

Double-Loop Knowledge Management

A White Paper

By Mark W. McElroy, Principal
IBM Knowledge Management Consulting Practice
August, 1999 (v.3)

Introduction

Knowledge management has now evolved into two distinct, if not *competing*, schools of thought. I, and many others, have begun to differentiate between the two as *first-* and *second-generation KM*, respectively. In summing up the major new themes of *second-generation KM* (SGKM), I recently made the following remarks in another paper of mine entitled, *Second-Generation KM (Knowledge Management Magazine, September, 1999)*:

- Second-generation KM schemes emphasize knowledge production (demand-side thinking) without discounting the importance of first-generation KM codification and sharing (supply-side thinking)—a new more balanced view has emerged;
- This emergent emphasis on *knowledge production* points to a much higher-value proposition for KM than has been proffered to date: *the prospect of increasing an organization's rate of learning, and hence, its rate of innovation*;
- SGKM has identified both the form of organizational knowledge (declarative and procedural rule sets in *knowledge structures*), as well as a process-based life cycle that must be

- nurtured and cared for if healthy organizational learning is to occur;
- Lastly, SGKM convincingly makes the powerful connection between KM and *organizational learning* (OL), out of which comes recognition of the role that KM can play in helping organizations, not just individuals, learn faster and more effectively than their competitors. The value-proposition implications, here, are extremely important.

While first-generation KM schemes are decidedly technology-centric—perhaps even obsessively so—second-generation thinking gives priority to process-based organizational learning, with or without the use of technology. The advent of second-generation KM, then, can be seen as a convergence in thinking between the *organizational learning* and *knowledge management* communities. In effect, second-generation KM has emerged as an implementation strategy for organizational learning—a practitioner's model for how to help organizations increase their capacity to learn, innovate, and adapt in their own operating environments.

While the preceding paper of mine did a sufficient job (I hope) of describing the differences between first- and second-generation KM, it made no attempt to convert *SGKM theory* into

practice. That, then, is the purpose of this paper—a *what-to-do-about-it-on-Monday guide* to second-generation knowledge management.

Two Levels of Learning

The title of this paper is taken from the concept of *single-loop* versus *double-loop learning*, an idea hatched by *Harvard Business School* professor, Chris Argyris, several years ago. In an article of his entitled, *Teaching Smart People How to Learn* (*Harvard Business Review*, May/June 1991), Argyris explained the difference between the two in the following way:

“To give a simple analogy: a thermostat that automatically turns on the heat whenever the temperature in the room drops below 68 degrees is a good example of single-loop learning. A thermostat that could ask, ‘Why am I set at 68 degrees?’ and then explore whether or not some other temperature might more economically achieve the goal of heating the room would be engaging in double-loop learning.”

Argyris, therefore, makes the important distinction between two levels of learning. The first, single-loop learning, can be thought of as part of the process one goes through when attempting to function successfully in the real world. As people encounter discrete conditions, or events, during the course of normal experience, internally-maintained rules are invoked in response. Rules, in this context, means *knowledge*. Business processes, for example, are nothing more than codified

procedural knowledge—that is, business rules that inform workers of what to do in defined situations. As workers experience elements of these scenarios (i.e., as they “learn” of these events in the single-loop, sensory-perception sense of the term), they choose their responses in accordance with what their business rules tell them to do.

By contrast, double-loop learning not only references the rules, themselves, but constructively challenges the single-loop reflex to invoke rote responses. In the Argyris example above, the double-loop response to the below-68-degrees condition was to first reflect on the value of the prescribed single-loop response by considering whether or not an alternative to simply *turning the heat on* might fetch *better* results.

In the human mind, this kind of double-loop thinking would lead to active constructions of alternative scenarios in which the learner would play out likely outcomes in a lookahead fashion. Promising ideas could then be tested, in which case learners might actually override the prescribed response and temporarily replace it with a new one. Depending on how well the new rule fares in practice, the old one would either be reinstated for reuse again or would be replaced.

The point here is that living “agents” (e.g., people, animals, insects, communities, markets, economies, etc.) maintain active *rule sets* that are variously invoked from one moment to the next, as warranted, by conditions or events encountered in the outside world. Here again, the use of the terms *rules* and *rule sets* is meant to refer to an agent’s knowledge. In the lexicon of second-generation knowledge management, *rules* and *knowledge* are

synonymous terms. Moreover, the composition of rules is constantly changing—evolving in response to the creation of new rules and the loss of old ones. As organisms innovate by experimenting with *new* rules in practice, older, less effective rules (knowledge) give way to newer, more successful habits. Knowledge evolves accordingly.

The extent to which an organism engages in healthy rule-making and knowledge innovation will, to a large degree, determine its success in life. An agent whose rules are rarely tested will tend to perform more poorly in practice than one whose rules are constantly being challenged, upgraded and refreshed. The same thing is true for human organizations. A business whose approach to the marketplace, or whose operating processes are rarely revised, will tend towards ossification and atrophy. On the other hand, companies that engage in healthy levels of rule-making and *rule-revising* are inherently more adaptive and capable of adjusting to spontaneous changes in the marketplace. Indeed, organizational agility depends, to a very large extent, on just how well an organization's learning system is working.

That, then, is the principal aim of second-generation KM—to take whatever steps are required to enhance an organization's ability to engage in constructive levels of double-loop learning. In a sense, what we're talking about here is *double-loop KM*, an OL practitioner's framework for helping organizations, not just individuals, learn.

Double-Loop KM

Understanding Argyris' notion of single- versus double-loop learning is an

important early step in appreciating the fundamental differences between first- and second-generation KM. It's not just the relative difference in the extent to which technology is used in either case. Indeed, the distinctions go much deeper than that. Only first-generation KM schemes, for example, begin, and end, with the assumption that current knowledge in practice is valid! Their goal, then, is to optimize the delivery of currently-held organizational rules, or knowledge, to workers so that they can function successfully in their normal, operating environments.

This is why technology has played such a conspicuous role in the complexion of knowledge management schemes to date. After all, computers and telecommunications are unparalleled in their ability to deliver more information to more people, where it's needed and when it's needed, compared to any other medium we know of. Conventional knowledge management practice, then, boils down to little more than *getting the right information to the right people at the right time*. Think, "*single-loop learning*."

This conclusion is further supported by the analysis of other conventional KM practices, including contemporary attempts to build *communities of practice*. After all, human communities are the sociological equivalent of computer networks. *Bring people together in groups of defined interests*, first-generation practitioners tell us, *and better knowledge sharing will follow*. Once again, the target of this kind of intervention is enhanced *day-to-day operations*, or single-loop learning. The purpose of sharing knowledge is to distribute organizational rule sets as widely as possible so that

“best practices” can be employed pervasively on the job front.

First-generation KM schemes, therefore, are solely devoted to improving *knowledge operations* with the performance of day-to-day business processes in mind. This is why so many contemporary KM methodologies begin with a business-process orientation. They start by asking two very telling questions: *What knowledge do people need to do their work? And how can we help them get it?* Both questions reveal first-generation KM’s narrow preoccupation with business operations and the role of knowledge in supporting them—hence, the term *knowledge operations*. While this focus is in no way inappropriate or of little value to the organization, it completely side-steps the question of where organizational knowledge comes from to begin with—in other words, *how is organizational knowledge created?*

If a first-generation KM practitioner were asked to explain the role of knowledge management in business, an example of the following sort might be used:

A knowledge worker is sitting at her desk performing a task, then suddenly develops a need for information to complete her work. Where does she turn? Is the knowledge readily available? How long does it take to get it? Does she tap her relationships with other workers? Has technology been effectively placed at her disposal? Is her knowledge source current? Is it complete? Was the task successfully carried out? These are the kinds of questions we

wrestle with in knowledge management—it’s all about getting the right information to the right people at the right time so they can do their jobs more effectively.

This is vintage first-generation KM thinking in action. It’s all about single-loop learning. And it’s all about individual performance in the field. The target of all investments in first-generation KM is the individual worker and the extent to which he or she has access to, and can leverage, information needed to get the job done—where and when it occurs. Nowhere in this proposition is *organizational learning* mentioned, and *not once* is there any discussion of *knowledge creation* or rule-making. Only with the arrival of second-generation thinking do we see an application of knowledge management to these issues. What second-generation knowledge management offers, then, is an implementation strategy for *organizational double-loop learning*.

Knowledge Structures

One of the fundamentals of second-generation KM is the concept of *knowledge structures*. This is one of the more important ideas in the new field of second-generation KM, and so a brief review of what we mean by the term would be useful (see Figure 1).

Knowledge structures in human organizations are codified expressions of organizational knowledge. Business processes, for example, are codified expressions of procedural knowledge (*know-how*), while business strategies can be seen as codified expressions of declarative knowledge (*know-what*). Most of an organization’s shared

knowledge is expressed by being embedded in its behavior (i.e., in its practices). Propagating this knowledge from one generation to the next is accomplished either through explicit training or by the creation of *knowledge artifacts* such as data bases, documentation, and stories.

But none of this new. Human civilizations have been codifying knowledge for millennia in the form of myths, rituals, dance, and other cultural artifacts. This gives new meaning to phrases like, “etched in stone” or “cast in concrete.” Both refer to the codification

ways of making commonly-held rules more explicit—ways, that is, of making what a *few* of us know more available to the *rest* of us. The vast majority of what any organization knows (i.e., its knowledge, or rule sets) can be expressed in a formalized way, such that it can not only be shared with others, but can also be carefully *managed* as well (as in, *periodically updated* and

refreshed). Let’s take a quick look at the processes by which these evolutions in knowledge occur.

Figure 1: Organizational Knowledge Structures

	Declarative Knowledge	Procedural Knowledge
Business Strategies	X	
Products and Services	X	
Business Processes		X
Organizational Structures	X	
Policies and Procedures	X	X
Culture and Values	X	X
Information Systems*	X	X

*(Note: Includes automated, hard-copy and other *knowledge artifacts*)

of organizational knowledge. We might have thought knowledge management was new, but with the advent of second-generation KM, we now know it is not. Anthropologists have been studying the subject for years, now. It turns out that knowledge management is as old as the hills.

Thanks to anthropology, then, human organizations not only have a better appreciation *now* for the role that knowledge plays in guiding collective behavior, but we’ve also discovered

Knowledge Processes

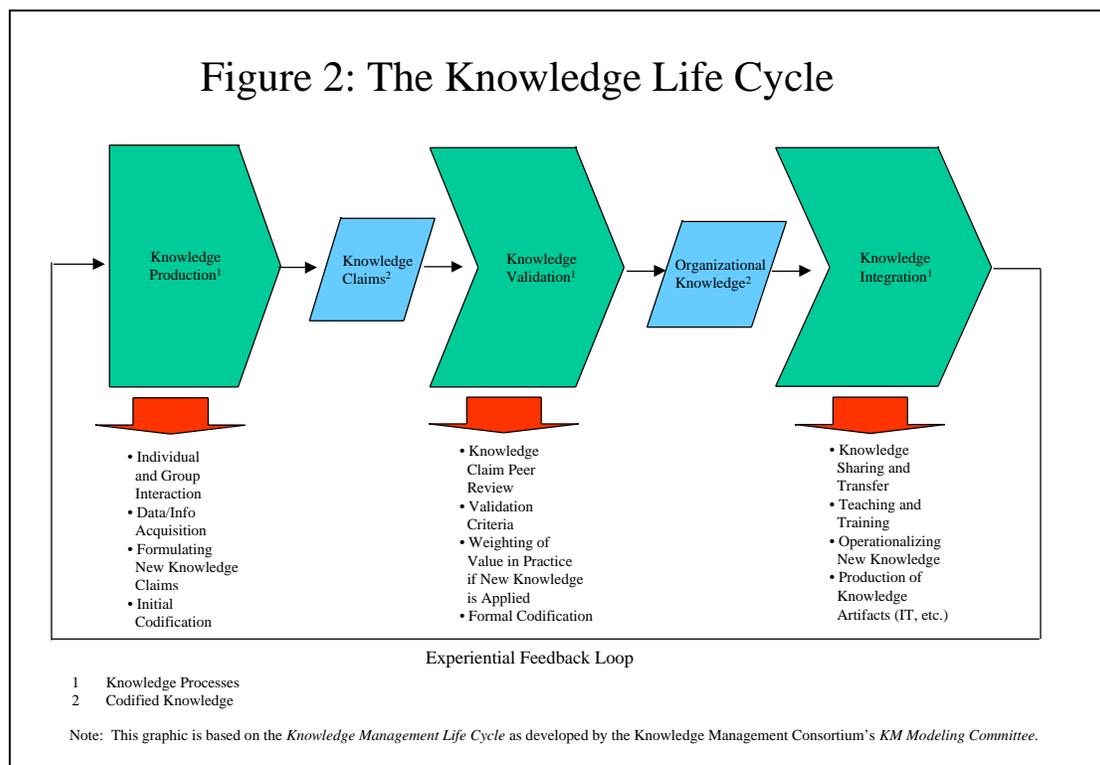
Since first-generation KM schemes are not at all concerned with knowledge creation at the organizational level, there are no *knowledge production processes* to be found in the prevailing style of KM practiced to date. Rather, first-generation initiatives have tended to concentrate only on knowledge sharing and transfer. This is usually attended by technology-centric solutions, a good portion of which invariably entail

codification and indexing. Indeed, most so-called knowledge management systems amount to little more than yesterday's technologies re-cast in today's more fashionable clothes. These include *document management, imaging, data warehousing, data mining and information retrieval.*

These technologies are all designed to support job performance in the field—knowledge delivery systems for knowledge workers in their day-to-day operations. Even corporate *Intranets,*

considerably more evolved. As discussed in my earlier paper on this subject, the knowledge management community (especially the *Knowledge Management Consortium*) has developed a three-phase model for second-generation KM. The three phases are, *Knowledge Production, Knowledge Validation, and Knowledge Integration* (see Figure 2).

So as not to spend too much time here on the life-cycle view of knowledge, suffice it to say that



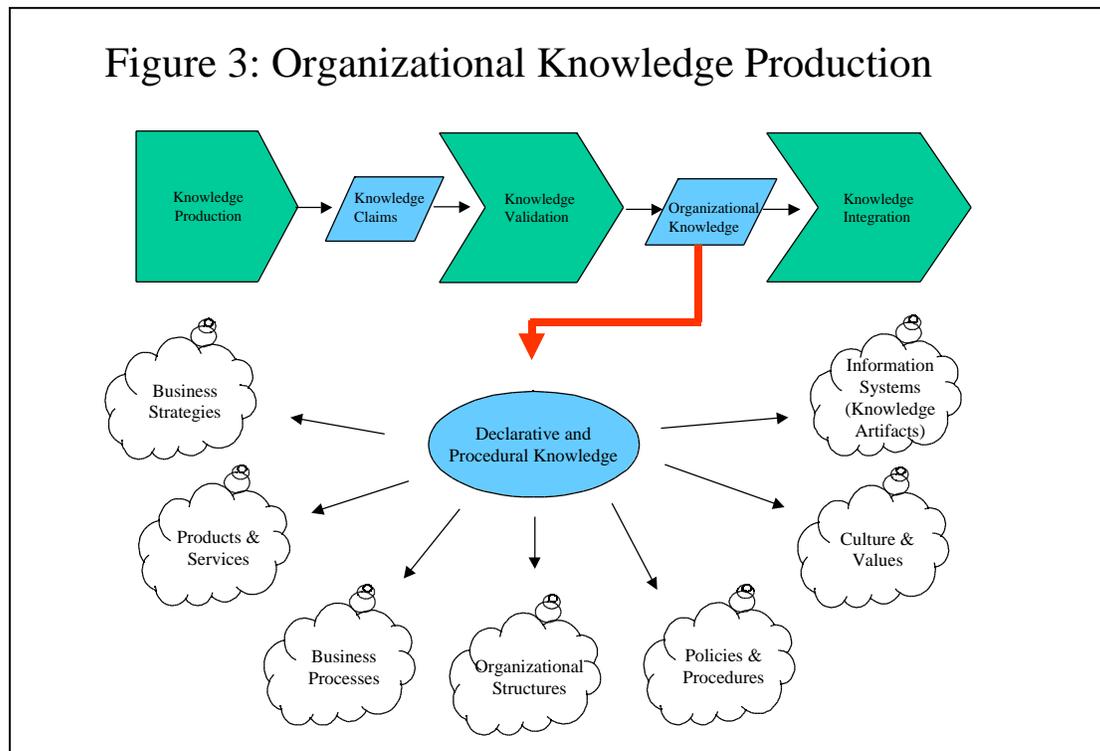
including the new, KM-spun variety of “knowledge portals,” are nothing but first-generation KM/single-loop learning tools—very useful in the performance of day-to-day tasks, but of little value when it comes to real *organizational double-loop learning.*

By contrast, since second-generation KM is concerned with organizational knowledge production, its perspective on knowledge processes is

organizations literally create new knowledge through a process of nonlinear interactions between individuals and groups, during which new “knowledge claims” are formed (*Knowledge Production*). Once they have been defined, these claims can be escalated for formal and informal validation (*Knowledge Validation*). The satisfaction of validation criteria often then leads to the formal adoption of *new*

knowledge at the organizational level. Operationalizing new knowledge, including first-generation style codification and indexing, is what the third phase, *Knowledge Integration*, is all about (i.e., *integrating* organizational knowledge into an organization's day-to-day operations).

structures) and the knowledge processes by which they come about (see Figure 3). Working to formalize and maintain these cognitive systems, then, should be the primary goal of the Chief Knowledge and/or Learning Officer in today's knowledge-based organizations. Let's explore this further.



Practice Implications

The emergence of *second-generation KM* and its vision of *double-loop organizational learning* have significant implications for practitioners of both KM and OL. Chief among them is the assertion that an organization's knowledge and learning environments can both be made explicit, and can be managed explicitly. The basic components of a well-managed knowledge/learning environment, therefore, consist of expressed organizational knowledge (knowledge

Tools and techniques for expressing human knowledge in syntactic terms have been out there for some time now, but only recently have attracted serious attention from the KM and OL communities. One such tool, the *Knowledge Harvester*, from LearnerFirst in Birmingham, AL, provides a systematic way of reducing both declarative and procedural knowledge into a form that can be converted to multiple expressions, or simply kept as a record of knowledge in practice. The value of these tools is easy to understand when faced with, say, the impending

departure of a valuable employee whose knowledge is either tacitly held or has never been captured for organizational diffusion. The one-time conversion and codification of an expert's tacit knowledge not only makes it easier to share that knowledge with others in the future, but also helps prevent the total loss of knowledge to the organization as a whole, should the expert choose to leave.

Organizational knowledge must be codified at the level of the organization, whether it be embedded in culture, procedures or information technology, if it is to have any chance of surviving the test of time. It is therefore in an organization's best interests to at least create *knowledge artifacts* that capture the sum total of an organization's knowledge in a formalized way. This, then, leads to a *first principle* for the practice of second-generation KM:

Organizational knowledge can be found in the various knowledge structures of an organization. The rules embedded in these structures (i.e., knowledge in practice) can, and should, be deciphered and managed, accordingly.

It therefore falls to the knowledge, or learning, function (K/L function) within an organization to systematically capture, in a formalized way, the essence of an organization's collectively-held knowledge. But this is for baselining purposes only. On the assumption that most companies have historically done no such thing, the initial codification of existing organizational knowledge will be a very time-consuming, slow, and priority-driven process—one that will likely take years to complete. Over

time, however, the total complexion of an organization's knowledge *can* be accumulated, thereby rendering the maintenance of existing rule sets, and the production of new ones, manageable propositions.

In addition to the codification and maintenance of organizational knowledge, the corporate K/L function must also take whatever steps are required to ensure that all three knowledge processes are active wherever they *should* be. More specifically, every *knowledge structure*, or knowledge domain, should be accompanied by a well-defined, well-running, explicitly visible knowledge life cycle that makes it possible to not only see the rules in use, but to challenge them as well. This points to the need for a standard approach to the implementation of knowledge life cycles in all organizations—a way that *we* produce knowledge, a way that *we* validate knowledge, and a way that *we* operationalize it.

How many organizations can claim that such a standard approach exists in their own domain? Very few, I think. And yet all knowledge is born of these processes, whether we formalize them or not. They're utterly natural and can be found in all living systems, human organizations included. I say it's better to embrace and exploit them to our advantage. This, then, leads to the second principle in the practice of *second-generation KM*:

Organizational knowledge is the product of natural knowledge learning processes in all human organizations. These processes should be formalized and managed, accordingly.

Here again, it arguably falls to the corporate K/L function to ensure that the *knowledge processes* that lie behind the formation of *business processes* and other *knowledge structures* are fully up and running. After all, knowledge in practice is only as good, or as current, as the last time its content was reviewed or refreshed.

This, then, leads to a third principle in the practice of *second-generation KM*:

Know what you know and why you know it!

Only by having the chance to scrutinize the basis of *why we know what we think we know* can real understanding and confidence unfold in our work. Organizations that reserve all of their rule-making processes for senior-level management, or *committees on high*, are only asking for trouble downstream. Subject to legal and competitive constraints, then, all organizational knowledge should be open to scrutiny by workers whose organizational lives are affected by such knowledge. After all, knowledge held by most organizations is ephemeral. One need only look at the frequency of corporate reorganizations, or business process redesigns, to clearly see that our best laid plans are usually short-lived.

The last principle of interest is one that speaks to the balance of emphasis needed in second-generation knowledge management programs. It is critically important that second-generation schemes avoid falling prey to the myopic obsession with activity at only one or two of the three stages in the KM life cycle. All three are essential, and all three are necessary, but insufficient, elements of the larger *whole life cycle*. Moreover, while second-generation KM

stresses the importance of *Knowledge Production and Validation* (typically *not* stressed in first-generation models), the kinds of activities found in the third phase, *Knowledge Integration*, including those that first-generation schemes *have* addressed, are no less important. Second-generation schemes, then, call for a more balanced view, one that the following mantra is intended to convey:

Innovate, Validate, and Integrate

This phrase captures the essence of all three phases in one simple expression. By actively maintaining knowledge processes that make all three functions possible, and by clearly expressing organizational knowledge—that is, knowledge held both *in fact* and *in practice*—the K/L function of an organization can go a long way in ensuring that true double-loop [*organizational*] learning occurs. This is the essence of *double-loop KM*, a *what-to-do-about-it-on-Monday* practitioner's guide to second-generation knowledge management.

Measuring ROI

No self-respecting practitioner's guide would be complete without taking on the obligatory question of how to measure, let alone express, returns on investments in KM. In short, there are two ways to do this. I will attempt to summarize both without going into too much detail.

The first way to measure ROI from investments in KM is a traditional one. Proposed expenditures in people,

processes and technology investments are accompanied by a forecast of what the business benefits will be based on the effect that KM will purportedly have on either lowering costs, improving productivity, or increasing revenue. This leads to a fairly straightforward business case using conventional performance metrics, which then lends itself to easy measurement at the end of the day. KM initiatives that successfully lead to their intended effects (i.e., lowered costs, higher productivity, or increased revenue) are declared victorious, and the numbers are tallied up accordingly.

The second way to measure ROI from investments in KM has nothing directly to do with business results, per se. This, in my humble opinion, is actually the more useful way to express ROI from investments in KM. Having said that, I admit that I hold the minority view in this ongoing debate.

Strictly speaking, investments in knowledge management can only *directly* affect the production, validation, or integration of knowledge in an organization (i.e., the three knowledge processes). Related business events including lowered costs, improved productivity, or increased revenue can only *indirectly* be attributed to the effects of investments in KM. A better measure of success, then, in judging

whether or not KM investments have had their intended effects would consist of measuring their impact, not on business events, but on *knowledge events*, instead.

By following the principles set forth above, a practitioner, for example, could literally measure the number of new knowledge claims (*innovations*) produced by an organization both *before* and *after* an investment in KM has occurred. Further, similar measurements could be made on the number of *validation events* experienced in the same time frames. Of even greater interest would be the number of innovations actually created, validated, and *put into practice* after investments in KM have been made, compared to the same measures beforehand. After all, one of the strongest propositions to emerge in the field of KM, to date, is the assertion that KM—in ONLY its second-generation form—can actually help an organization increase its rate of innovation.

That's the kind of claim that leads me to favor the second style of measuring ROI from investments in KM, since the first approach could lead one to completely overlook evolutions in knowledge events, like innovations, altogether. Perhaps a combination of the two makes the most sense.

* * * * *

About the Author

Mark W. McElroy is a Principal in IBM's Knowledge Management Consulting Practice. In addition, he is an active member of the New England Complex Systems Institute (NECSI), and serves as Chairman of the Knowledge Management Consortium's KM Modeling Standards Committee. Mr. McElroy is also a board member of the Sustainability Institute, a think-tank in Vermont that applies system dynamics tools to the study of economic and social issues in business and industry. He can be reached at mmcelroy@vermontel.com.