

The Knowledge Management Theory Papers

A series of papers on the theories, processes and practices behind Knowledge Management

A Framework for Characterizing Knowledge Management Methods, Practices, and Technologies

Brian (Bo) Newman
bo.newman@km-forum.org
Kurt W. Conrad
conrad@SagebrushGroup.com

First Publication, January 1999, in support of *The Introduction to Knowledge Management*, George Washington University Course EMGT 298.T1, Spring 1999.

Subsequent Publications and Presentations based on this paper,

February 1999, panel paper - *Choosing Knowledge Management Technology* Panel, Documation '99, Santa Clara, California.

November 1999, *Characterizing Knowledge Management Tools*, Documation '99 Toronto, Canada

The Knowledge Management Papers are a publication of Knowledge Management Forum. This series is published to promote dialog and peer review of the thoughts and work of Knowledge Management theorists and practitioners striving to bring definition to the practice of Knowledge Management.

Inquiries regarding subscriptions to this series and submissions for publication should be directed to:

TheoryPapers@km-forum.org
or to:
Director of Publications,
The Knowledge Management Forum,
4600 Mallard Court, West Richland, WA 99353.

The thoughts and positions expressed in this and other papers in this series are those of the authors, who maintain and reserve all rights.

Done

Abstract

Knowledge management is not one single discipline. Rather, it is an integration of numerous endeavors and fields of study. This paper provides a framework for characterizing the various tools (methods, practices and technologies) available to knowledge management practitioners. It provides a high-level overview of a number of key terms and concepts, describes the framework, provides examples of how to use it, and explores a variety of potential application areas.

Introduction

Over the past several years, a number of authors have proposed a variety of approaches for classifying the tools (methods, practices and technologies) that typically comprise knowledge management systems. This is not the first attempt to develop a framework for organizing and understanding knowledge management tools.¹ And, given the emerging practices and changing understanding of knowledge management, it will not be the last.

As with any discipline that lacks a recognized unifying paradigm, various views will emerge, each based on what can be readily observed or what can be applied from practices associated with other disciplines. Likewise, as individuals encounter particular phenomena, they tend to describe and interpret them in different ways (Kuhn, 1996).

The following working definition of knowledge management frames the discussion: knowledge management is a discipline that seeks to improve the performance of individuals and organizations by maintaining and leveraging the present and future value of knowledge assets. Knowledge management systems encompass both human and automated activities and their associated artifacts.

From this perspective, knowledge management is not so much a new practice as it is an integrating practice. It offers a framework for balancing the myriad of technologies and approaches that provide value, tying them together into a seamless whole. It helps analysts and designers better address the interests of stakeholders across interrelated knowledge flows and, by doing so, better enables individuals, systems and organizations to exhibit truly intelligent behavior in multiple contexts.

The classification framework presented in this paper can be used in several ways:

- to organize and classify knowledge management methods, practices and technologies by relating them to distinct phases of the targeted knowledge flows
- to examine knowledge flows to understand the interactions and dependencies among pieces of information, communicators and their associated behaviors.

This paper is organized into two sections. The first defines key terms and concepts. The second describes the knowledge management framework, its uses and its benefits.

¹ This characterization framework was initially developed by Mr. Newman for the Introduction to Knowledge Management course taught at George Washington University by Dr. Arthur Murray. More information about this course can be found at <http://www.seas.gwu.edu/faculty/amurray>

Key Terms and Concepts

The characterization framework described in this paper is based on and integrates a number of conceptual models and frameworks. This section introduces those and their related terminology.

Knowledge Flows and Their Associated Activity Areas

There are those who believe that it is impossible to truly manage knowledge, only behaviors. When individuals examine business processes, events and activities, they also tend to use a behavioral focus as the organizing framework. Accordingly, most people find that behaviors are the most comfortable frame of reference for understanding the relationships between business processes and knowledge flows.

Knowledge flows comprise the set of processes, events and activities through which data, information, knowledge and meta-knowledge are transformed from one state to another. To simplify the analysis of knowledge flows, the framework described in this paper is based primarily on the General Knowledge Model. The model organizes knowledge flows into four primary activity areas: knowledge creation, retention, transfer and utilization (Figure 1).




Figure 1. The General Knowledge Model

Knowledge Creation. This comprises activities associated with the entry of new knowledge into the system, and includes knowledge development, discovery and capture.

Knowledge Retention. This includes all activities that preserve knowledge and allow it to remain in the system once introduced. It also includes those activities that maintain the viability of knowledge within the system.

Knowledge Transfer. This refers to activities associated with the flow of knowledge from one party to another. This includes communication, translation, conversion, filtering and rendering.

Knowledge Utilization. This includes the activities and events connected with the application of knowledge to business processes.

The General Knowledge Model sequences the activity areas in a deterministic fashion. In reality, though, all but the most rigorously automated knowledge flows comprise complex systems that are built mostly from asynchronous processes. The model is valuable precisely because it relates the individual, highly dynamic behaviors and processes to general activity areas and, by association, to each other. Various theories of learning, problem solving and cognition may imply specific activity patterns, but they are usually not required to organize the key relationships and dependencies among the activity areas. The model allows analysts to trace individual knowledge flows by helping them to examine and understand how knowledge enables specific actions and decisions.

Within each activity phase exists other, smaller knowledge flows and cycles. These layers span a wide range of macro- and micro-behaviors, ranging from broad organizational and multi-organizational processes to discrete actions and decisions, and include all the various intervening layers: activities, tasks, workflows, systems, interfaces and transformations.

Knowledge Artifacts

Artifacts come in a variety of forms, including documents, files, papers, conversations, pictures, thoughts, software, databases, e-mail messages, data sets, winks and nods, and whatever else can be used to represent meaning and understanding. Said another way: knowledge artifacts flow among and form the linkages between the activities and events that comprise knowledge flows.

Most people's involvement with a knowledge stream is through various artifacts. Artifacts are what we deal with every day. We write reports, send e-mail, read books, remember bits and pieces of old thoughts, engage in conversations and follow procedures.

The term *knowledge artifact* does not specify the form of the artifact (e.g. information, transformation, metadata or meta-knowledge) but it is very specific as to the process that gave rise to the artifact. This makes the term valuable for explaining such things as the importance of knowledge artifact retention, establishing provenance and enabling reusability.

Knowledge artifacts differ from one another in several ways: their form of codification, the way in which they are rendered, their degree of abstraction and their ability to enable actions and decisions.² Knowledge artifacts also vary in their degree of articulation; simple knowledge artifacts can be explicit, implicit or tacit. Most artifacts, however, are not simple but complex, and contain a combination of explicit, implicit and tacit components.

Explicit Knowledge Artifacts. These are knowledge artifacts that have been articulated in such a way that they can be directly and completely transferred from one person to another. This normally means that they have been codified so it is possible to touch, see, hear, feel and manipulate them (e.g. books, reports, data files, newsreels, audio cassettes and other physical forms).

Implicit Knowledge Artifacts. These are knowledge artifacts whose meaning is not explicitly captured, but can be inferred; in effect, the codification process is incomplete. Explicit artifacts can be interpreted totally on their content. Interpreters of implicit artifacts must rely on previously retained knowledge.

For example, the knowledge that a given phrase is a book title tends to be implicit. Rarely is there anything that specifically tells someone that they are reading a book title, as might be the case in an SGML or XML system when <BookTitle> tags explicitly communicate semantic meaning. In most cases, the reader infers the meaning of the

² In many circles, it is still common to refer to the level of abstraction and the potential role of knowledge artifacts by differentiating among data, information, knowledge, understanding and wisdom. While such distinctions may still prove helpful in some cases, problems in definition and interpretation often arise from any attempt to maintain rigid lines of demarcation. These problems can be avoided through the use of the collective term *artifact* without any significant loss in the effectiveness or validity of the framework.

6

Tangible
Form.
Intangible
Content.

words from their position (on the cover of a book), formatting (big, bold and centered) and content (lacking formal subject and predicate).

The potential for ambiguity is one of the characteristics of implied knowledge artifacts. Most readers of the sentence, "Ann put on her heavy coat and locked up her classroom," implicitly understand that it is winter and Ann is a teacher, but there are other inferences that could be made as well. For consistent interpretation, both the person making the statement and the person interpreting it must share some common frame of reference to understand when heavy coats are worn and who locks up classrooms.

The underlying knowledge embedded in processes can also be considered as an implicit artifact. For example, a manual detailing the safe way to handle corrosive materials might include a statement such as "This material should not be used on polished or anodized aluminum services. If swallowed, immediately rinse mouth and drink a glass of milk or water. Do not induce vomiting." The implicit knowledge contained within these warnings, combined with what the reader might recall from high school chemistry, tells the reader that the material is likely to be very caustic.

Implicit knowledge artifacts can also be found in process-specific software. In developing the software, the designers had to conceptualize the processes that the software would be supporting. That knowledge will show in the way the software is intended to be used and in the range of behaviors it directly supports. Even if not explicitly apparent, these implicit knowledge artifacts will effectively constrain users' actions. This is often referred to as implicit policy making by technologists (Conrad, 1995).

Tacit Knowledge Artifacts. These may be the most insidious and powerful of the three. Michael Polanyi referred to tacit knowledge as "knowing more than we can say" (Polanyi 1966). Simply stated, tacit artifacts are those that defy expression and codification.³ This is not to say that tacit knowledge artifacts are without influence. The most vivid example is the old saw about what would happen to the centipede if she were to stop and think about how to walk.

It is important to note that, for the most part, artifacts are passive. While they can change (or, more accurately, be changed), they can't act. Has anybody ever seen a financial report make a decision or a book on aerodynamics build an airplane?

Agents

Knowledge artifacts do not perform actions and make decisions. Actions and decisions are undertaken by agents: people, organizations, or in some cases, technology. Agents carry out all the actions and exhibit all the behaviors within a knowledge flow.

³ The problem of *tacit knowledge*, its acquisition and epistemic status has been the focus of considerable philosophical investigation by such people as Ludwig Wittgenstein, Edmund Husserl, Hilary Putnam and, most significantly, Michael Polanyi.

Often, analysts attempt to apply the same behavioral models to all agents in a system. More appropriately, agents can be placed in three categories:

- Individual agents
- Automated agents
- Organizational agents.

Individual Agents. These agents sit at the center of almost every knowledge flow. For most analysts, the individual (human) serves as the prototypical active force for affecting change. In this paper, the term *individual* is used in the collective sense and is not meant to imply that every specific individual is capable of the full range of behaviors attributed to this class of agent.

Individual agents are capable of working with knowledge and knowledge artifacts in all degrees of abstract articulation. They are limited, however, in their ability to deal with artifacts that are codified in ways that fall outside the range of human perception (radio waves, for example). The individual agent is the only agent capable of performing all aspects of knowledge development, retention, transfer and utilization without the need for intervention by either of the other two agents.

Automated Agents. These agents can include any human construct that is capable of retaining, transferring or transforming knowledge artifacts. They are not exclusively computerized processes, as is often assumed in discussions of knowledge management. A conventional camera that encodes a representation of the visual world through chemical changes to the surface of a film could act as an automated agent, supporting knowledge creation and capture.

Organizational Agents. These agents exist in situations in which knowledge retention and transfer cannot be fully attributed to individuals or specific automated agents. In these cases, the organization itself serves as an agent in the retention and dissemination of knowledge. As with tacit knowledge artifacts, current tools and concepts do not account very well for the roles of organizational agents in knowledge flows.

Organizational value systems provide strong evidence for the existence of organizational agents. Much has been written about the ability of organizations and communities to establish value systems that outlive the involvement of specific individuals and the power that these value systems have to influence the behavior of individuals and groups (Krogh and Roos, 1995; Kuhn, 1996). The principles and practices that make up these value systems are almost never codified.

In fact, when individuals attempt to describe the organization's value system, the descriptions are usually incomplete, reflecting either an interpretation of the organization's values or a blending of organizational and individual values. The common use of the terms *unwritten rules* and *organizational culture* is a reflection of the difficulties involved. The terms acknowledge that organizations are repositories of tacit knowledge.

Individual, organizational and automated agents have different behavioral models. Unlike computerized agents, for example, most individuals don't perform a given task exactly the same the way every time. If human-based knowledge transfer processes are designed to work as software processes do and the designers fail to leave sufficient room for the factor of human variability, the system is unlikely to perform as intended.

Individual and automated agents also differ in their ability to handle implicit knowledge artifacts. For example, the ability of individuals to infer meaning of book titles usually allows them to accept a wide variety of formats and styles and even recognize titles inside streams of text (for example, *The Bible*). Anyone who has built filters to convert documents knows that automated agents are not skilled at supplying context.

Agents also differ in the how well they use tacit knowledge. Individual and organizational agents can handle tacit knowledge, but because automated agents can only deal with codified artifacts, and tacit knowledge by definition defies codification, automated agents seem destined to be unable to follow suit.⁴

The Characterization Framework and How to Use It

The characterization framework is the application of the concepts described above to activities such as those associated with tool selection, development and deployment.

The framework is easy to use when represented as a table (for example, Table 1, below). In this form it allows a given tool to be described in terms of its interactions with the various elements of knowledge flows and their associated subtypes.

Tool				Activity Phase			Activity Level			Agent Type			Artifact Type			Focus				
Name		Practice	Method	Technology	Creation	Retention	Transfer	Utilization	High-level Process	Mid-level Process	Decision or Action	Individual	Automated	Organizational	Explicit	Implicit	Tacit	Agent	Artifact	Process

Table 1. Sample Rendering of the Framework

⁴ More on the ways in which tacit knowledge can be addressed by knowledge management efforts can be found in *The Siamese Twins: Documents and Knowledge* (Newman, 1997).

This is not the only way the framework can be displayed. The framework is a general-purpose tool that can be applied to a variety of problems and solutions and adapted to individual work styles.

For example, you could use the table as a simple guide or checklist to make sure that you have examined a tool or situation from all of the suggested aspects. You might use it to record primary and secondary characteristics when comparing similar tools. Or, you could expand the cells to contain short statements that reflect what you know about the interaction of the tool with the target element. All of these are valid approaches and could be used separately or in combination. This flexibility is intentional; it can be traced back to the framework's theoretical foundations.

The framework's theoretical roots focus on the role of knowledge in complex systems and fundamental knowledge interactions. This focus provides a solid foundation that can be built upon, applied and adapted to different contexts.

So, while this framework can be used to do highly formalized analysis, it also works for simpler, back-of-the-envelope analysis, or even to sort out a couple of facts. It works for engineers and psychologists and can be used to discuss and describe information policies in neutral language that is neither business-centric nor technology-centric.

The important thing is to keep it simple. By using the framework, you will learn more about your problem and make explicit other things you had known implicitly or even tacitly. As new facts become apparent, the tendency might be to start drilling deeper and deeper into a facet of the problem. It will not take long to realize that the deeper you go, the more interrelated that particular facet becomes with the other areas of the framework. Before you know it, the problem will have become very complex.

To avoid this trap, take a high-level look at the problem or situation from all the vantage points offered by the framework. This way you develop a balanced view of the situation and are in a better position to understand the interrelationships that occur as you extend your analysis.

It is also important to note that the effective use of the framework is not dependent on selecting just the right starting point. Whether you choose to start with an examination of the agent, the artifact, the activity phase or the activity level, the interrelated nature of these elements will end up leading you into the rest.

When using the framework in a group setting, or when two people are using it to examine the same situation, it is important to keep in mind that the framework cannot make everyone see a given problem in exactly the same way. If different people or groups use this framework to look at a single event, odds are, they will come up with different results. That does not mean the framework is flawed. What it means is that the different

observers have applied their own experience and personal knowledge⁵ to the interpretation.

The approach of using a single record for all of the elements associated with a specific tool is intended to focus attention on high-level analyses. This is just one application of the framework. If you need to perform more detailed analyses, other application approaches are possible. For example, you might want to construct smaller matrices that contrast individual elements, such as activity phase and activity level and repeat the analysis at each intersection point. This could help clarify the location of critical interactions to better identify targets of opportunity for improving knowledge flows and associated agent performance.

Using the Framework to Classify Knowledge Management Tools

The framework was designed primarily to support tool classification, hence the term *classification framework*. It helps individuals identify and differentiate among the roles different tools can play in a knowledge management system. One of the easiest ways to evaluate a tool is to describe its characteristics in terms of its interactions with each of the element subtypes of the framework.

Once tools have been characterized in this fashion, analysts are in a better position to do the following:

- Relate the various ways that methods, practices and technologies can impact the flow of knowledge within an organization
- Identify if the primary role of a given tool is to manipulate artifacts, influence agent behavior or establish behavioral patterns
- Distinguish between the types of agents whose behaviors will be most influenced by a given tool
- Distinguish the level of organizational behavior the tool will most likely affect.

Using the Framework to Identify Knowledge Flow Elements

Underlying the application of the framework in tool classification is the ability to use the framework to identify and classify the individual elements of knowledge flows and their supporting knowledge management systems (i.e. agents, artifacts and behaviors). In effect, the framework provides a way to subdivide knowledge flows into more manageable components.

It is one thing to subdivide, but it is another to subdivide and maintain meaning. Mapping knowledge flow elements to agents, artifacts and behaviors raises the level of abstraction and, in effect, simplifies the elements. At the same time, care must be taken not to oversimplify and damage usability. The further subtyping of agents, artifacts and behaviors (for example, by activity phase and level) maintains enough richness and

⁵ Prior personal knowledge can be explicit, implicit, or tacit, or a combination of all three. More on the nature of prior retained knowledge can be found in *The Siamese Twins: Documents and Knowledge* (Newman, 1997).

context for the framework to be usable. For most purposes, this level of subtyping maintains a healthy balance between simplicity and sophistication.

Using the Framework to Analyze Knowledge Flows and Identify Gaps

Because the framework illuminates the nature of the relationships among knowledge flow elements, it allows analysts to take the myriad of data points associated with complex knowledge flows and put them into a structure in which the relationships and dependencies become far more apparent (or even explicit). This structuring process, in turn, helps illuminate both patterns and the gaps that result from missing or unarticulated elements.

Often, gaps in understanding drive the phenomenon of overloading, when the distinctions among artifacts, agents and behaviors are blurred and the nature of the relationships becomes confused. With the growth of advanced data standards, such as dynamic HTML, people seem far more likely to confuse artifacts and behavior. When faced with an HTML <Blink> tag, for example, it is not uncommon for people to say that the document blinks, when in reality the interactions, dependencies and transformations are far more complex. The document (artifact) is a repository for instructions (artifact) that conform to an agreed upon specification or protocol (artifact). It is actually a browser (agent) that is responsible for making the text blink (behavior).

Application Areas

Internal Development of Knowledge Management Solutions. The framework can be used in the following ways to support internal development efforts:

- Mapping specific tools and technologies according to their potential roles in knowledge flows
- Identifying functional gaps
- Determining integration points
- Validating the scope of development efforts that seek to extend base technologies with application-specific functionality.

The framework has been used to successfully improve the reliability of collaborative decision-making processes and the quality of resulting decisions. It has helped to illuminate the relationships among new information, known facts, prior leanings and value systems. This has allowed decision making and governance patterns to be identified and translated into both general strategies for improvement and specific designs for decision-making processes and formal governing bodies.

Within a number of projects, the framework has been used to map end-user behaviors to specific metadata requirements and document designs. The framework has proved useful precisely because it focuses attention on the interactions of multiple agents and processes. This helps individuals identify and differentiate the metadata and other knowledge artifacts most appropriate and valuable to each of the ever-increasing number of agents and processes that seek to interact with such artifacts.

Selection of Knowledge Management Products. Like the document management market that preceded it, the market for knowledge management tools and technologies is a confusing one. When used to support market analysis, the framework helps to articulate the organizational context in which the tool will be used and therefore illuminates previously unrecognized gaps. These understandings can then be leveraged in the form of more complete and formal specifications to aid in the selection of tools that best fit the organization as a whole. Finally, the framework helps to illustrate that the true value of a tool results ultimately from the decisions that are made about how to deploy it in the context of specific knowledge flows.

The framework also helps to explain the subtle but noticeable shift away from traditional file management systems to more interactive and highly granular component management systems that support personalization and dynamic content. With increasing demand to support a broader range of context-specific behaviors, information management systems are being asked to provide ever more sophisticated metadata and relationship management services. Such emerging metadata management systems are well suited to providing just the right content to the right person at the right time and customizing artifacts to better enable that person's actions and decisions.

Marketing of Knowledge Management Products. The framework can benefit companies contemplating new offerings, as well as those actively engaged in the marketplace. Because knowledge management serves as an integrating discipline for many existing practices, organizations have had a hard time distinguishing just where both existing and new tools fit into the picture. One of the more common complaints from customers is that vendors are just re-labeling existing products as knowledge management tools with little or no change in the underlying functionality.

For vendors in the knowledge management market, the framework offers a well-grounded way to differentiate products and services. The examples that follow show that existing technologies, such as e-mail, and methods, such as facilitation, have meaningful roles in supporting knowledge flows and are valid pieces of a more comprehensive knowledge management system. This framework gives the vendor a way to describe how their product or service fits within the broader context of the knowledge management solution space. As well, it can help identify strategic opportunities for product evolution and increased customer value.

Examples

Below are two examples of how you can apply the framework to assess how two tools that may not normally be associated with knowledge management could help with knowledge management activities: e-mail and a facilitation method called AtStake.

E-mail

You are looking to improve communications and are exploring e-mail systems. The first question is whether e-mail is a practice, method or technology. And there's no doubt: we are definitely talking about a technology.

Activity Phase — Which Activity Areas Does E-mail Support? E-mail doesn't generally contribute to knowledge creation. It does not matter what kind of editor you are using to draft a message because the primary purpose of the tool is not to help you synthesize new knowledge. Still, you might decide to compare tools in terms of their knowledge capture capabilities.

Likewise, e-mail does not have much to do with knowledge utilization. The real focus of e-mail, as with most office automation tools, is knowledge transfer and, depending on how you use it (for example, whether you keep all of your old messages), possibly retention. A few tools, such as modeling and decision support tools, focus on creation and/or utilization, but most of the software applications associated with "management" (for example, information management, document management and image management) tend to focus on retention, transfer and their associated transformations.

Activity Level — On What Does E-mail Tend to Have the Most Impact? E-mail has an impact on three activity levels:

- On low-level decisions and actions because it is one of the ways (sometimes a primary way) that people engage in one-on-one communication with others, decide priorities, allocate tasks and exchange the small bits of information that drive individual actions
- On mid-level activities because it is not uncommon to see business processes at various project and program levels designed around specific e-mail capabilities and/or specific protocols established for the use of e-mail within the organization
- On high-level business processes because of its well-documented impact on organizational culture, openness, knowledge sharing and structure. For most organizations, the impact of e-mail on strategic processes is fairly low. For businesses with virtual organizations or Internet-based sales and marketing components, however, e-mail is likely to be a critical enabler of core competency.

Agent Type — What Types of Agents Interact with E-mail and How? E-mail tools are automated agents. The primary interfaces tend to be with individuals and not organizations, keeping in mind that organizations cannot type or read. However, e-mail tools can and often do interface with such automated agents as data mining, security, the firewall and a variety of attachment-specific tools.

Artifact Type — How Does E-mail Interact with Each Type of Artifact? E-mail systems inherently accept and reject certain forms of codification and rendering. Some of these codifications represent communication protocols that specify the way that e-mail messages are to be encoded and packaged. An e-mail tool, for example, is not expected to render music. Likewise, e-mail cannot process machine code. In general, e-mail tools only actively interact with textual material and the most complex behaviors are usually associated with a limited set of textual representations.

Focus — Is E-mail Optimized for Interactions with Agents, Artifacts or Processes? Although there is typically quite a bit of interaction with individual and automated

agents, e-mail systems do not direct or influence agent behavior. Instead, most of their functions are associated with the manipulation of e-mail artifacts.

By now your analysis has gone full circle. It started with the type of tool and ended by looking at the impact of the tool on artifacts. You could, of course, start anywhere. By the time you have completed a row in the table, the tool has been examined from the standpoint of process, agents and artifacts. Table 2 shows this analysis in a fully populated rendering of the framework.

AtStake

The framework can also be used to characterize methods. In this example, a stakeholder-focused strategic planning process called AtStake is evaluated using the characterization framework.

As with e-mail, there are many methods (including AtStake) that are not normally considered to be knowledge management tools. However, the activities that comprise an AtStake session, the facilitation approaches that it is based on, and the artifacts that are produced can, be characterized from a KM perspective.

AtStake is considered a method, because it is based on a series of repeatable steps that produce predictable results. Although it is a fairly general tool that can be used in a variety of ways (including conflict resolution and the structuring of negotiations) it is not sophisticated enough or used by enough people to be considered a practice. Also, it does not rely on enough automation to be considered a technology.

Activity Phase — Which Activity Areas Does AtStake Support? AtStake's most significant contribution to knowledge management is knowledge creation. AtStake sessions typically result in the creation of a new, shared understanding among the participants. This shared understanding, in turn, functions as a context for aligning individual behaviors. Knowledge capture is usually done with flip charts, personal notes and memories.

In terms of retention, there does not have to be a strong emphasis on generation and retention of explicit artifacts. In many cases, some form of follow-up documentation is produced. However, often the only form of retention is the tacit knowledge of the participants. The participating organizational agents often retain as tacit artifacts the shared values that are synthesized.

A typical AtStake session involves considerable knowledge transfer among the participants. The facilitation model is designed to amplify and focus such transfer activities through a series of small and large group exercises. Along the same lines, the most apparent form of knowledge utilization occurs in real time within the facilitation process. The new shared understanding also drives knowledge utilization as individual behaviors align while consensus is being reached.

Activity Level — On What Does AtStake Tend to Have the Most Impact? AtStake is often used to provide direction to high-level business processes and contributes to the development of consensus among multiple organizations (and even multiple governments). Many organizations in the governmental, quasi-governmental and private sectors have used it to define high-level processes and organizational structures.

It has also proved to be quite useful for integrating stakeholders' knowledge requirements into the design of mid-level business processes, activities and projects. In this context, AtStake can be used very effectively to define policy parameters, performance objectives and specific action plans.

The AtStake process, especially its facilitation model, is weakest at the level of individual decision and actions. The underlying concepts can be and are used to enable decisions and actions, but the process as a whole is not designed for this.

Agent Type — What Types of Agents Interact with AtStake and How? Individual agents are the participants that create, retain, transfer and act upon the knowledge flows associated with an AtStake session. No automated agents of any consequence are involved with the process; standard office automation software can be used in preparation and documentation, but it is not required.

Organizational agents are certainly involved because the primary function of AtStake is to help groups to think collaboratively and produce a tangible product, if needed. Also, the impact of an AtStake session is usually felt most directly at the organizational level. One of the primary outcomes is a creation of a shared reality that mobilizes and provides focus to an organization's diverse (and sometimes autonomous) components.

The role of the organizational agent is so important to an AtStake session that special consideration should be given to including individuals whose concerns fall clearly outside the scope of the targeted organizational agent(s). These external stakeholders (e.g. customers, clients, information and technology suppliers, regulators, auditors, approval bodies and financiers) must be active participants to ensure that they come to consensus at a point that not only suits them, but is also consistent with the broader community of interests that will ultimately determine the success or failure of the enterprise.

Artifact Type — How Does AtStake Interact with Each Type of Artifact? The most important artifacts associated with the process are not explicit. Participants rarely bring explicit artifacts into the sessions. Flip charts and final reports are explicit, but they are of secondary importance to the process.

While the conversations are explicit, much of their value is derived from the context that is established. Stated another way, the facilitation process frames explicit speech in a way that amplifies its value by using it as a vehicle for transferring implicit and tacit knowledge.

Implicit knowledge is also captured by the facilitator and can be used to interpret the group's flip charts and prepare the final report. The explicit artifacts are used to trigger implicit knowledge about their context and both are re-synthesized in the final report.

Manipulation of tacit artifacts is the primary focus of an AtStake session. The participants bring tacit knowledge to the table. It is then made explicit through conversations that are prioritized and re-codified into a written record. Once it has been transferred through a variety of written and spoken forms, it is internalized as new set of tacit artifacts whose content is the new shared value set.

Focus — Is AtStake Optimized for Interactions with Agents, Artifacts or Processes?

AtStake sessions focus on agents, specifically influencing the behaviors of individual and organizational agents. Secondly, AtStake produces a set of tacit artifacts (the shared value set).

Tool				Activity Phase			Activity Level			Agent Type			Artifact Type			Focus				
Name		Practice	Method	Technology	Creation	Retention	Transfer	Utilization	High-level Process	Mid-level Process	Decision or Action	Individual	Automated	Organizational	Explicit	Implicit	Tacit	Agent	Artifact	Process
E-mail				X		X	X		X	X	X		X		X				X	
AtStake			X		X				X	X		X		X			X	X		

Table 2: A Sample Populated Framework

Conclusion

Selecting knowledge management technologies is often a daunting and risky task. Without an independent frame of reference, attempts to compare knowledge management technologies can be very confusing and fail to drive needed decisions. By providing a means to differentiate technologies according to their impacts on agents, artifacts and behaviors, the characterization framework described in this paper provides just the kind of neutral reference point organizations often need.

The framework also adds value to supporting analytical, design, development and deployment activities by guiding the analysis of knowledge flows and construction of a usefully comprehensive picture. The framework provides a mechanism for developing a balanced, high-level view that can be used to set the stage for deeper analysis, identifying the compelling and critical issues that warrant more careful examination. Once the picture is complete, the framework can be used to identify the specific needs that can be met with off-the-shelf technology, localized customizations or change-management programs.

By using the same framework to relate technologies, methods and practices back to targeted knowledge flows and their associated behavioral goals, it becomes easier to balance technical and non-technical approaches. This allows project teams to take a more rational, whole systems approach to development and deployment, improving their ability to develop tools and approaches that target and resolve root problems and not just symptoms, improve organizational performance and lower overall life cycle risks.

References and Suggested Readings

- _____. (1998). *Harvard Business Review on Knowledge Management*. Boston, Harvard Business School Press.
- Conrad, K. W. (1995) *SGML, HyTime, and Organic Information Management Models*, Second International Conference on the Application of HyTime, August 1995 (<http://www.SagebrushGroup.com/organic.htm>).
- Boisot, M. H. (1998). *Knowledge Assets: Securing Competitive Advantage in the Information Economy*. New York, Oxford University Press.
- Dunlop, C. E. M. and J. H. Fetzer (1993). *The Glossary of Cognitive Science*. New York, Paragon House.
- Fahey, L. and L. Prusak (1998). "The Eleven Deadliest Sins of Knowledge Management." *California Management Review* 40(3): 265-276.
- Jensen, M. C. (1998). *Foundations of Organizational Strategy*. Cambridge, Massachusetts, Harvard University Press.
- Krogh, G. V. (1998). "Care in Knowledge Creation." *California Management Review* 40(3): 133-153.
- Krogh, G.V. and J. Roos (1995), *Organizational Epistemology*. New York, St. Martin's Press.
- Kuhn, T. S. (1996). *The Structure of Scientific Revolutions*. Chicago, The University of Chicago Press.
- Lamberts, K. and D. Shanks (1997). *Knowledge, Concepts, and Categories*. Cambridge, Massachusetts, The MIT Press.
- Leonard, D. and S. Sensiper (1998). "The Role of Tacit Knowledge in Group Innovation." *California Management Review* 40(3): 112-132.
- Moss, E. (1995). *The Grammar of Consciousness: An Exploration of Tacit Knowing*. New York, St. Martin's Press.
- Newman, B. D. (1997). *The Siamese Twins: Documents and Knowledge*. West Richland, WA, The Knowledge Management Forum.
- O'Dell, C. and C. J. Grayson (1998). "If We Only Knew What We Know: Identification and Transfer of Internal Best Practices." *California Management Review* 40(3): 154-174.
- Pinelli, T. E., R. O. Barclay, et al. (1997). *Knowledge Diffusion in the U. S. Aerospace Industry*. Greenwich, Connecticut, Ablex Publishing Corporation.
- Polanyi, M. (1966). *The Tacit Dimension*. Garden City, NY, Doubleday.
- Ruggles, R. (1998). "The State of the Notion: Knowledge Management in Practice." *California Management Review* 40(3): 80-89.
- Taylor, S. (1989). *Conceptions of Institutions and the Theory of Knowledge*. New Brunswick, Transaction Publishers.
- Turner, S. (1994). *The Social Theory of Practices*. Chicago, The University of Chicago Press.
- Wiig, K. M. (1995). *Knowledge Management Methods: Practical Approaches to Managing Knowledge*. Arlington, Texas, Schema Press.

About the Authors

Brian (Bo) Newman has more than 20 years' experience offering services in the areas of knowledge management and project management. As the founder, host and moderator of the internationally recognized Knowledge Management Forum, Mr. Newman has long worked to establish improved models for understanding the ways knowledge is developed, stored, transferred and used within organizations.

Kurt Conrad is one of the original members of the Knowledge Management Forum. As a project architect of enterprise-wide document production systems, he blends knowledge management methodologies with SGML and XML technologies to integrate the needs of multiple stakeholders, resolving the organizational and cultural problems that are common to such initiatives.