall target was established by extending historic trends for dynamic random-access memory (DRAM) bit count by a factor of 4 every 3 years until 2010, implying a 64-gigabyte chip in 2010. From this follow the requirements for supporting technology to enable such devices to be designed, manufactured, and tested.

In the upstream exploration and development sector of the oil and gas industry, a simply stated target is to double economic hydrocarbon recovery from the current typical levels of 30 to 40 percent to 60 percent and beyond (Smith and Maitland 1998).

For knowledge management, such a crisp target is elusive. As a result, we have suggested the goal of achieving the knowledge-powered enterprise.

The road map is shown in figure 5. *Today* is the current state of the art in knowledge management. *Long Term* is the possible state of the art a decade or so hence. The boxes represent important points along the way.[1]

In the following subsections, we examine each of the boxes on the road map, starting with today and proceeding to the long term.

First-Wave Knowledge Management

Figure 6 shows the current state—first-wave knowledge management.

People: The people focus is on the creation, support and nurture of CoPs.

Process: The process focus is on the capture, validation, integration, and dissemination of best practices (see The Best-Practice Knowledge Management Process and Example: Drilling Best Practices).

Technology: The technology focus is on the portal, backed up by the sorts of content discussed in the earlier example. Access to the intranet is assumed.

A portal is an application that gives users a single gateway to the information and applications they need to do their jobs. It draws together on the desktop all the important information from both inside and outside a company—everything that is needed to make a business decision and take action. Microsoft uses the term *digital dashboard.*

The word *portal* was first applied to knowledge management in 1998, following the observed success of internet sites such as Yahoo! in enabling users to find information (Shilakes and Tylman 1998). Building an intranet portal is now a standard first step in knowledge management.

To be sure, there are inflated expectations and hype associated with portals (called *portal mania* in the July 1999 issue of *Knowledge Management Magazine,* devoted to the subject [KM 1999]). Portals have been put forward as the silver bullet for a number of information technology problems, such as revamping enterprise resource planning (ERP). In response to the excitement, a number of software companies have relabeled their offerings, and many more startups are offering portal technology to support knowledge management.

However, despite the hype, there is something of real value at the heart of the interest—the portal might be the killer app for knowledge management.

Today, a portal provides unified access to all the organization's information, both unstructured and structured. The *unstructured information* is mostly web pages and documents, including product plans and marketing brochures. The *structured information* is usually stored in databases. In the case of Schlumberger, the structured data include reservoir data, semiconductor test data, and simulation data of one sort or another.

Portals deliver information to the desktop by way of a "thin client," usually a web browser. They typically offer two ways to find information, following the original internet portals:

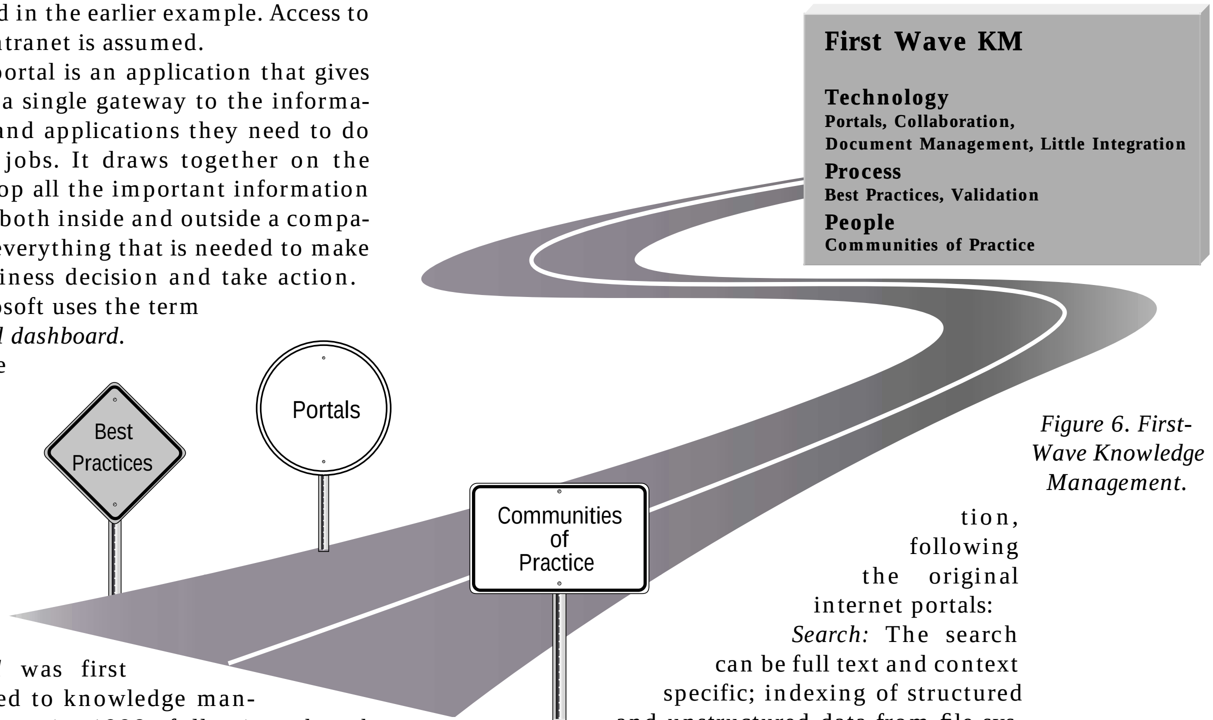

Figure 6. First-Wave Knowledge Management.

Search: The search can be full text and context specific; indexing of structured and unstructured data from file systems, web servers, e-mail, and so on. The search engine is often the "big-ticket" item in a portal solution.

Browsing through a hierarchy of categories: Creation and maintenance of a common vocabulary and one or more category hierarchies (ontologies) are key tasks associated with portal construction.

At Schlumberger, we have found that information seekers are evenly divided between finding information by searching and finding it by browsing. In a recent survey of over 1000 people, 10 percent claimed to find information by browsing, 23 percent mostly by browsing, 37 percent by both searching and browsing, 20 percent mostly by searching, and 10 percent by searching. The symmetry of the distribution is striking, as is the divergence of preferences, which shows how crucial it is to design portals that support distinct information-seeking paradigms.

The design, construction, and maintenance of the category structures used to organize information in knowledge management systems is an area where AI technology could help substantially. The task is similar to that of building large-scale knowledge bases or ontologies. Projects such as ONTOLINGUA (Farquhar, Fikes, and Rice 1997) have begun to develop tools and methodologies for distributed collaborative ontology construction. This sort of approach is critical for knowledge management, in which communities of practice are often geographically distributed.

In addition, there is considerable leverage to be had by using the ontology explicitly as a guide to search, allowing the user to search in the "semantic neighborhood" of a concept (Clark 2000).

Increasingly, portals are also expected to provide services such as the following:

Presentation-visualization: This includes web display, data visualization, abstract information visualization, and presentation in context. The web pages displayed by a portal are often generated dynamically to match the task of the user.

Subscription-notification: It is usually possible to subscribe to a series of channels. Agents learn an individual's interests, actively go through news feeds, select the relevant news, and "push" it to the user by e-mail.

Collaboration: This includes threaded discussions, chat room sessions, and video interactions. This technology is established, although the interfaces for integrating existing collaboration tools into portals are not yet well established. Technology that supports synchronous interaction (for example, chat, video) is, of course, more useful in regional communities, where the members work in roughly the same time zone.

… information seekers are evenly divided between finding information by searching and finding it by browsing.

Personalization: Presentation of information is customized for individual users; agents filter information. LDAP authentication is sometimes used to reduce the number of separate user names and passwords needed by people.

Publishing and distribution: A document management subsystem is a vital component of a portal solution. The support for publishing is primitive in most current portals, which place a greater focus on access than on publishing. We expect to see more emphasis on publishing in the near term.

Data feeds: News, stocks, weather, and so on, are standard. However, a number of company- or task-specific real-time data feeds are of interest (for example, semiconductor manufacturing yield data, flow, and temperature data from producing oil wells).

Security: High-end portals must support fine-grained authorization and access control, which becomes increasingly important as an organization accesses, stores, and organizes more of its knowledge using the portal. For example, employees, customers, and suppliers require access to overlapping information packages but must be prevented from accessing information to which they are not entitled. Similarly, within the organization, some information is for managers only; some is restricted to the members of a project team; and so on.

Integration: Today there is relatively little integration between portals and data management systems, back-end databases, or ERP systems.

Over time, intranet portals can be expected to follow another trend that has occurred in their internet ancestors. They will become places to do something as well as places to find something. That is, they will offer a variety of application services. Indeed application service providers (ASPs) have already begun to integrate portals with the applications they offer on demand over the internet.

Returning to security, one advantage of the portal approach is that a user can sign in once and continue to navigate through the knowledge space without having to sign in again to access each individual application for which he/she is an authorized user.

A major cost of knowledge management comes from the tasks of generating, collecting, organizing, and storing information. A major risk comes from using information of poor quality—out of date or inaccurate. By reusing information for multiple tasks and multiple audiences, the cost and risk can be reduced. For example, both employees and customers can use technical information about a company's products and services.

Using the same information for multiple purposes and audiences provides substantial benefits. It enables amortization of the cost of knowledge management and information technology support. It improves quality because more people review, access, and use each piece of information. It can also have substantial benefits for a company because the employees are used to seeing the same information organized in the same way as their customers—they are on the same page as the customers.

Of course, sharing is more subtle than showing the same information to everyone in the same way. Each audience might need a different view of the underlying information—a view that is filtered, organized, and presented in a slightly different way. For example, a good practice might be appropriate to show to engineers with the experience to evaluate its applicability for their activities and understand the context in which it should be used, but the same information might not be appropriate for trainees or customers.

Similarly, although a company's products and services should be presented to customers at the top of the list of information categories, perhaps information about customers should be at the top of the list for employees.

As an example, the Schlumberger hub uses a common ontology for employees and customers. We have developed and deployed technology to allow customers and employees in internet, intranet, and extranet settings to use shared information that is filtered, organized, and presented through custom views. It is this technology that enables the success of vignette 4.

First-wave knowledge management processes in CoPs typically require assistance by support staff. This might include journalists who capture and report news or highlight people who have participated in best-practice activities. Librarians can help employees codify and disseminate information. Knowledge champions can animate the knowledge management processes and validate best practices. An alternative that can work in some CoPs is self-service knowledge management. In this scenario, CoP members are expected to find what they need by themselves as well as publish and evaluate information themselves.

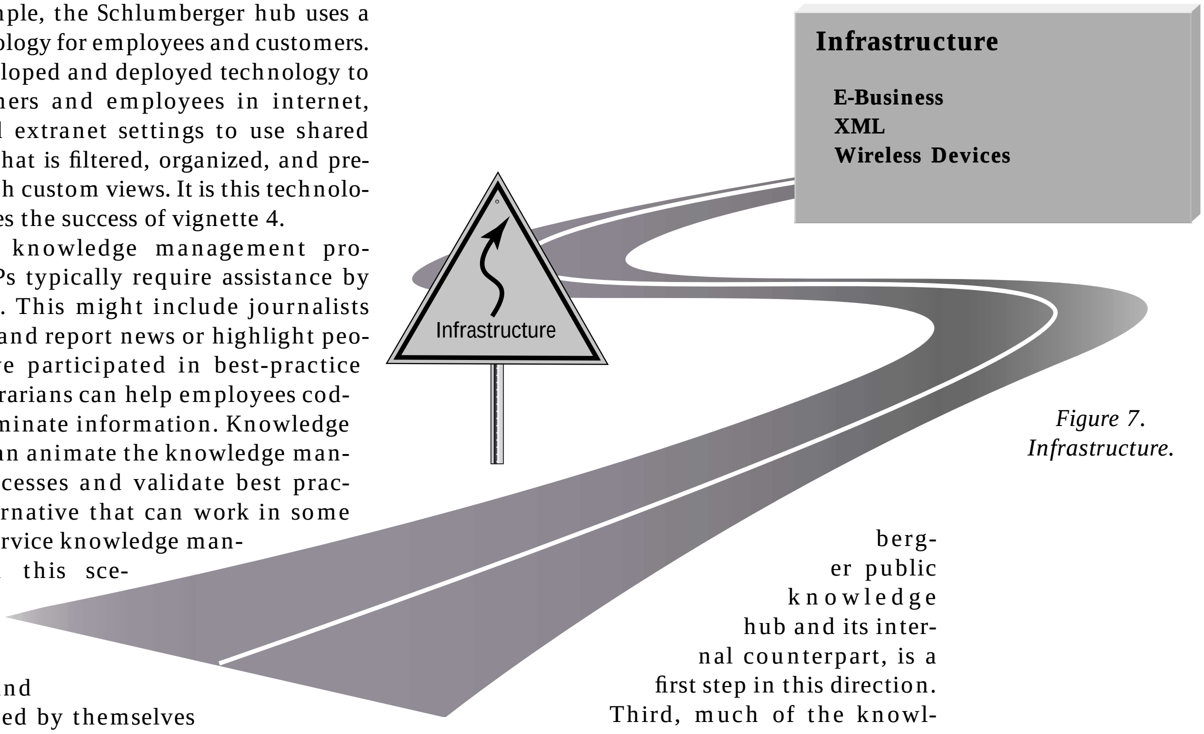

Figure 7. Infrastructure.

Infrastructure

As we proceed along the timeline from today, we arrive at the first near-term way point. It highlights some infrastructure developments in progress (figure 7). These do not necessarily involve AI directly but are nonetheless extremely important.

e-business: It is not just about selling products online, as e-commerce was originally defined. It covers all aspects of relationships with customers, suppliers, and partners (for example, sales, marketing, technical support, news, tracking and reporting progress, linking customers and employees). It also includes service; infrastructure; and multiparty, multidirectional (business-to-business) transactions (Karlenzig 2000).

Knowledge management projects lead naturally to e-business projects over time. Why?

First, the benefits that an organization reaps from knowledge management activities (for example, improved practices, access to up-to-date technical information) will also be beneficial to the organization's customers, suppliers, and partners.

Second, the portals and other knowledge management systems that ensure employees have access to critical information can provide a similar service to customers, suppliers, and partners—the requirements are almost identical. In the previous subsection, we noted how portal technology, common to the Schlumberger public knowledge hub and its internal counterpart, is a first step in this direction.

Third, much of the knowledge that companies are trying to reuse is generated beyond the bounds of the single enterprise. It is generated by the so-called "extended enterprise" (customers, suppliers, partners, and alliance members) and beyond (Dawson 2000). It is essential to have access to this external knowledge.

Some organizations use their knowledge as a competitive differentiator, some sell it directly, and others offer services and technology to enable their customers to improve their own knowledge management. Three other possibilities for direct exploitation of knowledge management are customer-specific (extranet) portals, industry portals, and online industry product catalogs.

The *customer-specific* (extranet) *portal* provides direct information and support for products and services that are required to support an organization's relationship with a specific customer.

For example, the Schlumberger coiled-tubing drilling (CTD) library is an extranet portal set up to support BP/Schlumberger CTD work. It is available globally to personnel from the two companies. The trend in applications such as this is clearly moving from finding to doing. We expect direct access to services, such as data management or seismic data analysis services, to become widespread.

The *industry portal* draws traffic and revenue by providing a variety of things to see and do. Such portals exist to serve many industries (for example, plastics, metals, medical). We expect this trend to continue (along with consolidation in the number of portals that serve any particular industry). An example of an industry portal that serves the energy industry is IndigoPool.com. In addition to news and information, it provides an online oil and gas asset-trading service, complete with the data management and economic simulation tools necessary to evaluate the fit of the assets with a potential buyer's portfolio.

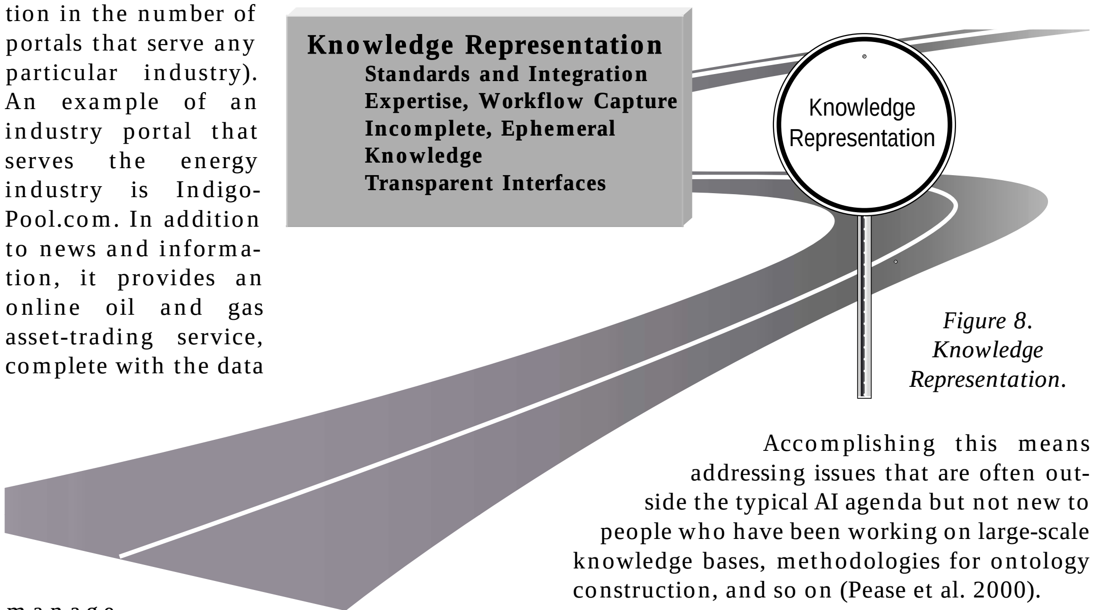

Figure 8. Knowledge Representation.

XML: The extensible markup language is used today in the catalogs of B2B industry portals. We referred earlier to some weakness in publishing and editing in today's portals. Knowledge management involves a continuous content update process. We anticipate that XML will play a significant role, as it has already started to do in some of the e-business offerings.

Knowledge management is essential for e-business.

Wireless devices: By 2002, it is expected that web-enabled wireless device shipments will surpass PC shipments. By 2005, it is expected that there will be 2 billion wireless internet devices. One effect of this explosion will be that people will expect to have the information they need to do their jobs not only anywhere, any time, but also delivered by any means to the device(s) of their choice. There is interesting work to be done to develop general methods for intelligently selecting-synthesizing the appropriate information to be delivered to different device types.

In addition, means by which wireless devices can communicate with intelligence, that is, with awareness of the contextual environment, are needed. Developing these means in the near term is a good agenda item for the AI community (du Castel 2000).

Knowledge Representation

Knowledge representation is clearly important to knowledge management. The challenge is to put knowledge representation tools in the hands of the people who have the knowledge (figure 8).

Accomplishing this means addressing issues that are often outside the typical AI agenda but not new to people who have been working on large-scale knowledge bases, methodologies for ontology construction, and so on (Pease et al. 2000).

Transparent interfaces: The interfaces to knowledge representation tools will have to be transparent. Busy people with "regular day jobs" are the front-line knowledge management practitioners. These are the people who build knowledge hubs. By and large, they are not computer specialists. This means that the tools must be extremely quick and easy to use, even on an occasional basis, and that the tools must be embedded in a user's normal workflow. The activity cannot be seen as an optional extra. We believe that there is exciting work to be done by the knowledge representation community to enable those who have the knowledge to use knowledge representation tools.

Workflow capture: Knowledge management applications frequently revolve around workflows defined for particular tasks (for example, drilling a horizontal oil well). The workflow is typically a graph of steps, loops, decision points, and so on, associated with performing the task. Representations and tools for capturing workflows (Lee et al. 1998) will be valuable, especially if they integrate well with other information (for example, best practices, training material, financial applications) and with other applications.

Expertise: When a CoP member has a problem to solve and is unable to find a solution, the next step is to find an expert—a person who can help. Today, this is done by keyword or full-text intranet search—some sort of structured search—for example, in an LDAP directory or a database of people profiles. We believe that a knowledge representation approach has the potential to produce more effective techniques for matching those who have problems with the experts who can help.

Incomplete, ephemeral knowledge: The knowledge of an organization is often messy, ephemeral, and incomplete. It might not be accurate, and it is rarely neatly packaged. As an example, each time a new product is introduced, the field community must learn a new set of best practices. We believe an interesting challenge for the knowledge representation community is to develop technology and methods to deal with this type of knowledge, representing the dynamics of changing knowledge (Heflin and Hendler 2000). In addition, improved methods and tools are required to assist the CoP in following its knowledge management process. On an ongoing basis, it must keep up with what's new, what's working, what doesn't work, what should be retired, and so on.

Integration: For the knowledge representation community to have an impact on the future of knowledge management, it is important that the tools and ontologies are plug-and-play compatible with the other knowledge management technologies. Key among these are portal and search engine technologies. Although there have been some efforts to provide standard application programming interfaces (APIs) to support knowledge base interoperation (Chaudhri et al. 1998), these have not provided practical integration with existing non-AI technologies.

Standards: By standards, we mean for representation as well as standard ontologies. XML is important here. Of course, XML and related protocols and standards will evolve on their own, whether the AI community gets involved with it or not. However, the knowledge representation community is well placed to drive the evolution—as it has in emerging standards such as the resource description framework (RDF) (Brickley 2000; Decker et al. 2000).

One-Stop Search / Case-Based Reasoning

One-stop search: Today, one uses full-text search with a web browser, database search for structured data (for example, seismic data), Find on a local PC, and so on. Each of these searches requires the user to learn a different interface. Furthermore, one must know which system to search to find the desired information.

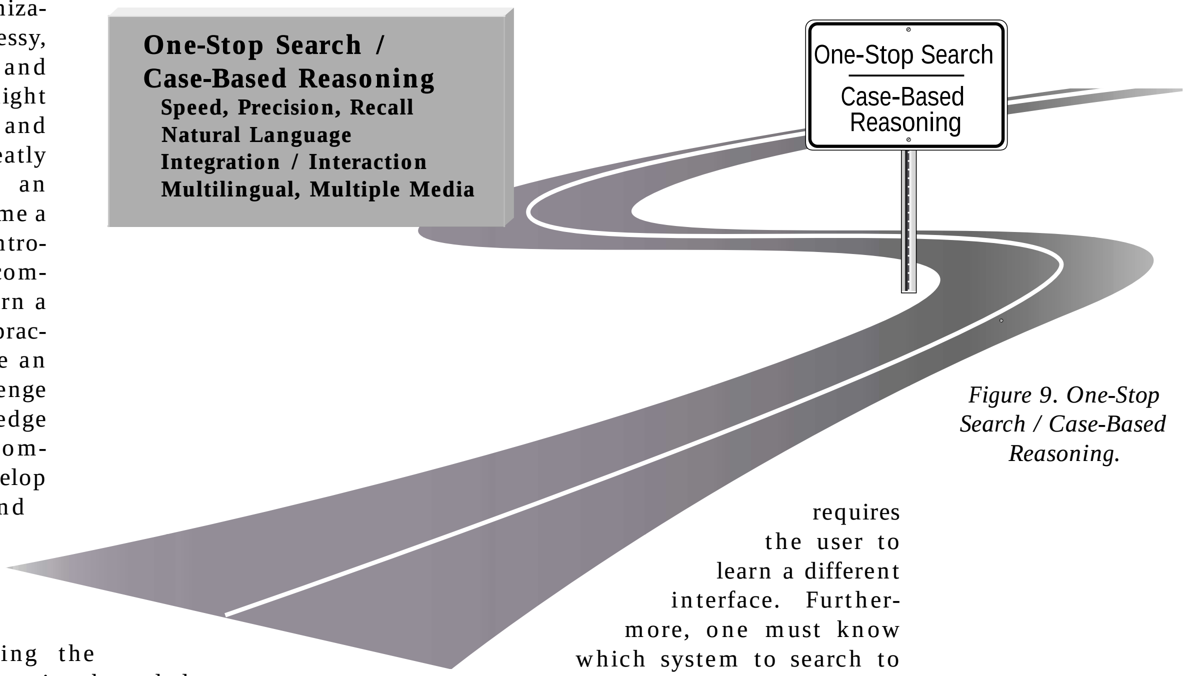

Figure 9. One-Stop Search / Case-Based Reasoning.

We envision a one-stop shop for search with a single interface, one search box if you will (figure 9). We suggest that the portal approach be applied to one of the major components of any portal, the search engine itself.

There are at least three challenges: First, the search portal must understand a user's query well enough to direct it to the appropriate index, database, or file system. Second, the interface differences of the various subsystems must be hidden from the user. Finally, the search portal must retain the power of the previously separate search systems and deliver it to the user.

Case-based reasoning (CBR): CBR appears in figure 9 because it shares many common elements with search (for example, indexing, information classification, and extraction). We expect case bases to be included in the repositories accessible by a search portal (for example, the project archive given in the example earlier).

Speed, precision, recall: We assume that the search engine companies will be hard at work to make substantial progress on speed, precision, and recall.

Natural language: There are many opportunities to make a difference here. Among these is evolving the ideas used in "Ask Jeeves" or Dell's version, "Ask Dudley." This is an area where there is interesting work to be done and where successful approaches will continue to be adopted quickly in commercial applications.

Integration-interaction: We are thinking about integration with other applications and systems. By *interaction,* we mean changing the concept of

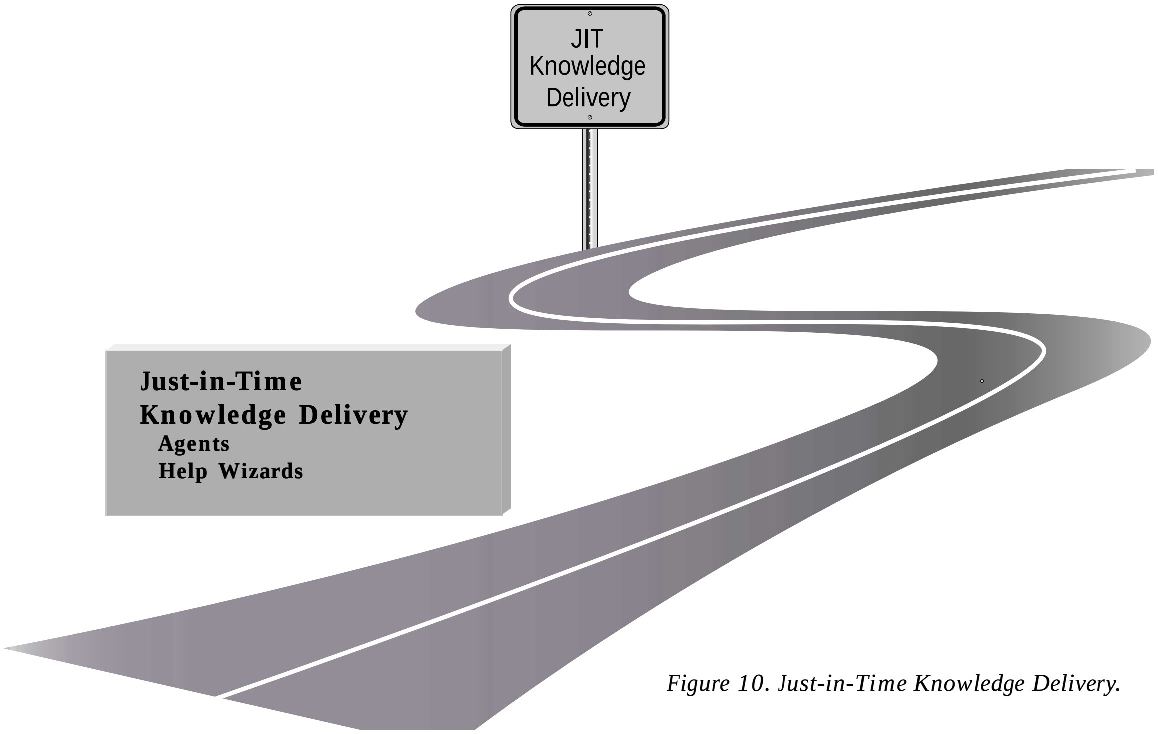

Figure 10. Just-in-Time Knowledge Delivery.

search so that the search engine doesn't simply return information; it bundles that information with an application so that the user can interact with it (for example, update a three-dimensional model or run a simulation). This starts to happen in the comparison shopping systems of today. We expect that this is going to be an exciting area over the coming years.

Multilanguage: Support for multiple languages is obviously important for organizations that have diverse customers and employees around the world. We expect much improved approaches to multiple language search to be developed. We also look forward to improvements to support multiple-language authoring and content maintenance.

Multimedia: Today, most knowledge hubs are text oriented. This is not by desire; it is apparent that much of the knowledge of an organization is not captured in text but in other media. However, the technology to support indexing and searching of images, sound files, and video is immature. We include multiple media in figure 9 to emphasize that the knowledge management community is anxious to use new technology in this area if and when the AI community can deliver it.

Just-In-Time Knowledge Delivery

Just-in-time knowledge delivery means that the knowledge needed to perform a task is delivered when it is needed and in the context in which it will be used (figure 10).

Agents: We expect significant advances in this technology. Agents will find and bring to our attention (for example, by push) the things that we didn't know we needed to know. Today, agents exist to monitor news feeds and other real-time data streams and return only the items indicated by a user's profile. Over time, these agents are going to get much smarter about finding and summarizing the relevant information. We expect they will be able to personalize the summaries for individual users. This means that two individuals might receive different summaries of the same source document, depending on their profiles. We also expect advances to be made in the ability of agents to select and synthesize information for display on wireless devices. Finally, we expect that it will be standard for users to solve large problems by contracting with a number of agents. Each "entrepreneur" agent will act independently, bidding for the work and delivering a piece of the overall solution.

Help wizards: Online assistance in the context of solving a problem, integrated with the workflow, will become the norm for applications. As an example, in oil-well–drilling planning software today, it is possible at each step in the process to view training material to assist in making decisions. We expect technology to be developed that will allow these training materials to be turned into active help wizards. Imagine having access to active, multimedia

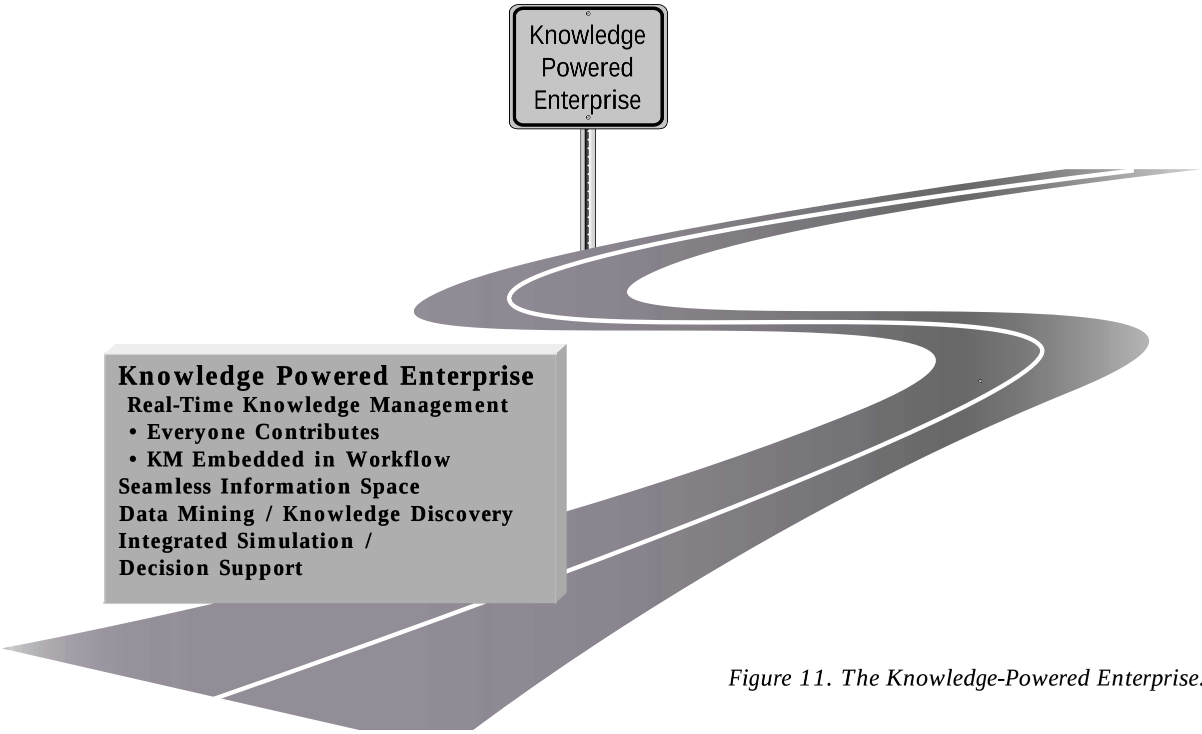

Figure 11. The Knowledge-Powered Enterprise.

best practices, explained by the experts themselves, when you most need help.

e-learning: Just-in-time knowledge delivery means learning anywhere, any time, on demand, and in context. Today, self-learning is enabled by online CoP knowledge (for example, best practices) and interactive computer-based training materials. Asynchronous distance learning is also used, wherein a remote instructor is available by e-mail. Finally, synchronous distance learning is carried out, wherein a course is offered by net broadcast, and the instructor is available during the session. This is an area close to the expertise and experience of the AI community. There is a good deal of work to be done, and a receptive knowledge management community is ready to apply it quickly if it is robust and can be integrated with the other knowledge management systems. A simple example would be a "training wizard" that observes the browsing and searching done by a user on the knowledge hub and makes inferences about the training courses that may be of interest.

The Knowledge-Powered Enterprise

Farther forward, beyond the first chicane, the next visible road sign is "knowledge-powered enterprise" (figure 11). In such an organization, knowledge management is "organic" (Shevlin et al. 1997). It happens everywhere—in the background, in real time. Everyone contributes. Knowledge management process and behavior are embedded in the workflow as part of the normal day-to-day job. Knowledge management functions are embedded in core business applications and employee productivity tools.

Knowledge management specialists may still be present to assist, for example, librarians. However, the specialists will be part of a small team that operates across the company.

There are three important components from a technological point of view:

Seamless information space: Everyone has rapid access to the information needed to do his/her job. The portal concept will be extended across the organization. CoP members will not need to know that the information is stored in a particular application, which is on that server. It will not matter. They will get the information they need when they want it, where they want it, and how they want it delivered.

Data mining and knowledge discovery: Today, most enterprises are still at the stage of trying to capture the data in an organized fashion, let alone trying to mine it to influence business decisions. However, we expect three threads to come together: First there will be continued evolution of data management technology. Second, powerful new data-mining and knowledge-discovery technology will be developed. Third, these two types of technology will be coupled with portal technology to deliver the information needed to perform a task just in time for it to be of use (Mitchell 1999).

Integrated simulation and decision support: Vignette 3 dealt with an oil company well engineer talking to a potential investor. The investor was able to compute the likely effect of a particular oil well operation, drilling a sidetrack, from a geoscience point of view. He then used those results to compute the rate of return and risk associated with the proposed investment. This kind of combined geoscience, economics, and risk simulation can be accomplished today, but it is not commonplace, nor does it happen in real time.

By including integrated simulation and decision support in figure 11, we do not mean to suggest that the AI community will introduce the necessary technology. Indeed, in a company like Schlumberger, a sizable fraction of the entire research and development budget is devoted to building the tools that can be used by our clients to solve problems such as the one in the vignette. However, we do believe that the AI community can make a useful contribution to the design of these tools. It is important to ensure that they can be coupled in useful ways to the seamless information space by means of the intranet portal (including its case bases, best practices, and other information) and with the data-mining and knowledge discovery systems.

Looking Further Forward

In the knowledge-powered enterprise, knowledge sharing and application are standard. As a result, the enterprise has substantially improved its performance and reduced its costs.

The new focus is the end game for knowledge management—fostering knowledge creation and innovation by continuous learning to replenish and renew its stocks of knowledge. Data-mining and knowledge-discovery tools play an important supporting role (figure 12).

The new challenge for knowledge management will be reinventing the organization as a provider of products and services that are only possible because it is able to leverage the collective knowledge of its people.

Although we cannot yet imagine the complete set of technology packages that will be required to enable this new state, we are confident that AI technology will play a role.

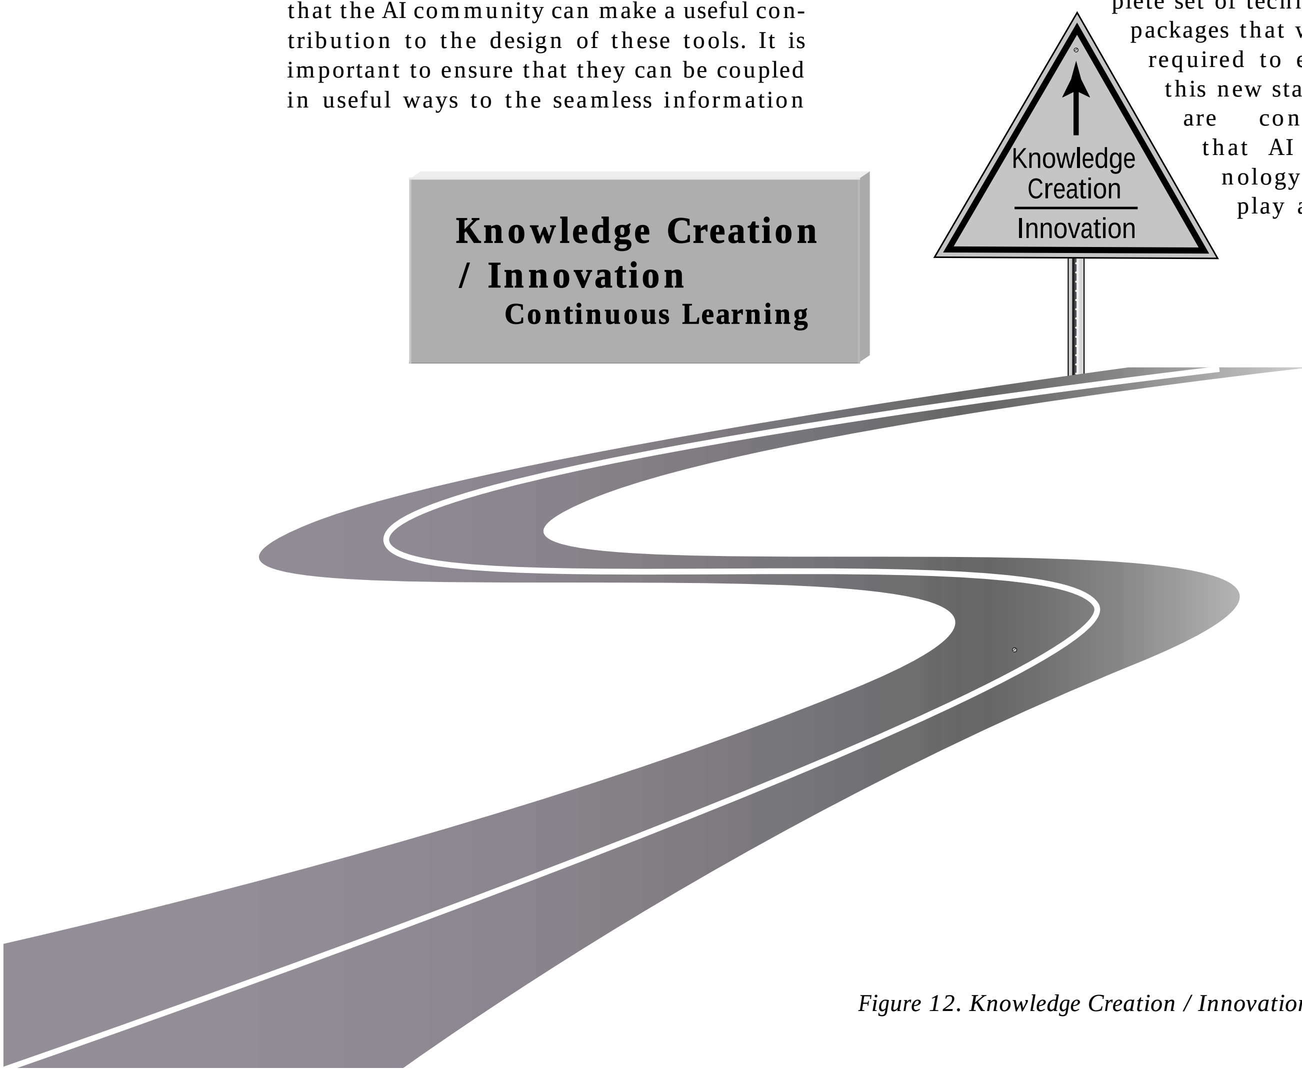

Figure 12. Knowledge Creation / Innovation.

Driving down the road ahead is not a leisurely experience—it is a race.

As a simple example, observing what CoP members are searching for—but not finding—could be used to infer areas where the knowledge stocks are weak, and new research projects and/or alliances are required. This intelligence would be valuable to the knowledge champions in the Example: Drilling Best Practices section and to the line management.

Summary

We attempted to summarize existing practice in knowledge management and suggest possible scenarios for the way that the enabling technology and engineering might evolve over the next decade or so. We have done this by presenting a technology road map and vignette future scenarios for technology that has the potential to enable a new era in organizational performance. It is envisaged that the enterprise of the future will be knowledge powered. Proactive, real-time knowledge management will become the norm.

Such a summary must inevitably be incomplete and selective, but it is hoped that the major targets have been identified and that ways in which the AI community, in particular, can contribute to shaping the future of knowledge management are clear. Meeting the challenges described here is beyond the scope of any single organization. It will require synthesis of the work of many contributors—in industry, academia, and government. Hopefully, this article will stimulate further dialog and collaboration among these groups to develop a technical agenda that turns the vision of the knowledge-powered organization into a reality.

Finally, let's remember that driving down the road ahead is not a leisurely experience—it is a race—and the competitors are lined up on the grid.

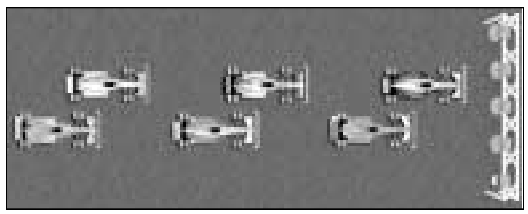

Acknowledgments

We want to thank our colleagues for their insight and support, in particular, Euan Baird, Harry Barrow, Claude Baudoin, Tom Beddingfield, Eric Brown, Pete Edmonds, Henry Edmundson, Ed Feigenbaum, Daniel Guermeur, Jim Kent, Chuck Martin, Amy Price, Claire Vishik, and Stephen Whittaker.

To paraphrase the motto of the worldwide learning program of Holderbank—one of the world's leading suppliers of cement, aggregates, and concrete—our colleagues have *shared with delight*, we have *stolen with pride* (Bock 1998).

Note

1. This article presents only a sketch of the technology forecast aspect of a complete technology road map. To be used by a firm as the basis for a plan of action for product development, the complete technology road map is rather more complicated. It must take into account business context, technology roadmap matrix (time versus product technologies and product characteristics), quality, allocation of resources, patent portfolios, product descriptions, and so on.

References

Baird, D. E. 1997. Knowledge Management and E&P Productivity. Presented at Offshore Europe, 9 September, Aberdeen, Scotland.

Bock, F. 1998. The Intelligent Organization. *Arthur D. Little PRISM,* Second Quarter.

Brickley, D., and Guha, R. V. 2000. Resource Description Framework (RDF) Schema Specification 1.0. W3C Candidate Recommendation 2000. Available at www.w3.org /TR/rdf-schema.

Brown, J. S., and Gray, E. S. 1995. The People Are the Company. *Fast Company* 1:78–82. Available at www.fastcompany.com/online/01/people.html.

Chaudhri, V. K.; Farquhar, A.; Fikes, R.; Karp, P. D.; and Rice, J. P. 1998. OKBC: A Programmatic Foundation for Knowledge Base Interoperability. In Proceedings of the Fifteenth National Conference on Artificial Intelligence, 600–607. Menlo Park, Calif.: American Association for Artificial Intelligence.

Clark, P., Thompson, J., Holmback, H., and Duncan, L. 2000. Exploiting a Thesaurus-Based Semantic Net for Knowledge-Based Search. In Proceedings of the Twelfth Conference on Innovative Applications of Artificial Intelligence, 988–995. Menlo Park, Calif.: American Association for Artificial Intelligence.

Davenport, T. 1996. Knowledge Roles: The CKO and Beyond. *CIO Magazine,* April 1, 1996. Available at www.cio.com/archive /040196_davenport.html.

Dawson, R., 2000. *Developing Knowledge-Based Client Relationships: The Future of Professional Services.* Boston: Butterworth-Heinemann.

Decker, S., Melnik, S., Van Harmelen, F., Fensel, D., Klein, M., Broekstra, J., Erdmann, M., and Horrocks, I. 2000. The Semantic Web: The Roles of XML and RDF. *IEEE Internet Computing* 4(5).

Denning, S. 1999. Using Stories to Spark Organizational Change. Storytelling Foundation International, Jonesborough, Tennessee. Available at www.storytellingfoundation.net/articles/business/denning.htm.

Dixon, N. M. 2000. *Common Knowledge: How Companies Thrive by Sharing What They Know.* Boston: Harvard Business School Press.

du Castel, B. 2000. Intelligence in Artificial Wireless. Paper presented at the Seventeenth National Conference on Artificial Intelligence, 30 July–3 August, Austin, Texas.

Farquhar, A.; Fikes, R.; and Rice, J. 1997. The ontolingua Server: A Tool for Collaborative Ontology Construction. *International Journal of Human-Computer Studies* 46(6): 707–727.

Hagel III, J., and Armstrong, A. G. 1997. *Net Gain: Expanding Markets through Virtual Communities.* Boston: Harvard Business School Press.

Hauswald, E., 1999. Mobil Manufacturing's "Best Practice Networks." Paper presented at Knowledge Management in the Oil and Gas Industry, 15–16 March, London, United Kingdom.

Heflin, J., and Hendler, J., 2000. Dynamic Ontologies on the Web. In Proceedings of the Seventeenth National Conference on Artificial Intelligence, 443–449. Menlo Park, Calif.: American Association for Artificial Intelligence.

KM. 1999. *Knowledge Management Magazine (*Special issue on Portal Mania) 2(7).

LDAP. 2000. LDAP Central Web Site. Available at www.ldapcentral.com/.

Karlenzig, W. 2000. What KM Really Contributes to E-Business. *Knowledge Management Magazine* 3(4): 14.

Lee, J. M.; Grunninger, Y.; Jin, T.; Malone, A.; Tate, G.; and Yost, G. 1998. The PIF Process Interchange Format and Framework. *Knowledge Engineering Review* 13(1): 91–120.

Mitchell, T. M. 1999. Machine Learning and Data Mining. *Communications of the ACM* 42(11): 31–36.

O'Dell, C.; Elliott, S.; and Hubert, C. 2000. *Knowledge Management: A Guide for Your Journey to Best-Practice Processes.* APQC's Passport to Success Series. Houston, Tex.: American Productivity and Quality Center.

O'Dell, C.; Hasanali, F.; Hubert, C.; Lopez, K.; and Raybourn, C. 2000. *Stages of Implementation: A Guide for Your Journey to Knowledge Management Best Practices.* APQC's Passport to Success Series. Houston, Tex.: American Productivity and Quality Center.

Payne, L. W., and Elliott, S. 1997. Knowledge Sharing at Texas Instruments: Turning Best Practices Inside Out. *Knowledge Management in Practice* 6.

Pease, A.; Chaudhri, V. K.; Lehman, F.; and Farquhar, A. 2000. Practical Knowledge Representation and the DARPA High-Performance Knowledge Base Project. Paper presented at the Seventh International Conference on Principles of Knowledge Representation and Reasoning, 11–15 April, Breckenridge, Colorado.

Shevlin, R.; Maney, R.; Sawyer, J.; and Edward, B. 1997. *Vision: Managing Knowledge.* Forrester Leadership Strategies Research Overview. Forrester Research, Inc., Dallas, Texas.

Shilakes, C. C., and Tylman, J. 1998. *Enterprise Information Portals.* Merrill Lynch, New York.

SIA. 1994. The National Technology Roadmap for Semiconductors. Semiconductor Industry Association, San Jose, California.

Smith, R. G., and Maitland, G. C. 1998. The Road Ahead to Real-Time Oil and Gas Reservoir Management. *Transactions of the Institution of Chemical Engineers: Chemical Engineering Research and Design* 76A: 539–552.

Stewart, T. A. 1997. *Intellectual Capital: The New Wealth of Organizations.* London: Nicholas Brealey.

Willyard, C. H., and McClees, C. W. 1987. Motorola's Technology Roadmap Process. *Research Management,* Sept-Oct., pp. 13-19.

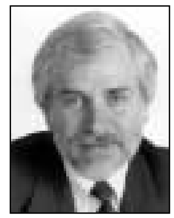

Reid Smith is vice president, knowledge management, at Schlumberger. He is responsible for leading the effort to improve organizational performance through processes and technology to capture, share, and apply the overall experience and know-how of people in the company. Smith has worked for Schlumberger since 1981 in a variety of positions in the United States and the United Kingdom. He received his Ph.D. (1979) in electrical engineering from Stanford University for his work on distributed problem solving. He also holds a B.Eng. (1968) and an M.Eng. (1969), both from Carleton University.

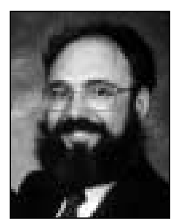

Adam Farquhar is a research scientist at the Schlumberger Austin Technology Center, where he is involved in developing knowledge management and intranet portal technology. Prior to joining Schlumberger in 1998, Farquhar was a research scientist at the Stanford University Knowledge Systems Laboratory, where his research focused on tools and processes for supporting large-scale ontology development and use. Farquhar received his Ph.D. (1993) from the University of Texas at Austin in computer sciences for his work on qualitative simulation and modeling. He also holds a B.A. (1982) in philosophy from St. John's College.