

CHAPTER 1

The Human Touch

The mind is not a vessel to be filled but a fire to be kindled.

—Plutarch, AD 46–120, Greek essayist

Sydney, Australia: Brian, a young programmer who recently started with the company, opens up the contribution form for ToolPool, a global system for sharing technical knowledge. He enters some text describing a program he recently wrote based on his knowledge of a programming language he had learned at the university. His program extends one of the core company products in a smart and unusual way.

Madrid, Spain: Isabel, an experienced consultant, is working on a project at a Spanish bank, where she faces an interesting requirement. She visits ToolPool and after a quick search finds and downloads Brian's program, as it will help fulfill the requirement quickly and elegantly. After using it, she goes back to ToolPool and rates Brian's entry with five stars and adds a comment about how much it helped her.

Cary, North Carolina, United States: Mary, the development manager for the product that Brian extended, scans the Monday morning e-mail from ToolPool, finds Brian's program, and adds a link to the wiki page used for planning the next release of the product.

These are examples of what has been known as knowledge management. Unfortunately, very often the analysis of this situation would now go on to talk about what ToolPool is, what technology it was built on, how much it cost to implement it, and how many information technology (IT) people are currently needed to run it.

But what is really happening here is not that Brian's knowledge is being managed. If anything is managed in this process, it is the flow of Brian's knowledge to other relevant parts of the organization. And ToolPool is only one way that this could have happened. Equally, it could have been that Isabel met Brian at an international technical workgroup and found out about his program.

This book is not about knowledge management technology. It is about ways to influence organizational knowledge flow. Technology does play a role as an enabler, and I mention aspects of it, but the focus is on the human side of making knowledge sharing work. How can you motivate people to share their knowledge, if at all? How can you ensure they will continue to participate? What type of incentives should you use? What are some of the barriers inhibiting the flow that you will have to overcome? What can you do to retain the knowledge that exists only in the minds of those leaving your organization?

While I include examples and case studies from an IT company, many of the principles equally apply to any type or form of organization, whether a government agency, a hospital, or a loose group of physiotherapists exchanging their knowledge in some organized fashion. So the word *organization* is to be seen as wider than a single legal entity or company.

WHY THIS BOOK?

If you currently search for books on knowledge management (KM),¹ you will find a lot of them out there. Amazon.com returns about 16,000 results when you search for the combined term. These books range from highly academic ones to hands-on manuals. So why would you need another one? Why did I even consider writing one with all that coverage out there?

Over the years, I have had many discussions on the topic of knowledge sharing and how to make it work in an organization.

When I started my first knowledge exchange initiative (ToolPool) back in 1997, it was not specifically labeled knowledge management, but after a couple more years and through my involvement with the IBM Institute for Knowledge Management (IKM), it became clear that what we were doing would fit into one of the definitions of KM.

ToolPool, this first initiative, is used as the main case study in the book. It is still running very strongly with global participation internally to SAS (the business analytics company for which I work). At almost 13 years, it might be one of the longest-running KM initiatives. By definition, ToolPool is about technical knowledge, but the principles that make it work are highly nontechnical, as you will see. The contrast of technical and nontechnical elements makes it quite suitable as an example.

ToolPool is only one out of a whole range of different KM initiatives, but it was the major experimentation playground for many years. It was the one to observe, adapt, and analyze. ToolPool has a clear focus on sharing technical tips, tricks, tools, and program code. That topical focus made it easier to deduce learnings and lessons learned from it than from a big-bang all-encompassing knowledge base. As it turns out, this focus is already one of its key success factors. As discussed in Chapter 2, big-bang approaches will have a harder time surviving.

Learning through experimentation was paired with learning through interaction with those responsible in other organizations for getting knowledge to flow. Many insights of what works and what does not came through interaction with others, such as colleagues or those whom I met at external organizations, such as the IKM, the Harvard Learning Innovations Laboratory, or the Babson Working Knowledge program. Key notes from KM pioneers such as Larry Prusak and Tom Davenport influenced my thinking as much as break, lunch, and dinner conversations during those events or with colleagues at SAS.

Most of the time it was not about getting tips but sharing experiences and the reaction and discussion that followed. In one case, I was not sure whether I should add ratings to contributions and discussed some of my thoughts around it. It would have been almost impossible

to really share *everything* I had on my mind about the issue, as it included a considerable amount of context that was tacit and only in my head.

As always with like-minded people, the learning was reciprocal. I shared key information that others would integrate into their context to come to a new level of understanding, and I received feedback that would take my thinking to the next level and help me realize better strategies to approach complex issues. In the case of ratings, I concluded that they would make sense to explore as long as I created the right environment and made them very practical.

I think it is very important to draw the line between knowledge and information. Knowledge is connected to all the prior experiences and exists only in the context of the mind. It cannot be managed. What can be managed are ways to enable the flow of that knowledge to others. What can be passed is information (data in context), not knowledge.

I was inspired to put my ideas on knowledge sharing into a book by those who experienced the passion that I developed for the topic, whenever I got into one of the frequent discussions around it.² And some specifically suggested sharing my recent ideas around the flow of knowledge to a wider audience by publishing them.

Adding to my motivation to write this book was the realization that KM as an organizational discipline has been around for almost two decades, is still acknowledged to be a key factor for organizational success in the future, but often just does not work. A pattern seems to be that those driving it are sometimes doing so based on incorrect assumptions.

Reading and scanning books, articles, or just Twitter messages on KM, I also felt a growing frustration that too much of what has been written focuses on technology. Typically authors talk about “KM Systems” as if you could manage knowledge in a system. And if you look closely, authors frequently mix the words *knowledge* and *information* as if they were synonyms. Often it seems that technology is the only part the authors really understand very well. Other topics that are much closer related to humans are largely neglected. While the author might acknowledge that humans play a central role in KM, often the focus remains on technology.

It is a little bit as if you have a hammer in your hand and then suddenly everything looks like a nail to you. But without understanding and acknowledging the basic difference between the concepts of knowledge and information, you are very likely to use a hammer when you need a razor-sharp knife. In the best case, you just do not get the full benefits from your initiative; in the worst case, you are actually wasting huge amounts of money doing so.

So what is different about this book? I am raising the question of why, with so much expertise (and thousands of books) on the topic, there are still many organizations that struggle with KM. Why can they not make it work? Why is KM not embedded into the everyday practice of every organization if it is so strategic? Why does almost everybody that I talk to tell me that their organization is struggling with making use of existing knowledge? Why is it so hard? I am not claiming that at SAS we have solved all these issues, but we are definitely ahead of the game in many respects.

This book provides some answers to those questions. It might not give you solutions to all potential problems, but it will provide some important reasons why KM might have not worked in your organization and help you with some proven ideas that will make success a lot more likely going forward. According to the Economist Intelligence Unit,³ knowledge management is one of the five key trends that will determine competitiveness in the coming decade. The other four trends are globalization, demographics, atomization, and personalization.

Some of the ideas and lessons presented here might prove priceless, as they will help you avoid some simple traps and focus on elements to improve the organizational knowledge flow that you might not have thought of or tried in the past.

The remainder of this chapter sets the stage by introducing some terms and basic principles to be discussed later in the book.

I do not provide an extensive set of models or research. Enough books out there cover that.⁴ The next chapters contain pragmatic tips and tricks extracted from real-life experiences. The information comes from the front, where initiatives really worked and produced extensive value. The stories and examples presented here come from

initiatives that survived the critical starting stages and are continuing to prove themselves after more than a decade.

To set a base-level understanding, I start out with a short discussion of the term *knowledge management*. You will notice that the title of the book contains the term *Knowledge Flow* instead of *Knowledge Management*. I strongly believe that one of the main reasons why KM projects fail is somewhat due to the use of the term knowledge management and the misunderstanding it creates in the mind of stakeholders. The approaches that people take are often guided, or should I say misguided, by starting out with the wrong frame of mind.

Who is the intended audience for this book? For one, it is aimed at those who have been challenged to bring a new organizational knowledge management program to success or revive an existing underperforming one. The stakeholders might be from IT, from a human resources function, from a business unit, or in a strategic role already focused on knowledge like a chief knowledge officer. They might be line function or sponsoring stakeholders like a chief information officer or the head of personnel.

Because executive buy-in and leadership is a major success factor in driving organizational knowledge flows, it is also important that chief executive officers (CEOs) have the proper understanding as they get involved with strategies. After all, the CEOs are the ones who put the topic high on the future agendas of their organizations.

Knowledge management in the current understanding is often seen as a very technical, software-oriented area, and some people see it as relevant for high-tech organizations only, exclusively for those knowledge workers who spend most of their time online.

With the wider view I am taking, I claim that managing knowledge flows is something that can be applied and used in almost any type of organization. If you detach yourself from the idea that it is about storing “knowledge” in a database, you will see that it is applicable to you, even if you work in an environment that sees itself as being highly nontechnical. Some principles will even work for a group of physiotherapists sharing their experiences in various ways, such as in workshops, expert circles, and online forums.

TERMINOLOGY AND DEFINITIONS

About 15 years ago, the term *knowledge management* was starting to be used in organizational environments, and although I had been dealing with activities that would fit a number of the many KM definitions, we did not call it that at the very start. The first initiatives around exchange of knowledge at SAS were dubbed “supporting sharing of what people know”; we did not use the term *knowledge management* until about 1998. Out in the industry the term had quickly been adopted by management consultants and certain software organizations. Suddenly a “database” was a “knowledge base,” and any product only remotely connected to helping with influencing the flow of knowledge was given that new “cool” label.

But the hype created a number of issues with the term, and in the end the term got “burned” to a certain degree. The problem was that knowledge is a wide concept, so it was easy to drop anything into it. But using a term for something too unspecific has a number of effects—for instance:

- Everybody makes up their own definition of it.
- It ends up encompassing elements that were never meant to be covered.
- It creates a wrong sense of understanding, and people will use unsuitable approaches to solve issues connected with it.

This is precisely what happened with the term *knowledge management*. So let us have a look at the term in more detail.

First, there is a problem with those two words in combination. If you take a puristic view, it describes something impossible. As Larry Prusak and other KM experts have pointed out, knowledge is actually connected to people, it cannot be managed outside of people’s heads. It exists only in the context of prior human experience. So correctly spoken, it is not possible to “manage” that knowledge.⁵

Second, knowledge is actually tacit (implicit) by nature. Nonaka and Takeuchi had talked in their SECI model about ways of externalizing knowledge,⁶ but still, once it is outside of people’s heads, it is mere information, not knowledge anymore. It actually needs another human being to interpret, internalize, connect, and apply it to actually

become knowledge again. Along those same lines, it is not possible to “transfer knowledge,” at least not in the direct sense of transferring an entity from one person to another. What actually happens is that person A shares some information, which is then used by person B and combined with prior (tacit) knowledge and experiences to *create* new knowledge. The knowledge is never transferred directly, as that would indicate that it is moving unmodified. It will always change, however. The knowledge that person A had while sharing information and experiences might be somewhat similar to what person B recreates out of that shared information, but it will always be different, because the framework and the context of prior knowledge and experiences will be different. The word *transfer* indicates the movement of an entity, but that is definitely not what is happening.

A third situation where the word *knowledge* is actually out of place is in the connection with systems, or knowledge bases. The use of the word in that context seems to indicate that knowledge can be stored outside of humans, for example, in a computer system.

One could argue the difference is marginal, but in my mind the fact that knowledge is often seen as an entity that is external to human beings is the number-one reason that so-called knowledge management projects have failed. It easily leads to people using the terms *knowledge*, *information*, and sometimes even *data* as synonyms.

Based on the previous discussions, you can very easily spot articles or books that talk about knowledge management without the proper understanding. I usually stop reading once I find that authors are mixing the terms *information* and *knowledge* as if they were the same thing. For me that is a clear indication that they do not understand what knowledge is. Try it for yourself: The next time a proclaimed KM expert mixes the terms interchangeably in the introduction to an article, a blog entry, or even a book, I advise you to be careful with the rest of it.

Along the same lines, terms like *knowledge management software* or *knowledge management vendor* are somewhat dangerous. Yes, in the holistic process of managing the knowledge flow, certain tools (be it software or other) play an important enabler role, but to say that you can manage knowledge with software would be similar to talking about *motivation software* or a *motivation vendor*.

So what about Business Intelligence and Business Analytics? Since Business Intelligence and Business Analytics are core offerings of SAS, I have looked into that question for quite some time. Personally I do not regard Business Intelligence or Business Analytics as being part of knowledge flow management. I see those at the feeding end of the knowledge flow. They are technologies that are important in today's organizations in the knowledge discovery and knowledge creation phase. They provide the basis for individuals to develop the type of knowledge that is worth flowing through the organization. Technologies that are getting growing attention include data and text mining. Those are not only increasingly important in the knowledge discovery field but are also used for building ontologies and categorizations, which can be helpful for information structuring. These are all important related topics, but for the sake of success on the key of knowledge flow I would not include them into knowledge flow management.

I hope you have followed my line of thinking so far. I acknowledge that there are alternative ways of looking at these issues. Products and solutions being offered by "KM vendors" can provide considerable value. But they are not managing knowledge. They are enablers to the knowledge flow. The information they process, store, and provide can be used to create new knowledge. Information stored in systems and repositories can be seen as representing "pointers to the one who knows." If those using them do understand it in that way, they will be much more likely to actually go beyond the system and see the value of the knowledge that is behind that information, connected to the human who contributed the "pointer."

Many people do not get the subtle difference and immediately jump to the conclusion that you can manage the actual knowledge in those systems. As a result, the selected approaches to drive the initiatives around them (if there is more than a system) are insufficient and likely to fail. To reduce this danger, I recommend moving away from the idea of "managing knowledge" when we are actually talking about managing its *flow*.

Some of the pioneers of KM have for quite some time disputed the use of the term *knowledge management*.⁷ But for lack of a better alternative and because it is widely used by consultants and "KM software vendors," the term has persisted. It still represents a possible

answer to that one big problem that almost all organizations would like to solve: How can I make the best use of the knowledge (present, past, and future) in the heads of the people in my organization?

After a long time of playing with alternative terms, the one that actually fits best with my understanding is *knowledge flow management*, because the thing that you *can* manage is the flow of knowledge. You can speed it up by providing tools and technology as a foundation. You can enable a flow by creating an environment that people find safe, attractive, and efficient and that motivates them to share their knowledge. This could be either face-to-face or by recording relevant information that can be used by others to re-create knowledge in their frame of reference. The flow can be influenced with the help of certain individuals and their actions and behaviors. Chapter 3 discusses the different roles that those individuals will have to play.

The term *knowledge management* will probably stick for a while longer, but if you want to get to the heart of the problem, I advise that you also start using new language. The focus of knowledge flow management is different, and the investments will be different. And as resources are inherently scarce, putting resources into the wrong activities can be a major reason for failure.⁸

TAKING A HOLISTIC VIEW

In 1998 I attended an early KM conference at the Chicago Pier Conference Center. Organized by DCI, the conference was packed with presentations, many of which had the term *knowledge management* in the title. Most of the presentations started with “KM is 80 percent people and 20 percent technology,” and then presenters went on to talk 100 percent about technology. But not all followed that pattern. There were a few exceptions, such as a presentation by Etienne Wenger talking about communities of practice that was the big eye-opener for me in discovering what KM is really all about.

The worst of the presentations was a keynote by a high-level Lotus executive that began with “I am not going to talk about technology,” only to follow five minutes later with a video that turned out to be a Lotus Notes technology commercial. I went home largely disappointed by the conference in general but exhilarated by Etienne’s presentation.

He sparked something in my mind that grew to a major fire over the last decade.

The formula in Chicago was usually:

$$\text{KM} = \text{People} + \text{Technology}$$

And some went as far as extending it to:

$$\text{KM} = \text{People} + \text{Process} + \text{Technology}$$

When I was driving the topic within SAS, I found that it made a great target for three-ball juggling. As it happened, I was asked by the organizers of our company sales kick-off meeting if I would dare to put up a little entertainment during that event. My response was “Sure, as long as I can pick any topic I like to talk about while I am doing it.” So I created a short juggling routine that was essentially a presentation on KM concepts the way that I understood them. And it all started out with juggling three balls representing the key ingredients: people, process, and technology. It was easy to show visually where things can go wrong. I got great response from the audience, and what made me especially happy was that people remembered more than the tricks. After the show, I was able to dive deeper on the topic with a number of people who approached me, and as a result I definitely increased awareness. A lack of awareness of the importance of sharing knowledge is one of the top barriers, as discussed in Chapter 6.

I repeated the kick-off juggling presentation for several years, with a different focus every time. It was a little bit like a yearly live-blog entry for my colleagues. After a couple of years I added a fourth ball, which represented “culture,” as I realized that without the right culture, the other three elements are very hard to bring together.

In the juggling routine, the balls representing the elements are basically all equal in size, but I do point out that the focus needs to be disproportional. Culture and people account for about 70 percent, process 20 percent, and technology 10 percent. One thing to clarify here is that those percentages represent the main effort that goes into the initiative to get it started successfully. Once everything is fully up and running and integrated into your organization, the proportions that those components play may shift. For example, technology might

play a bigger role. But when you start out, the emphasis should be on the human side.

Technology is often the easy part, and that is why too often the efforts are focused on it. “We will deal with the other things later” is the common thinking. Chapter 7 discusses this *technology trap* in more detail.

If you want to take only one thing from this book, then let it be the fact that successful knowledge flow management needs a holistic approach and that working on the hard stuff (the people issues) will be your most important task. By *holistic approach* I mean an approach that covers all elements according to their final impact on success. Many people still believe that technology is the biggest portion of that equation. But with many organizations considering that their KM initiatives have failed, it seems clear that the other two probably did not get the attention they deserved.

What do I mean when I talk about focus on people? Especially within the area of knowledge flow management, it is essential that people are fully involved, motivated, and prepared to share high-potential information based on knowledge they built. Otherwise, any initiative that attempts to enable knowledge to flow through the organization will provide only a fraction of the potential value. And to take it to the extreme, every cent invested in technology could be a pure waste of money. Without proper real attention to the other components for success, you might as well invest the money into something more worthwhile. In fact, if you get presented with any proposal for a “KM” initiative that will spend 95 percent on technology and only 5 percent on (ongoing) costs for support functions (and this includes not only technical support, but marketing, training, strategy, etc.), you might as well take the 95 percent sum and donate it to a good cause of your choice.

Within the holistic view, there is another way of looking at humans and technology. It is the degree of automation that you should strive to create. This is where process comes into play. Process is the piece that brings the other two elements together. In my mind, the degree of automation (the human and technology pieces) are two parts on a continuum, the technology–human continuum (see Exhibit I.1). And the big question is: Where is the cut-off point? How much technology




Exhibit 1.1 Technology–Human Continuum

should we have, how much process should be automatically embedded in the technology, and how much do we leave it to people and their ability to create, follow, and adapt processes? The cut-off point moves frequently. With new technologies, the cut-off point moves to the right, but with the greater sophistication of people and increased complexity and expectations, it moves back to the left.

One great example of a movement of the cut-off point back to the left is the way that people share knowledge via some of the social media tools that came with the Web 2.0 wave. Twenty years ago, everybody was talking about technology agents that deliver all the information to us automatically. We would not have to do anything other than turn on our computer and ask some agent any question in natural language. Contrary to the predictions, I do not think many organizations are even getting close to that scenario. The more sophisticated ones have business intelligence technology in place that will automatically highlight key events and trigger automatic pushing of information, but when it comes to a lot of our daily questions in our professional and personal lives, these days people ask their peers, call an expert, or post a question in some social community like Twitter.

Basically they are turning to “human” agents. The big difference is the scale at which we have those human agents at our fingertips these days. Space and time are not really limitations anymore. I can post a question on a Twitter stream and might get responses in minutes from literally any corner of our planet.

Following blogs can also be seen as using a human agent. Instead of scanning the Web for news and tips on photography, I can follow two or three expert photography bloggers. Apart from being experts, they also spend the majority of their professional or personal lives scanning the Web and consolidating, collecting, and positioning what might be relevant to me. As long as I trust those bloggers and the

sources they consolidate, this process is a lot more effective than if I tried to spend the time myself.

So instead of taking a “me-central” view of my computer and some electronic agents that go out and try to make sense of the information that is available, by following blogs I am using human agents who each represent an entry point into a larger community. That community practices what I am interested in and represents a body of knowledge on that topic.

This is a great example of an interesting shift in the technology–human continuum. The turn to trusted human agents is also something that plays a major role in communities of practice.⁹

One other reason for using a holistic approach to manage the organizational knowledge flow is what Peter Senge calls system thinking. In his famous book *The Fifth Discipline*, Senge introduces system thinking as the most important discipline that lets you look beyond snapshots or isolated parts of your system.¹⁰ For success, you must look at the different components that make up your knowledge flow and how they influence each other. Concentrating on one of them (i.e., technology) alone will be highly insufficient.

But even under the holistic view, the framework needs to be open and simple. So it is not about prescribing everything to the lowest level with extensive and perfect process descriptions. You should create simple rules that are easy to understand and follow, and within those rules you provide degrees of freedom to get innovative adaptation. The rules themselves need to be followed consistently. But within the boundaries of those rules, groups and individuals have considerable flexibility.

Some powerful examples that follow this successful pattern are:

- **The Toyota Production System.** Individual workers have the freedom in their small groups to change processes very quickly and autonomously, but they operate in a strict framework of conduct, collaboration, and feedback procedures to make sure that successful processes travel to other parts of the organization.¹¹
- **The Web.** The basic rules for the Web were simple. The addressing scheme and cross-linking functionality built the

basic framework to an unprecedented technological and social revolution that in its global reach has surpassed anything like it before. Within those simple rules, individuals and groups have the possibility to produce any type of content they can think of. Not all of it is desirable or legal, but those are the side effects that a dynamic system is likely to produce that need to be dealt with.

- **Wikipedia.** As a subset of the Web, Wikipedia conquered a large portion of the encyclopedic market based on some simple principles. Anybody, including nontechnical people, can define terms via simple-to-edit Web pages. The set of rules to follow started out extremely simple and needed to be refined and made a bit more sophisticated. But Wikipedia produced a way for the masses to participate by defining any term that comes to their minds, and for many it was their first encounter with Web 2.0.
- **Twitter.** This microblogging service is showing extraordinary growth and makes simplicity the main mantra. Limiting each post to 140 characters forces people to be concise. Twitter offers extremely easy ways to knit networks among truly global participants via that little “Follow” button. The rules are simple, the scaling is large, and the effects are amazing. It appeals to people interested in any imaginable topic, whether it is the exchange of technical information or swimming tips. As with the Web in general, there are certain behaviors, such as spamming, that might create serious challenges to the system, but the community will counter those behaviors if they still see value in the system.
- **The iPhone.** The principle behind the iPhone is a simple, attractive, and appealing interface that provides a framework for a range of applications. The device combines a strong brand and excellent marketing with innovative and appealing technology.

The actual functionality of the basic device was simple, and users are provided only a handful of applications to start with. But Apple made it very easy to obtain as many additional applications as you like. The key was involving thousands of external developers to create any type of smart or not-so-smart

application based on a common set of rules (programming standards). Another key was a platform to easily share and sell them to millions of users. And Apple did not start from scratch, as there were already millions of users of iTunes before the iPhone hit the market. Because iPods had been around for a while, managing media in iTunes is something that many from children to adults is quite familiar with.

As typical prices for the phone applications are only a few dollars, the hurdle to obtain one of them is low, but good ones might well be bought by a million users. This represents the scaling effect that I discuss further in Chapter 5, because understanding this type of pattern can also help with driving your knowledge flow management initiatives.

I want to say a few words on Web 2.0 or what it usually stands for. Some of the technologies and social implications that were introduced by the Web 2.0 era have really helped with the recent comeback of knowledge management, as they represent an easy way to get in personal contact with others to share knowledge. But even in a Web 2.0 or 3.0 or N.0 world, a lot of the principles introduced in the following chapters still hold. This book is specifically not bound to a given technology but touches on the collaboration elements of social media that are great enablers for person-to-person interaction and introduce unprecedented degrees of scaling. Without some guiding principles, some passionate drivership, and some sponsorship, though, social media implementations often do not live up to their potential, especially in organizational settings.

One of the common misconceptions is that successful Web applications are done in a build-it-and-they-will-come fashion, which neglects the very strong elements of strategy and passionate drivership behind them. Without ongoing strategy and care by dedicated (in any sense of the word) supporters, most Web applications would not have been able to reach their current success level.

To acknowledge the growing role that social media and other Web X.0-type technologies play in organizational knowledge flows, Chapter 9 discusses how knowledge flow management applies under those rapidly changing environments.

As many knowledge management projects focus too much on the technology, the holistic view will need to add the nontechnology elements and specifically human and motivational aspects. Chapter 4 discusses human support and drivership components.

GETTING INTO THE FLOW

When I started playing with the notion of knowledge flow, the analogy of knowledge flowing through the organization like a river flowing through its bed seemed to fit for a number of reasons. Flows find their own way, but they can also be guided and stopped by barriers. You can have some individuals steering the direction of the flow on a daily level and others providing the main bed of the river by setting strategic goals for the longer run. Connected to those strategic goals, you need some metrics that will drive initiatives toward reaching those goals. Chapter 8 spends some time on the topics of setting realistic goals, measuring success, and ongoing analytics needed to steer an initiative to success. I also discuss some of the limitations and misconceptions of what can be measured based on my experience.

Even at the start of our KM program, I looked at driving success from two sides: the active side with directed actions and the passive side where you remove barriers that prevent knowledge sharing from happening.

One good early KM study came from the Fraunhofer Institute in Germany.¹² It asked roughly 400 organizations about knowledge management, and one result was a list of the top barriers for KM that people encountered. Removing those blocking stones from the flow of knowledge within an organization often can be more effective than trying to influence people or processes directly.

As knowledge is connected to humans, it is up to them to decide whether they want to share it. Some people actually think that they can *make* them do it, but as David Gurteen pointed out in a video interview,¹³ this is a fundamental flaw in thinking about knowledge. And as Chris Riemer discussed in a recent editorial for the *K-Street Directions* newsletter,¹⁴ people generally enjoy passing on what they know, so if they are not sharing, it is mostly because something is hindering them. As a result, managing the flow is just as much about

creating conditions that will make sharing more likely as it is about trying to have a direct influence on people.

The good news is that it does not take a rocket scientist to remove some of the barriers, but you will need to know what they are and tackle them with the right approach. In Chapter 6 I discuss the major barriers, position them properly, and suggest some approaches to remove them.

CASE STUDY: TOOLPOOL

At the start of this chapter, I mentioned a specific initiative named ToolPool that I use as the major case study throughout the book to show what a successful knowledge flow initiative might look like. I give some examples and share some stories that should clarify what factors were key to its success. To set the stage, let me quickly introduce ToolPool to you and share some of its history.

When I joined SAS Institute, I worked in a development/consulting combination role at its European headquarters. As part of my position, I worked with a range of consultants locally in the different European offices. One thing that occurred to me while I spent time on various projects at customer sites with my colleagues was a certain degree of overlap in the tools, technologies, and approaches they used. A number of the local consultants were aware of this and had built small repositories of things they could reuse from customer to customer. But the degree to which those reusable assets were shared across the offices was largely based on coincidence. Sometimes I was the one who told people about assets I had seen elsewhere.

So after looking at this for a little while and also finding some of those collections appearing on our growing intranet, I decided to look into ways that we could improve that situation. As I mentioned, there were some collections of tools¹⁵ out there already. One of those collections was a toolbox with a small set of reusable, standardized plug-and-play tools.

Looking at all of this in combination, I ended up with a proposal that I took to senior management. The key points of the proposal were these:

- Tools can range in:
 - Size, from a few lines of text to an application with thousands of lines of programming code.
 - Quality, from a raw concept to a plug-n-play component.
 - Location, from staying local with the author to kept in a central place (pool).

Instead of choosing to focus on certain types of tools, the recommendation was to cater to all of those but use proper categorization. My leading motto was “Don’t discriminate, but categorize,” which differed from the idea of a standardized toolbox that focused on plug-n-play tools only.

- An initiative to support the sharing of tools had to:
 - Be available to everybody within the organization in a very simple way (I had our global intranet in mind here) and without additional cost that could inhibit or exclude any local staff.
 - Have a dedicated support person over an extended time frame. I immediately asked for at least one extra person, as I knew it would be tricky for me personally to stick to only that one initiative over time.

There were additional parts to the proposal, but I will hit on those factors in the more general discussions where they apply in later chapters.

I got the go-ahead to implement what I had proposed in May 1997. By July 11, ToolPool, as I named it, was ready to be launched. It started out with a very simple Web application that basically represented a registry with a descriptive Web page per tool that followed a common basic structure of information provided. Between the end of May and July 11, I spent some time building the simple Web interface for that registry. But considerably more time was spent building up a launch collection of entries. On launch day, I had a base of 150 entries, a number of which came from the collections that were out there on the Web and the plug-n-play toolbox. When people came to ToolPool for the first time, there was something to find.

The audience to which I launched ToolPool was mainly some of the consultants who had provided the first tools as well as those on

specific e-mail lists. It was actually a comparatively small audience to start with. In August I hired the owning support person, and together we spent considerable time ensuring that there was a constant flow of new tools going in. On average we managed to publish about five to six entries per week in those early days.

We found new tools using several ways. Early on, we contacted consultants who had indicated in a mailing list that they might have candidates. We also actively scanned the intranet for candidates, but more and more, as people got to know ToolPool, we received unsolicited contributions from the field. Download numbers in those first months rose quickly to about 300 to 400. A key element of the application that we built in and constantly adjusted to feedback was an analytics component that would give us some feeling how much usage we had and where it came from, something we extended considerably over the years and that I explain more about in Chapter 8.

Year after year, ToolPool has been growing in a number of ways. Being on our internal Web, it quickly went to global usage. The number and range of contributions increased based on marketing we did to other internal audiences. The download numbers increased over the years and are currently between 75,000 and 80,000 downloads per year (i.e., 6,500/month). SAS employs 11,000 people, of whom I would consider maybe 3,000 to 4,000 employees to be prime candidates for reusing a ToolPool entry.

But what is also important is that ToolPool is regarded as *the* place to go for these types of tools. It is an established brand within the technical community of SAS, which includes field-facing programmers, technical support personal, as well as product developers in our development facilities around the globe.

It has seen and survived organizational restructurings and several major software releases, which also represents the evolution at SAS from primarily providing technologies and tools toward the range of customer solutions offered these days. In some cases ToolPool has been a key factor in being able to make those types of shifts.

SAS reinvests roughly 25 percent of its revenue back into research and development, so there is of course a good breeding ground for technical knowledge. What we did with ToolPool was to ensure that knowledge can flow not only within local organizations

but across the whole SAS enterprise, which spans over 400 offices in 60 countries.

ToolPool is not only synonymous for reusing technical components; it also drove other types of knowledge-sharing effects, among them an easier way of identifying those with specific knowledge and also helping some people to be recognized as subject matter experts in the organization.

Certain ToolPool entries turned into small open source communities in themselves connecting those who had a special need and interest in a specific technology or solution. And last but not least, there are numerous examples where ToolPool entries have ended up in a product or at least influenced product development, up to the point that locally developed intellectual property would make up a considerable part of a new solution. In these cases it not only saved development costs but also cut the time to market.

ToolPool continues to be a success story even after almost 13 years in use (at the time of publication). What I learned from it over these years and from other initiatives that I started with the ToolPool lessons in mind served as the basis for my recommendation on how you can raise your organizational knowledge flow to a master level.

Since ToolPool is only one out of several successful initiatives at SAS, in later chapters I introduce other examples, such as our skills management and global resource-sharing initiatives.

Based on the case studies, the remainder of the book shows how it is more successful to drive the flow of knowledge in an organization using a portfolio of methods and with a wider focus that covers a lot more than technology.

NOTES

1. *Knowledge management* is a term I use whenever I am referring to the external and current notion of it. See more in the terminology section later in this chapter.
2. Some of my friends and family can attest to the fact that it is easy to get me started but sometimes hard to stop me when I get into the topic. And it does not matter so much whom I talk to, as many are dealing with problems that are grounded in the fact that knowledge is not flowing the way it should be. Meet me at a party and you have a good chance of getting into some type of discussion regarding sharing knowledge and how it might relate to your environment.
3. Economist Intelligence Study Foresight 2020; see www.eiu.com/site_info.asp?info_name=eiu_Cisco_Foresight_2020&rf=0

4. Appendix B lists some recommended reading.
5. Larry Prusak made this the topic of a number of his keynotes at the Institute for Knowledge Management. The difference between information and knowledge has also been pointed out by T. D. Wilson, "The Nonsense of KM," *Information Research* 8 (No. 1) (October 2002).
6. Ikujiro Nonaka and Hirotaka Takeuchi, *The Knowledge Creating Company: How Japanese Companies Create the Dynamics of Innovation* (New York: Oxford University Press, 1995), p. 284.
7. See the interview with Larry Prusak and Dave Snowden done by Patrick Lambe, "Is KM Dead?": www.gurteen.com/gurteen/gurteen.nsf/id/km-dead-lambe.
8. In this book I still use the term *knowledge management* for cases where it is strongly related to the history of the field or represents the current common understanding.
9. For more on communities of practice, see the section on nontechnical tools in Chapter 7.
10. Peter Senge, *The Fifth Discipline* (New York: Doubleday, 1990/2006).
11. See more about the Toyota Production system in Jeffrey Liker's book *The Toyota Way* (New York: McGraw-Hill, 2003).
12. Hans-Jörg Bullinger, Kai Woerner, and Juan Prieto, "Wissensmanagement Heute," Fraunhofer Institut für Arbeitswirtschaft und Organisation, Stuttgart, Germany, 1997.
13. See www.gurteen.com/gurteen/gurteen.nsf/id/L004395/.
14. See www.knowledgestreet.com/About_Us/Directions/May_2009_Directions/may_2009_directions.html#StSmarts.
15. I will use the word *tool* in a wider sense here; that is, it can range from a document that describes on how to do something, a small tip, a small piece of programming code, or a reusable component, up to a full application or a small application that automates a process.