

Rethinking Knowledge Management: This Time It's Personal

By Stowe Boyd, Knowledge Capital Group

Unconventional Wisdom

Thomas Carlyle once wrote that "nothing is more terrible than activity without insight." Much of the knowledge management hustle-bustle over the past decade looks like a frenzy of activity with some crucial insight missing, meaning that the majority of knowledge management projects do not end in success, however measured.

Why? What is going wrong? Are the goals of knowledge management off base? Are we asking too much? Are people sabotaging these efforts willfully? Or is our problem much more fundamental? Perhaps "knowledge management" itself is to blame - perhaps "knowledge" cannot be "managed" in whatever the conventional understanding of that expression is, and so the basic, unquestioned premises about the practice of knowledge management are just unworkable. Is technology the problem? Are needed tools not available, or not yet invented? Or is technology just being misapplied, or incompletely integrated into the flow of everyday work? I want to start by taking one of these issues off the list: Effective technologies do exist, such as KnowledgeMail from Tacit Knowledge Systems, which I review in depth later on.

I strongly believe that the next wave in knowledge management is coming, and this time it's personal. The focal point of effective knowledge management initiatives today and in the future will start with the individual, and build upward from individuals to networks of people linked by collaborative learning relationships. Such knowledge networks can build up to the enterprise level, or even across enterprises. But the heart of these vital, active knowledge networks will be the individual, and the personal links between people based on their expertise, experience, skills, and purpose.

Last Time: Knowledge is Impersonal

The core premise of conventional knowledge management is that knowledge is a corporate asset, like money or real estate, and

should (and can) be managed like other tangible assets. The fact that knowledge sits inside a human head at some point is inconvenient, according to this viewpoint, and the best course of action for the enterprise would be to manage all critical knowledge in explicit forms - as documents, diagrams, drawings, formulae, and other static and stable forms.

This analogy has led to an explosion of 'knowledge capture' initiatives, where employees are asked to document and catalog what they know, and to submit their distilled wisdom into a knowledge repository. In practice this has turned out to be difficult, for a variety of reasons:

1. Not everything that everyone knows is worth capturing. The craze with converting human understanding into hard copy (because it is more easily managed in that form) is wrongheaded. It's like an obsessed diarist, attempting to catalog every event, every thought, in the chance that it will be useful in the future.
2. Even if you can limit capture to those people who have important knowledge (a tall order), what people know is not context independent: knowledge can be very situational. So documenting and cataloging general understanding does not necessarily allow a later reader to be able to make the decisions or take the appropriate action in a specific case that the expert might have done. Learning to think like another person requires more interaction than reading their observations.
3. People are often reluctant to work toward the development of an impersonalized system that makes what they do know less valuable. People's self-worth is closely tied to what they know, and even when a knowledge capture exercise is positioned as providing opportunities for all involved to increase their personal knowledge, each user may naturally draw back from full participation. The argument can be made that the company will benefit by having all contribute, collectively, to the knowledge repository. The whole is greater than the sum of the parts, after all. However, this does not mean each person is individually better off by yielding control of their personal sum of knowledge. For the equation to work both ways, the implementation of the knowledge system must give as much as it takes.

4. Serving up knowledge morsels from a central repository requires some search and retrieval mechanism - a portal or document management system, for example - that can handle all manner of queries and perspectives. This can prove to be very, very hard to design and build. As I stated earlier, knowledge about the world does not hang in space: it is very situational. Any extrinsic approach to categorizing and managing the information in a knowledge system - global taxonomic categories, keyword tagging of documents, and so on - will suffer from overgeneralizations. As anyone who has searched for something through a web search engine knows about this - 116,549 hits for "Qualcomm Results," for example. A radically different model is needed.

This Time: Knowledge is Personal

Technology approaches that are based on 'just-in-case' thinking - capturing all potentially useful knowledge today, just in case it will be needed in the future - go too far and not far enough. Some of what is captured will prove unhelpful later, and at least some of what will be needed later will not be captured.

Just as important, the application of 'just-in-case' knowledge requires context - consider a general dictate about corporate pricing strategy, for example, which must be particularized to market circumstances. It has been shown in many research studies that information is best retained when presented in a context where the information is relevant and timely. This means that a 'knowledge repository' will not generally be read end-to-end, just as people don't generally read the encyclopedia from start to finish.

In a perfect world, knowledge support technologies would operate just-in-time: the person looking for information - for a contact, for a pointer, or for a sanity check - would get that support exactly at the time needed, and in context. Guidance on pricing policy would be based on the current market conditions, not on abstract generalities. This has the added benefit of imposing a lighter burden on those providing the knowledge, since they are only providing information in response to actual - not presumed - needs.

And just-in-time is all about my time - the specific event driving me to look for guidance, advice, background material, contacts, whatever. It is not the phase of the moon, or the grinding of

impersonal market forces. It's about me, and my personal situation.

Unconventional knowledge technology should be based on personal possession of knowledge, transmitted interpersonally. Knowledge is found in the human mind, and only formalized, impersonalized versions of what people know can be found in documents, repositories, diagrams, and so on. While impersonal information is a useful surrogate for knowledge, it is only a surrogate. We can learn from a document, or a book, but we can learn more from the author.

While knowledge is intensely personal, it is also intensely emotional and social. People should be able to choose what to reveal to others - individuals or organizations - about what, and how much, they know. There is an implicit obligation to selflessly answer all sorts of questions once a person is seen as a guru in a field. As a result, various so-called 'yellow pages' techniques -- where the individual user does not control the specification of their own skill set -- do not accord with real world social interactions. For technology to conform to natural systems of knowledge creation and sharing, trust and respect must exist before rewarding collaboration can take place.

Expertise Automation Technology from Tacit Knowledge Systems

I recently had the opportunity to review the new Tacit product line, including the 2.0 release of KnowledgeMail, Tacit's flagship product, and came away convinced that the next generation of knowledge technologies are emerging, and at least one far-sighted company is opening the door to a new understanding of knowledge management. Tacit's model is person-to-person, and like well-known peer-to-peer technologies (a la Napster), offers a brand new take on how knowledge management goals should be played out in business.

KnowledgeMail is built on the principle that knowledge is personal, immediate, and social, and that it does not necessarily benefit the enterprise to attempt to laboriously build knowledge warehouses when people - who are the source of knowledge - already are pretty good at offering up what they know when asked to do so.

Email is one data source that KnowledgeMail can use to create a dynamically searchable directory of skills and expertise. Everything important - good or bad - is floating around in email first, and email is the bloodstream of many enterprises. KnowledgeMail automatically builds up your personal profile - based on what is being discussed in email - and allows a means for others to contact you when they are seeking advice in an area in which you have some interest, and are likely to be knowledgeable about. This is performed by an automated linguistic analysis of the content of your email and attached documents, yielding a long list of topics.

At the same time, people may want to keep some or all of their interests private, when they are immaterial or irrelevant to work. KnowledgeMail allows the individual to determine which interests form the public profile, which are kept private, and which should be deleted altogether. As a result, when others look at my profile, they would see that I am deeply interested in knowledge management, business strategy, and high-tech marketing. My personal obsession with Bugs Bunny memorabilia is my own business.

Note that KnowledgeMail goes the final step, and encrypts the individual's profile, so not even the system administrator - let alone management - can access the hidden data. This level of inviolability is intended to counter the obvious 'Big Brother is watching' concerns that arise - the system is 'reading' your email after all. But since the profile is encrypted and since no copies of the emails are retained by KnowledgeMail, users should be confident that their personal privacy is protected.

When a colleague is looking for guidance on business strategy, he would discover that I have that topic in my KnowledgeMail public profile, and that I spend a reasonable amount of time discussing that topic. However, the private side of the profile is also used, in the following way. Imagine someone in marketing is interested in Bugs Bunny memorabilia - a new line of business for the company, perhaps - and it turns out that no one is publicly knowledgeable. However, KnowledgeMail will inform them that a range of people - for example, somewhere between zero and ten people - have that topic in their private profile (note the vague hint on number - again intended to minimize triangulation on private information). And if the marketer wants to, he can ask KnowledgeMail to notify those people that he is

looking for some help on the topic. Until I respond, my private interest remains private. I can opt to add the now-material interest to my public profile, or to remain private on that matter. This feature, called Contact Request, offers the right balance between the individual's desire for anonymity and the company's business interests.

In this latest release the system can easily be populated with instantly searchable information by rapidly reading old emails or documents. Other performance improvements are core to this release, particularly with regard to scaling to tens of thousands of concurrent users.

Tacit now also offers the capability to integrate their expertise automation technology into other third-party applications, such as enterprise information portals, through the Tacit ESP (Expertise Search for Portals) suite or the Tacit Expertise Toolkit. The Tacit ESP suite includes newly developed interfaces as a result of close integration with several partners: Plumtree's Corporate Portal and Lotus' K-station. Both of these interfaces support KnowledgeMail style searching, and links the results of a topic search to the public profiles of proposed experts and interested parties.

Close

Knowledge is not a dead pile of facts, but on the contrary, the outcome of a dynamic interaction with the world at large, and most importantly, with the other people in it. Gregory Bateson once wrote "a business is best considered as a network of conversations" - and it is through the multiplicity of conversations going on in the business that knowledge is created, shared, applied, and built upon. Knowledge is not static, like a stack of twenties in a vault, but is an emergent property arising from the interplay of people joined together in a shared purpose -- it is both intensely personal and intensely interpersonal at the same time.

Promising new knowledge technologies - like Tacit's KnowledgeMail 2.0 - are moving past the publishing and repository models that dominate today's knowledge management initiatives, and offering an automated, non-intrusive, and intuitive approach to knowledge sharing. By ushering in a new, peer-to-peer style of interaction, where individuals directly interact through collaborative grassroots

networks, the impersonal, mass media publishing model is falling to the wayside. This time, it's personal.

Stowe Boyd is chief knowledge officer and vice president of Knowledge Capital Group, a strategic advisory firm (stoweboyd@knowledgecap.com; www.knowledgecap.com). Stowe is a contributing editor to *KM World*, and a frequent contributor to other well-known journals. He is the author of *The Way of Knowledge*, a reflection on the paradoxes associated with knowledge management.