KMERA

The KM Conundrum:
History Lessons Point to a Brighter Future for Knowledge Management

by Maggie Law, Staff Consultant, KMERA Corporation, August 2001

Introduction: The Birth of a Notion

In a knowledge-based economy, an organization's success or failure is directly tied to its ability to manage its collective intellectual assets effectively. From the very first industry analysis reports (*circa* 1997) trumpeting this announcement[1], the panicked race to implement the smartest, fastest and cheapest means to accomplish this goal produced a rush of new technologies, (many actually pre-existing technologies with a new spin), and a new consulting focus, all erected from lofty knowledge sharing ideals. The hype and its accompanying pattern of disappointing results resulted in a legacy of failure to deliver on the promise of *knowledge management*.

In this paper I will not only draw conclusions about why knowledge management (KM) has so far consistently failed to live up to its potential, I will also propose new approaches based on simple paradigm shifts and valuable history lessons.

I. A Fundamental Problem: The Terminology

What does it mean to *manage knowledge*? To answer this question, we must first agree upon a common definition of "knowledge", then a common understanding of how it can (if it can) be "managed." A close examination of this fundamental question exposes a serious semantic dilemma.

To define knowledge in succinct, generalized terms is to betray its complex and abstract nature, rich with cultural interpretations and traditions[2]. To illustrate how we distinguish knowledge from other degrees of the cognitive process, the following hierarchy[3] is frequently offered as a guide:

- *Data* = "raw facts"
- *Information* = "facts with perspective and context"
- *Knowledge* = "actionable information"[4]
- *Wisdom* = "an understanding of which knowledge to use for what purpose"

Given the complexities inherent in a serious discussion of the meaning of knowledge, one quickly realizes many abstract philosophical, cultural, theoretical and scientific angles. Yet, remarkably, when discussing knowledge in the context of KM, we tend to equate knowledge artifacts with knowledge itself. The simple question "Where is your organization's knowledge stored?" usually elicits a long list of responses that might include paper documents, books, digital files, hard drives, the company directory, the departmental bulletin board, and the like before someone finally thinks to mention "in our heads."

This confusion underscores the contradictions of the KM promise. To posit, as is frequently done, that the prime objectives of a KM initiative are to *create*, *capture*, *classify* and *share* knowledge throughout an enterprise is to mistake knowledge -- that intangible mental product of our distinctly human cognitive processes -- with the tangible (physical or digital) assets of our work environment. Knowledge, being an abstract concept, cannot itself be "captured" without first being converted into some other form. Nor, extending this literal semantic interpretation, can it be "managed".

---

[1] Some early examples include "If We Only Knew What We Know: Identification and Transfer of Internal Best Practices" (APQC, 1997), "You Know When You Build It" (Application Development Trends, October 1997), and "The Third Skill" (The Forrester Report, December 1997).

[2] Ikujiro Nonaka and Hirotaka Takeuchi provide a thorough theoretical treatment of knowledge in chapter 1 of their groundbreaking KM book, The Knowledge-Creating Company (1995).

[3] Much has been written about this hierarchy, so I will avoid yet another iteration. Thomas H. Davenport and Laurence Prusak offer a solid discussion in their 1997 volume, Working Knowledge : How Organizations Manage What They Know.

[4] "Actionable information" is meant to convey a sense that one has absorbed enough information to be able to act on it.

So I return to the original question: What does it mean to *manage knowledge*? The term "knowledge management" seems to have been coined without deep consideration for a truly complex definition of knowledge. Knowledge management, as I will define it in this paper, is about:

1) transforming knowledge into *knowledge artifacts* or *information assets*,
2) classifying and cataloging these knowledge artifacts, then
3) providing effective methods of storage and retrieval so that they might be internalized by others and made actionable (i.e. converted back into knowledge).

Ideally, there is a final step, which addresses KM's powerful potential as a catalyst for innovation:

4) creating new knowledge based upon the internalized knowledge assets of others.

The temptation to confuse *knowledge* with *information* also routinely undermines the successes of knowledge management initiatives. Just as the term "knowledge" invites inconsistent implications, "information" is frequently defined in overly simplistic terms in the KM context.

It is, for example, generally accepted that a document, such as this white paper, is a *piece of information*. The digital file in which it is stored is a self-contained entity, as is the paper printout version. Both can be stored in a shared repository, such as a database (for digital document files) or a file cabinet (for hard-copy documents). Many would contend that these are the ingredients of a perfectly self-sufficient KM system.

To agree unconditionally with this assertion is, I believe, to undersell the essential human aspects of information that go beyond the words on the page or bits in the file (even the metadata assigned to the file). Too often, this ingredient is missing from our KM systems.

One problem arises when the tools fail to deliver enough information -- or the right kind of information -- to the broad sets of users they serve. Remember our definition for information: "facts with perspective and context." How well would an average user, sifting through a database of thousands (or more) documents, appreciate the perspective and context spelled out in this white paper without first reading and understanding it? Assuming there is metadata assigned to this document file, how well could it convey meaning to the user? Search engines and other information retrieval tools do a fair job of connecting users to the best available information sources, but success depends on the accuracy of search techniques applied. (Novice tool users will retrieve disappointing results, even with the best available technology.) Search tools are also remarkably poor at offering critical thinking with the information they deliver.

A second problem is what might be described as KM's own digital divide. If the goal of knowledge management is to enable, even encourage, non-hierarchical information sharing throughout an enterprise by making available a set of knowledge assets to every worker, is it realistic to expect that everyone be equally adept at retrieving, processing and contextualizing this information? How much success might an average CEO achieve when navigating though a file cabinet -- or a bank of file cabinets -- to locate the latest board meeting minutes (assuming he or she has time to do this)? Could every marketing director string together effective Boolean search queries to retrieve and compile comprehensive competitive analysis? Does every sales manager know when to look for the latest price lists, and where?

In theory, as in practice, it is unrealistic to expect that everyone interface with a KM system in exactly the same fashion. Unfortunately, this is the typical approach of KM strategists and system designers.

II. Our Technology-Obsessed Blind Spot

The allure of technology in the modern digital age has narrowed our focus to the detriment of successful KM results.[5] The high tech generation has already witnessed unprecedented technological advances, dramatic gadget miniaturization, and increasingly affordable pricing. The United States and its major economic trading partners have seen the almost viral proliferation of PDAs, cell phones, and laptop computers.

It is tempting and perhaps naive to believe that technology offers miracle cures for all that ails us. It corrects our spelling, it recommends books, it enables near-instantaneous communication from half a world away; it is, indeed, expert at certain logical, repetitive and statistics-intensive tasks. From a KM perspective, though, computers should leverage the unique and powerful intelligence of humans, not think for us.

[5] In their book The Social Life of Information, John Seely Brown and Paul Duguid describe how narrow-minded "tunnel vision" leads to "tunnel design". This may, perhaps, describe why no truly remarkable KM product (or set of products) has yet emerged -- and why KM strategists so often fail to look beyond the tools for compelling value propositions.

The reality is that technology is not intrinsically accessible or easy to use; it proves as much a barrier for some as it does a bridge for others. In our modern society, this is a lesson we have had to learn over and over again. Just as the CD-ROMs of a decade ago never replaced our daily doorstep newspapers, e-books of today have yet to strike the devastating blow to traditional paper version books many have predicted.[6]

In short, technology that is designed around the shaky promises and enigmatic deliverables of knowledge management is almost surely destined to fail.

Technology's history of failures and fallen expectations is well documented. Yet plenty of elegant, highly successful approaches inform our evaluation as well. Take, for instance, the telephone: often cited as a remarkable (perhaps even the most effective) information-sharing tool, it offers immediate message delivery and is widely adopted, relatively simple to understand, time-tested and cheap to deploy and use. Not coincidentally, three recent software innovations -- email, instant messaging, and online chat -- share many of the same characteristics. Most importantly, all boast an incredibly effective facility to enable users to exchange not simply information, but also abstract thoughts, ideas, and perspectives, all in near-instantaneous delivery time. The rich context implicit in the definition of "information" is not lost on these powerful tools -- and their successes reflect this.

A KM strategy that ignores the human experience of knowledge creation and information sharing is doomed. Our blind spot as creators and admirers of modern technology is that we forget to see ourselves in it.

III. The Search for Intelligent Life in Our Digital Tools

Millions of years of human evolution should count for something. From birth, we conduct our richest exchanges with other people. Perhaps this explains why, seemingly by instinct, we infuse human elements into our online communication channels where otherwise there are none. Our emails are laced with ASCII smiley faces[7] and shorthand acronyms for conversational body gestures.[8] Most instant messaging applications offer cartoon "emoticons" -- cartoon versions of our ASCII faces. AskJeeves.com and MySimon.com offer excellent examples of our tendencies toward anthropomorphizing the Web.[9]

The growing popularity enjoyed in recent years by natural language search engines (such as AskJeeves) reflects our instinctual desire to communicate with tools that "listen" to our questions and "hear" our requests as another person might. But to most users, the shortcomings of natural language search technology are immediately apparent. Not surprisingly, the more abstract one's question ("Why are Mondays so dreadful?"), the more irrelevant the results are likely to be.

Banks understand well that not all customers feel comfortable interacting with machines (ATMs) over people (tellers), and so offer both options. Phone companies recognize that paperbound directories are useless to callers who cannot read, cannot see, or else simply prefer to speak with a live operator. And, as research librarians can attest, some library patrons are satisfied to use the online catalog while others will ask for help at the reference desk. With only very few exceptions, the tools and applications that populate KM's history have failed to adopt as flexible an approach that accommodates all users: the technophobes and technophiles alike.

To many, a lack of human context feels cold and unforgiving. When this is the case, users' trust, loyalty and interest typically wane. As an example, online shoppers are notoriously disloyal customers. They exploit the convenience of the Internet to serve a global-scale variety of options. But because the sites they visit lack a personal interaction, deciding factors are price, shipping costs, turnaround and availability. In contrast, even an expensive, poorly-stocked neighborhood store can earn a customer's repeat business for any number of reasons: convenient location, friendly service, pleasant atmosphere, community involvement, to list a few.[10]

---

[6] "By now, the conventional book has been pronounced dead so often that we shouldn't be surprised to find that it seems in excellent health," concludes Robert Darnton in his article "The New Age of the Book." (Quoted from the New York Review of Books, March 1999, http://www.nybooks.com/nyrev/WWWarchdisplay.cgi?19990318005F#top)

[7] :-)

[8] One popular example: "ROTFL" for "rolling on the floor laughing."

[9] Also noteworthy is the current motto of Internet utility company Netword.com: "Making the Internet human." (http://www.netword.com/)

[10] Amazon.com presents a telling exception to this rule. While it is nearly (if not absolutely) impossible to connect with a live sales representative or customer service agent, the company has invested much thought and development effort into its online communities.

Inappropriate delivery can spell failure too. It would be tragic if an announcement of a building fire were delivered to those inside via email. (Even more tragic if all building dwellers happened to be away from their computers at the time.) As conduits of information, our inanimate tools too often fail to recognize that the manner in which our messages are presented to us, and not simply the messages themselves, require consideration of the human experience.

IV. Reconceptualizing Our Systems

Public libraries offer rich examples of effective systems designed for the sharing and dissemination of information. Highly refined cataloging and storage conventions developed over many centuries are well suited for the modern generation of library patrons. Yet ask the reference librarian what specialized or proprietary search tools he or she uses and you may be surprised to learn that they are no different from those readily available to everyone else. The librarian's advantage (and it is significant) is experience and training. Leveraging the librarian's expertise, patrons who lack even the most basic research skills can still locate the materials they seek. For them, the librarian is a valuable part of the information retrieval process and system.

Knowledge management's techno-centric history has so far largely failed to recognize the indispensable human role in successful information sharing. Even the term "system" (as in "KM system") invariably conjures the image of a network of hardware and software. Rarely, if ever, are systems designed to include the users themselves (expert or otherwise) as a natural and essential extension of the devices.

Similarly, conventional KM theory dictates that an organization's existing tools be leveraged (to whatever extent practical) during the design phase of a KM initiative. Traditionally, this has implied data legacy systems (CRM, SFA, ERP, *et al*.), intranet or extranet sites, public web sites, document repositories, and other automated containers of information and data. Rarely, if ever, do we think to consider the goldmine of qualifications, expertise, core competencies, related experience and career aspirations found in the organization's own workforce.

The unfortunate result: KM systems that alienate users. Valuable perspectives are lost or excluded from an organization's collective intelligence purely because some workers lack the skills required to operate the automated knowledgebase, cannot make sense of the taxonomies, or cannot comprehend the information retrieved by the tools.

In reality, not all users will become expert at manipulating the tools available to them -- no matter how expensive or how elegantly designed those tools may be, and regardless of how many training sessions are held over time. Yet without strong user interaction, the tools will fail to deliver on their full value potential. In other words, except in the rarest of cases (if ever), a KM system that expects all users to master its tools is both unrealistic and an unfortunate waste of time.

What's needed is a more balanced approach to KM system design: one that not only accommodates all users, but also capitalizes on the supporting roles these users might play within the system. We must recognize that machines and people offer distinct, non-interchangeable sets of core competencies, and that a KM initiative that focuses on one set of specialized skills -- to the near or outright exclusion of the other – is out of balance. It is impractical to expect computers to provide an adequate measure of human reason, to apply complex critical thinking, or to make subjective contextual associations. Likewise, it is unreasonable to ask a person to identify and retrieve all records stored in a knowledge base of millions that satisfy a specified set of search criteria, within seconds and without the help of an automated tool.

To define a KM system as a dynamic, integrated sum of *all* of its parts, inclusive of human users and automated tools alike, is to concede the very human nature of knowledge and to appreciate the complexity of our information sharing needs.

V. Pioneering a New Form of Computer-Human Interaction

KM's focus should not be on how well humans interact with machines, but rather on how well those uniquely human contributions are balanced and integrated with the unique contributions of computers. Together, a new interpretation of "computer-human interaction" emerges.

For example, many IT shops today offer 24/7 helpdesk support. Users access IT experts via phone, email and instant messaging applications to respond to technical questions or system/tool failures. This service is designed around IT's highest priority: the stable and secure operation of the networked infrastructure (hardware and software) that IT supports. Knowledge management has much to learn from this model.

The highest priority of knowledge management is the effective enterprise-wide sharing of information and knowledge assets. Simply providing all users a set of tools, then mandating that they use them (and do so effectively), rarely achieves this objective. Because people access and share information in many different ways, staffing an expert KM problem solver to assist users with research projects and information contribution/retrieval, and to provide critical or contextual guidance around the stored knowledge assets of the enterprise, would serve KM objectives incredibly well.

The goal is to build systems that strike a good strategic alignment with business objectives and processes.[11] Our history with knowledge management reveals a tendency to define our tools, wrongly, as an *end*, when in fact they are purely an enabling *means*. As a general rule, the more transparent the tool, the better.

Perhaps future KM tools will do a better job of understanding this. A telephone is effective because it enables instantaneous vocal communication; email, chat and IM enable instantaneous typed communication. The human "voice" is front and center. So many KM tools, for all their potential for features and multimedia capabilities, too often forget to share the spotlight with their users.

KM strategists would be well rewarded for focusing on better ways of weaving together an integrated experience of both human and computerized elements.

VI. Visualizing the System: The KM Triad

Having established that a truer realization of a KM system consists not simply of a network of devices and the knowledge assets stored in them, but also of the people who interact with both, how might we visually represent this system?

One approach is to display the three most essential KM elements as a triad[12] -- first to understand them as discrete elements, then to appreciate them for their interactions with each other (as represented by their overlaps).

In its simplest terms, the triad consists of three individual components:

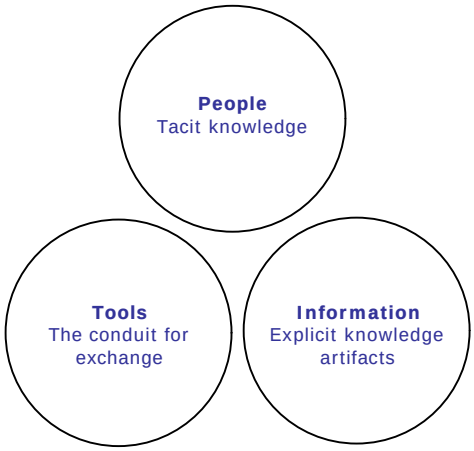

Fig. 1 - The Basic Components of a KM System

[11] Amrit Tiwana offers useful analytic advice and diagnostic tools to assist companies with KM strategic planning in his book The Knowledge Management Toolkit (2000, Prentice Hall). For example, it's important to determine if the needs of the organization suggest a "codification" approach (highly conducive to explicit knowledge documentation) or a "personalization" approach (requiring a means of drawing people together to share tacit knowledge).

[12] Stuart A. Robbins, CEO of Kmera Corporation and founder of the CIO Collective, deserves credit for the original concept of the KM triad. In this paper I offer detail and interpretation of my own, expounding upon Robbins' powerful and elegantly simple model. Additionally, Robbins' notion that IT systems are reflections of the teams that build them (articulated in his white paper entitled "The System Is a Mirror: Turbulence and Information Technology," published in 1995 by the Association for Computing Machinery) inspired my thinking around systems, themselves, as including both computer and human elements.

Some working definitions for the terms used in these representations:

People includes the entire population of end-users, potentially all of whom possess valuable knowledge to contribute to the larger organization. This knowledge is described as "tacit", referring to that which has yet to be documented explicitly and that which would be impossible to document.

Information covers all captured forms and instances of explicit enterprise knowledge artifacts (data or information), comprising both digital and physical objects.

Tools means simply the channeling devices that people employ to create, store, access and retrieve information, inclusive of anything from file cabinets to networked computers to the software running on them.

By definition, a *system* is comprised of interacting parts[13]. Examining how each component of the KM system interacts with each other helps convey the complexities of the knowledge sharing process.

The Embodiment of Information by People

The people-information overlap refers, quite simply, to those people who possess valuable knowledge in their minds that has not been explicitly captured for public consumption. They share this knowledge via non-technical means, such as verbal explanation or visual demonstration. Because no tool is involved, the methods for information exchange are entirely human.

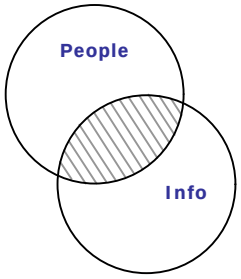

Fig. 2 - The People-Info Overlap

An apprentice gains valuable knowledge from a master craftsperson through observation and inquiry. Informal conversations around the company water cooler between employees from separate departments often yield vital information for some or all involved. Recognized experts in a particular field gain reputations as "knowledge fonts" or "gurus", and might frequently be visited by others seeking advice and insight.

When it is inconvenient, inappropriate, time-inefficient or simply impossible to convert tacit information into explicit knowledge artifacts, the experts themselves are living knowledge repositories.

The (Literal) Integration of Tools & Information

The fusion of tools and information occurs when the tool itself attaches new meaning or context to the information it is delivering; when the "packaging" or delivery method of a message can tell the user something beyond what the raw message conveys.

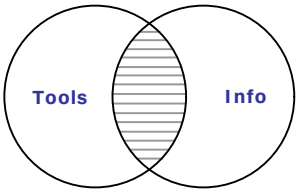

Fig. 3 - The Tools-Info Overlap

One example of the tools-information overlap is personalization, a technique whereby information-filtering tools automatically weed out content that fails to satisfy a set of criteria defined by the end-user. An information portal that scans thousands of news articles, then delivers a subset of "relevant" ones based on user preferences, has provided its own informative layer of context. While the strictly logic-based criteria are not immune to erroneous associations (it is, after all, a computer's brain doing the thinking) the overall result can be quite helpful to the end-user.

Other examples include metadata attribute assignment, relevancy rankings, push/pull agents, and just-in-place [14] communication.

[13] From the Merriam-Webster dictionary definition of "system": "a group of devices or artificial objects or an organization forming a network especially for distributing something or serving a common purpose."

[14] The term just-in-place communication (a play on similarly coined just-in-time communication) refers to the technique of strategic message placement most suited for the intended audience. In the physical world, a sign containing an important facilities announcement might be posted on a refrigerator or bathroom door -- locations most likely to be visited by all office dwellers. On a web site, an important announcement belongs on the most-visited page to ensure maximum visibility.

When tools employ logic to track user behaviors, model concepts, map patterns and associations, analyze communication channels and other trends on information and how it is being consumed, users can realize excellent contextual insight around the underlying information.

The (Figurative) Integration of People & Tools

The public library example described in section IV illustrates the concept of the people-tools overlap quite completely. Those library patrons, who (for whatever reason) do not access the collections via available directories or digital catalogs, rely on the librarian -- a trained tools expert -- to retrieve materials on their behalf. The librarian becomes an integral component of the patron's search and retrieval process.

From the patron's point of view, the librarian (a person) and the catalog search database (a tool) are, together, the complete instrument for accomplishing an information retrieval task. Furthermore, the librarian may possess relevant tacit knowledge, beyond what the search tool might retrieve, that will add further value to the results.

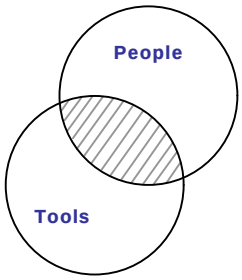

Fig. 4 - The People-Tools Overlap

It is generally the case that people who fall into the people-tools overlap are expert tool users, available to assist non-experts in some capacity. Telephone operators (a.k.a. "directory assistants") help callers locate and dial numbers. Scribes, photographers and A/V operators are specialists in the tools of transcribing, audiotaping, and photo documenting a meeting or event. Tech support agents guide novices through the pitfalls of software operation. Online discussion board moderators administrate BBS software to maintain topical relevance, answer questions from newcomers and enforce basic rules of etiquette, and so on.

In every case, people whose tasks and roles are tightly integrated with the tools they use are most effective when the service they perform for others aligns with their core competencies. (Someone who types 20 words per minute would be an ineffective scribe, indeed.) The roles described by the people-tools overlap support and enable knowledge sharing objectives, and so should be valued in any KM initiative.

VII. Assessing the Overlaps, Defining the Roles

The people-information overlap happens quite naturally. As humans, we have millions of years of practice teaching one another by showing and telling. The technology obsession of our modern digital age, fueled by venture capital and other sources of investment, cultivates the tools-information overlap. But as history reveals (and as I have so far argued), KM planning and strategy has tended to focus on technology that is disengaged from any meaningful human interaction. KM strategists typically forget, ignore or plainly undervalue the people-tools overlap component. But why?

For one thing, there is a strong tendency to expect our machines and software to solve more problems than they can. For another, it's unglamorous to concede that, given today's unprecedented abundance of available knowledge resources and so many sophisticated tools to help us make sense of this wealth of information, our best chances for success still require "old fashioned" human roles such as librarians, directory assistants, research assistants, and the like.

Yet in the case of knowledge management, a discipline that centers around a uniquely human mental process (knowledge), computers as yet offer no comparable alternative to the mind's profound capacity to reason, contextualize, judge, and abstract. Computers, ruled strictly by logic, are incapable of performing certain distinctly human tasks (i.e. those that require reason). Underrepresentation of either humans or computers can upset the balance of a KM system.

The knowledge management needs of each organization are unique and require individual evaluation. It's worth considering a handful of roles -- designed to address the fundamental truth that there's no substitute for human reason in a KM system, yet routinely overlooked (or poorly articulated) in most KM initiatives:

A system-wide librarian -- a problem solver, tools expert and facilitator for complicated research projects; oversees the information architecture (full system taxonomy) for all stored knowledge assets; integral in defining system policy; liaison to all user communities and SME librarians.

Subject matter expert (SME) librarians -- a subject-specific version of system-wide librarians; define and supply suitable metatags for content within a defined subject domain; oversee information architecture of

domain knowledge assets (subject-area taxonomy); liaisons between to system-wide and other SME librarians.

Scribes and recorders -- perform typing, transcription, audio/video digital transfer, file format transfer, language translation and other services; support the process of transferring knowledge into useful knowledge assets.

Content generators -- (potentially, though not always, the entire pool of end-users); qualified and devoted to the transfer of knowledge into knowledge assets; are optionally their own scribes or recorders.

VIII. The Core of the Triad

Studying the KM triad, one cannot help but wonder how to describe the center. Instinctively, we know that the place where all three essential components of knowledge management overlap must be special.

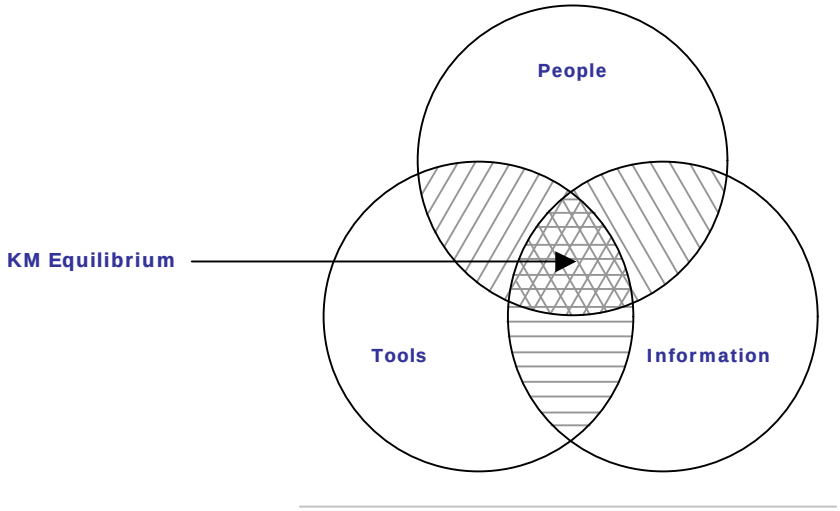

Fig. 5 – KM Equilibrium

At its foundation, the goal of any KM initiative is to remove obstacles that come between people and information. The center of the triad represents a balanced integration of both human and automated components -- *knowledge management equilibrium* -- providing an effective environment for accessing meaningful, context-rich information.

Clearly, this is no small feat. It implies an open, knowledge-sharing culture; it implies a set of widely adopted, non-intrusive tools; and it implies a body of rich, useful information. These are, in fact, the ingredients for success in any KM initiative. Once mapped to the KM triad, its success could be measured by the size of the center overlap -- the larger the better.

What drives KM equilibrium? The answer is communities.

IX. Communities

Consider our communities. There are many ways to define them: by a shared interest, a common practice, a neighborhood, a hobby, a belief, to name only a few examples. In a general sense, a community is a group of people who share something in common, feel a sense of belonging, take pride in their group, and are recognized as members of a whole. Communities are an ancient human tradition, and we define ourselves in no small part by our communal associations.

There are many similarities between virtual communities and their offline counterparts, but they are not entirely identical. The Internet demonstrates two powerful distinctions: 1) that a shared geographical location is no longer a requirement, and 2) that their member capacity is practically limitless.

MIT professor and founder of website framework technology company ArsDigita Philip Greenspun offers that an online learning community is “a group of people with varying degrees of expertise in which the experts attempt to help the novices improve their skills.”[15] By this definition, the end-users of a KM system could be similarly described.

Having acknowledged that knowledge management is as much about systems of people as it is about systems of automated devices, we can see that online communities -- like the computers that enable them -- offer models of dynamic networked interactions. Envision the members as a constellation of nodes, with messages and knowledge assets passing between them like so many data packets.

KM equilibrium describes the ideal environment in which online communities thrive. Essential characteristics of this ideal environment closely resemble those required for traditional (offline) communities. They include (but are certainly not limited to) trust, support, accountability, respect, free-flowing communication and self-governance; and are addressed in the form of firewalls, privacy policies, peer reviews, moderators, feedback channels and so on.

How we choose to define our online communities offers endless possibilities: by project, office location, team, department, profession, area of expertise, etc. And, because one individual may be affiliated with any number of distinct communities, our KM systems must flexibly support the variety.

Imagine, for example, an enterprise that consists of eight separate office locations (each represented by small numbered circles in Fig. 6[16]). Every location embodies its own group culture and identity. From a KM perspective, the information-sharing needs of each office are distinct as well -- facilities announcements in the Chicago office are irrelevant to the Helsinki office, etc. Yet all eight offices share a common set of enterprise-wide interests, represented by the shaded areas. Eight individual communities, defined by one criterion (location), collectively belong to a single community defined by another (enterprise-wide affiliation).

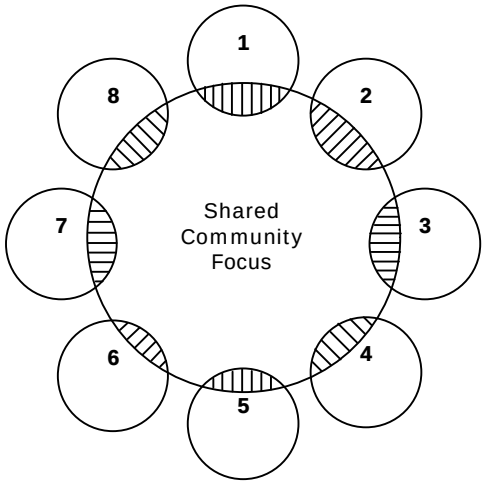

Fig. 6 – Communities: Individualities and Commonalities

Now imagine that each numbered circle in this same illustration (Fig. 6) represents a different functional department of the Chicago office (marketing, sales, finance, engineering, and so on). The shaded area they all share might now pertain to office-wide interests (building announcements, regional information, local birthdays, etc.).

In fact, consider all the intricate patterns our overlapping, interwoven communities might represent:

[15] Excerpted from chapter 3 of Philip and Alex’s Guide to Web Publishing, by Philip Greenspun (1999, Morgan Kaufmann Publishers). See http://www.arsdigita.com/books/panda/community.html for an online publication of this chapter.

[16] It should be noted that the basic design and concept behind Fig. 6 emerged from collaborative brainstorming among combined members of Kmera Corporation and the San Francisco location of design firm IDEO. Conveniently, this model provides a succinct framework for other centralized paradigms as well, such as a hosted computing environment or a shared file repository, where a multitude of separate parties or entities lend themselves, in some way, to a common middle ground.

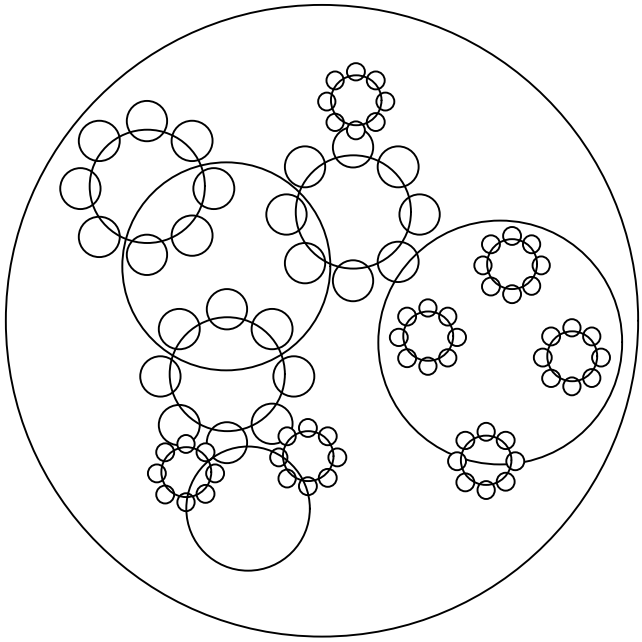

Fig. 7 – Communities of Communities

The job of a KM system is to link all of these communities -- each a "system" in its own right -- and enable effective knowledge sharing across them all. Boundaries that are drawn between or within these communities, which are meaningless from the KM perspective (e.g. a rigid org chart), should somehow be traversed. Meaningful boundaries (e.g. those that protect intellectual property or users' privacy) must remain intact.

By focusing on communities rather than technology, we embrace a flexible model that is boundlessly extensible. A familiar yet powerful pattern begins to emerge: the *extended enterprise,* consisting not only of internal members of an organization, but all external parties (partners, vendors, customers, advisors) that support the organization's successes.

And, because communities can be formed around *any* common thread, comprised of any combination of members, KM lends itself to any purpose imaginable. Communities of medical professionals, teachers, musicians, government leaders, entomologists, *ad infinitum*, might all convene effectively in a KM-type community setting. The tools become ancillary to the discussion.

X. The Challenge of Community Building

The technology behind online communities enables us to convene in much larger groups than could ever be achieved in person. It opens the door to unprecedented levels of innovation, collaboration and perspectives sharing. The biggest challenge of any online community, regardless of size, lies in sustaining member engagement through meaningful experiences and interactions.

Like atomic particles moving chaotically through the atmosphere, each of us conducts life blazing a unique path. By design or by circumstance, we make choices, pursue interests and follow instincts in a completely unpredictable fashion.

When we convene as a group, our individual interests and motives variously pull us closer into, or further from, the community focus. Figure 8, below, illustrates a perhaps more realistic snapshot of community member engagement than the uniformity presented in Figure 6. We see here that some members are more engaged than others, while one member (#2) has broken away entirely from the community. A snapshot of the same group taken a week, a month, or even a year later might reveal a very different picture. By then, member #2 might be very actively re-engaged.

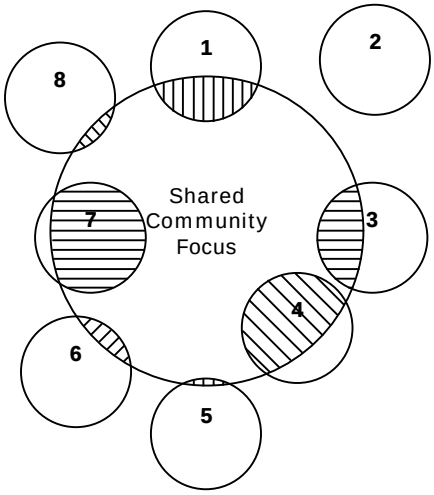

Fig. 8 – Typical Irregular Community Member Engagement

A community member's temporary disengagement is not necessarily a bad thing. He or she may, for example, introduce valuable outside perspectives back into the group consciousness. But prolonged or widespread member disengagement is a threat to a community's existence.

Itinerant curiosities and competing interests are typical of the chaotic pace of the modern information age. Online communities face the tougher challenge to earn prolonged member loyalty in the absence of face-to-face interactions.

For a community to sustain itself over an extended period of time, its members must recognize sufficient incentives for engagement. Meaningful rewards of membership provide the gravitational pull necessary to bind the group together around its focus. KM strategists must think carefully through the needs and motivations of the user groups they wish to reach, addressing the fundamental question that will inevitably enter every user's mind (over and over): "What's in it for me?" With so many competing demands on users' time, there must always be a compelling answer to this deceptively simple question.

XI. Conclusion

Since the earliest articulations of KM theories and principles, a variety of riddles have continued to confound experts and adopters alike. As a result, knowledge management has earned a stigmatized reputation in fairly short order, fueled by widespread frustration and a history of expensive mistakes.

Yet while failed KM initiatives have largely outnumbered successful ones, we can learn many valuable lessons from our experience and hindsight. Resounding perhaps most loudly has been our habitually disproportionate focus on technology in our KM systems, at the expense of essential human elements. We too frequently fail to accommodate the very human needs of our users, and to recognize the very human foundations of the knowledge we set out to manage.

Companies, academic institutions, government agencies and other types of working entities have begun to evolve beyond strictly isolationist thinking in favor of an extended view of the world. Virtual communities, powered by the World Wide Web, offer unprecedented flexibility and extensibility. We've only begun to realize our enormous potential to create, innovate, assist and empower ourselves and others.

But online communities don't just happen; users require motivation and clear reasons for participating, among other things. An open-minded approach that recognizes certain key elements of a successful knowledge sharing initiative, analyzed on a case-by-case basis and mapped to the KM triad, will provide valuable perspective toward a brighter future for knowledge management.

(9/13/01)