

Knowledge Maps: An Intellectual Infrastructure for KM

Beginnings are always delicate times. Decisions made at the start of a project take on an exaggerated importance. In Knowledge Management, the strategy phase with which KM initiatives are normally begun, typically consists of such activities as doing a knowledge audit, performing a cultural readiness review, and a knowledge opportunity analysis. It is also during this strategy phase that it is recommended that you align your KM strategy with your business strategy and that part of this alignment process is deciding on whether to emphasize a codification or personalization implementation of KM.

A second activity that is often undertaken at the beginning of a KM project, sometimes as part of a knowledge audit, sometimes as part of a requirements phase, is to create a knowledge map. This knowledge map is usually seen as something that needs to be done, but once done slips into maintenance mode where it is one of many tools available for the KM effort.

However, I am going to argue that for companies in fields like Finance, stock brokerage, and indeed, most companies outside the consultant world, that the decision between codification and personalization is not a really essential strategic decision and that if the creation of a knowledge map is seen as one of the core elements in KM, the result will not only be better KM, but a better, more integrated relationship between codification and personalization.

Briefly, building, evolving, and utilizing a knowledge map should be central to your KM effort and if it is, a much more interesting and fruitful relationship between codification and personalization can safely be developed.

Codification and Personalization

There was an article in the Harvard Business Review, March-April 1999 that has become a standard element in a KM strategy phase, deciding whether to emphasize a codification or personalization approach to KM.

Briefly the difference between the two is whether you emphasize capturing, codifying, and reusing information – connecting people with information – or you emphasize connecting people with people and providing in depth expertise as your product.

A table of differences can help illuminate the two.

<i>Codification</i>	<i>Personalization</i>
Reuse existing information & solutions	Experts provide custom solutions
Connect people and documents	Connect people and people
IT centric	People centric
Low profit margins, high volume	High profit margins, low volume
Large teams, junior level people	Small teams of senior level people
Reward contributions	Reward sharing

One of the essential points in the article is that it is, at best dangerous and at worst catastrophic, to try to straddle the two strategies. They recommend that you pick one and devote 80% of your effort to that one approach, while keeping the other side contributing at about 20%. The danger of personalization companies trying to add self-service knowledge bases, for example, is the alienation of customers who expect personal attention and don't want to be told to look it up. Codification companies that try to add expertise based solutions or in-depth personal attention risk losing their profit margin.

In addition, the two approaches call for different incentive structures that can cause confusion if both are in place. If your normal reward structure or basis for performance reviews is how much material someone

contributes to the company database, how do you capture and evaluate their informal tacit knowledge sharing?

The article uses consulting companies like Ernst & Young and McKinsey as raw materials for what the different approaches are and for illuminating the dangers of trying to straddle the two strategies. For consulting companies, their arguments are quite compelling, but when they try to generalize to other industries, the arguments are a lot less powerful

For example, in a company like Schwab, two factors argue against a strategic either / or. First, the company consists of multiple enterprises with at least four major KM constituencies: IT, Corp Admin, Retail, and Institutional. Each of these four areas has different KM needs.

The second factor weighing against an either / or decision is that Schwab has multiple products and services with different mixtures of personal expertise needed. For example, the authors of the Harvard Business Review article ask three questions to help determine if you should pursue a codification or personalization strategy. The first is "Do you offer standardized or customized products?" Our answer, is yes. Second, "Do you have a mature or innovative product?" Again the answer is yes. The third is "Do your people rely on explicit or tacit knowledge to solve problems?" The answer is both.

Schwab has been known as a discount brokerage and a brokerage technology leader, and we are now expanding the help and advice area of the company. Even if the decision concerning codification or personalization was such an essential strategic decision, we would be left with more than one answer for different enterprises. In that case, the decision would become whether or not we have a single unified KM initiative or department or create two or more KM departments with different answers to the codification - personalization question.

In this instance, I would argue that the situation at Schwab is a better model for industries in the financial sector to follow than that of consultant companies like Ernst & Young and McKinsey. Our early research indicates that the real issue is not which path to emphasize, but how to create a system that integrates the two in a flexible and systematic way. The relationship between codification and personalization has two primary intersection points, creating a support context around when to make the transition from codified information to personal and tacit knowledge, and enriching your information storage solution with more and more knowledge contexts.

The first intersection involves the question of when do you need to (or when does it make sense to) make the transition from information to knowledge.

As part of a knowledge audit we did an initial survey in which we interviewed branch employees who work in our the branch offices where customers come in for a face to face interaction with a representative. Many of the answers that these representatives provide are straight forward and can be quickly and easily answered out of either the knowledge in the representative's head or by quickly looking up the answer in an online reference application.

The majority of these transactions are information exchanges and rely on a codified information management system. However, there can and does come a time when the representative cannot answer the needs of the customer and then the essential question becomes if and when to escalate the situation to someone with more expert knowledge.

What KM brings to the situation is to help manage and support people deciding to escalate, training them on how and when to escalate, and ultimately, to enable this escalation point to be shifted so that more and more queries can be handled without escalation, in other words, to incorporate more and more knowledge into the information system layer and thus delay the need for escalation.

A knowledge map allows you to codify more knowledge elements and at the same time, supports the decision by human workers of when to escalate from information to knowledge, from a simple document to a human expert. An integrated KM system based on a knowledge map also provides a framework that

enables the transition to go more smoothly as the knowledge base, the initial contact rep, and the expert rep are all using the same vocabulary and categorization schema.

What I am going to show in the next section is that it is not only possible, but strategically sound, to build an integrated codification / personalization approach to KM if you build on the right foundation of knowledge architecture and a powerful and dynamic knowledge map and include a flexible KM team of internal consultants.

Knowledge Architecture: enriching codification and codifying personalization

Knowledge is a very difficult thing to define, but without some sense of the difference between information and knowledge, you run the risk of confusing the two and developing a confused approach to KM that falls prey to the hazards noted in the Harvard Business Review article. The danger isn't trying to approach KM with both a personalization and codification strategy, its creating a system that does not allow for the smooth integration of information and knowledge, codification and personalization.

Regardless of the definitional difficulties, we have an intuitive sense that knowledge is broader, deeper, and richer than data or information. Knowledge is information plus something - meaning, action, organization, patterns, or whatever. Rather than add to the list, I'm just going to say that the something extra is what we'll call context. Knowledge is information plus contexts of a variety of types.

There are two major kinds of contexts, intellectual and personal. The latter includes not just individuals, but also collections of individuals into social communities. It is the defining characteristic of Knowledge Management to model, organize, and support these additional contexts. Adding in these contexts is what differentiates knowledge management from information management.

In the area of intellectual contexts, when you go from information to knowledge, you are going from a relatively static and one dimensional area to a land that is dynamic and multi-dimensional. It is somewhat analogous to a Flatland progression from point (data) to lines (information) to an escape from flatland into the world of three dimensions (knowledge).

Because of these extra dimensions, the normal KM approach is to abstract these contexts from information, store the information and then focus on the process of a human brain converting the information back into knowledge.

KM can improve the conversion process by providing knowledge facilitators to work with employees both during the conversion of knowledge into information for storage and on the subsequent conversion back. The dangers of loss of context, distortion of meaning, and other losses attendant on any conversion process are particularly acute in this case since it is those contexts that are being converted twice that actually create knowledge out of information.

By adding facilitators at both steps, some of the loss and distortion can be alleviated, and more value can be added through another major component of any good KM system, the alliance between KM and learning. Training can empower the process of converting information in a database into knowledge by not only teaching good technique, but also by norming the results, that is, ensuring that the human side of the conversion is knowledgeable in what the company wide or community of practice wide consensus is.

However, no matter how good the conversion process gets, the less conversion that you need to do the better. This is where a powerful knowledge architecture effort can provide great benefit. By adding the ability to store more elements of the contexts that make information into knowledge, we make the systems smarter, minimize the loss or distortion of knowledge during conversion, and can integrate more fully with the training efforts.

Let's take an example:

We have a document that is a procedure statement. It contains pieces of information like, “When you finish X, the next step is Y.” There are a number of additional contexts through which this piece of information becomes knowledge.

History. We used to do Y before X, but it led to all sorts of problems.

Applicability. This procedure only applies to customers with an account over \$100,000.

Personal. You have never read this procedure document before.

Value. Following this procedure is important because if you don't, the company will be liable for a \$10 million fine.

Relatedness. This procedure is similar to procedure P in the Branch enterprise but differs in the following ways: X, Y, Z.

Some of these contexts can be captured and stored in a knowledge map of the enterprise. In some case, the critical knowledge can't be stored, but pointers to other documents, processes, and people can be stored. For example, if you have a smart knowledge base that knows whether or not an individual has ever seen a document (and how long they have read it), and knows that other documents in the knowledge base provide background or training for a novice reader of procedure X, then some of the contexts can be stored along with the information.

What is essential for an integrated approach to KM is that the intellectual infrastructure be consistent across topic or subject areas and across the personal and social tacit knowledge components. In other words, a knowledge map must deal with the integration of all types of contexts, intellectual, personal, and social. The nature of these latter two contexts can be seen in three example areas, personalizations, collaboration, and knowledge retrieval.

Personalization has been touted for a number of years as a high value proposition. Unfortunately, the payoff has often not materialized. One reason for this is a poor understanding of the need for a rich knowledge architecture to support personalization and place it in the overall context of communities.

This architecture must account for a variety of roles and functions, membership in a variety of communities, and be able to be mapped to a categorization of tasks and processes. It must include a temporal and historical dimension as well. For example, having a person categorized as a client service representative is useful, but having a system that knows the difference between an experienced client representative and a novice is even more useful.

Communities can be created around a variety of activities, interests, and channels of communication. Communities can have an established life span or can be created on the fly for a single meeting or the life of a project. KM support for collaboration must model and support all these communities and their interactions and must also integrate with the intellectual categorization schema.

A knowledge retrieval system for communities should include both the retrieval of information and surrounding knowledge contexts on the one hand and the location of humans that contain the knowledge the individual is searching for. The latter capability, an expertise locator, is one of the more exciting areas of knowledge management. However, without a proper foundation (a knowledge map), expertise locators could end up like many personalization efforts - where's the benefit?

Too often projects like developing an expertise locator, or an enhanced knowledge retrieval system, or a collaboration platform are approached either as just another technology project or even if they are recognized as belonging to knowledge management, they are seen as largely separate projects with one set of categories for experts, one for information or documents, and one for collaborative communities.

None of the three signature knowledge management projects will achieve their full potential until they are all integrated and rest squarely on the essential component of knowledge management, the knowledge map.

The Strategic Role of the Knowledge Map

The battle cry, "It's the Culture, Stupid!" represented a major advance in KM. Instead of pouring 100,000's and millions of dollars into technology infrastructure and wondering why there was little payoff, companies began realizing that culture was an essential ingredient. Some analysts claim that success in KM is 80% cultural.

While I agree the culture battle cry represented an advance, it is not enough. Currently companies are being told to find the right balance between technology and culture, between the technical infrastructure and organizational infrastructure. But this overlooks the most important infrastructure of all, the intellectual infrastructure. In other words, It's the Knowledge, Stupid!

During our strategic planning phase, we developed a three infrastructure model. The technical infrastructure includes the actual network with everything from central servers to the desktop. It also includes enterprise infrastructure elements such as a central database of employees. Most of this component was already in place and is needed whether or not a company is going to invest in KM. The one ingredient that we need to add is a single sign to support personalization and the metrics necessary to develop smart and adapting systems.

The technical infrastructure also includes such enterprise software as a powerful and customizable search technology, some way of developing and supporting a people based search - expertise locators, and a collaboration platform.

The organizational infrastructure includes the cultural elements of KM as well as such practical questions as how to staff a KM department, where to locate the department in the hierarchy, and which roles to bring into the KM team and how to plan for their integration into the other enterprises. These are the questions that have become the focus of much of the strategic thinking around KM and we are learning how to adapt that thinking to our own situation.

Where we are having to develop new ideas and strategic plans is in the area of the third leg of the infrastructure tripod, the intellectual infrastructure which largely consists of the creation and maintenance of a knowledge map.

What is a Knowledge Map?

A knowledge map is the intellectual infrastructure for KM initiatives. The basis for it consists of multiple taxonomies for content repositories, dynamic categorization of people, their expertise, and the communities they belong to, and finally a set of taxonomies for the variety of tasks that are performed within and by the company communities.

The taxonomies of content, people, and tasks then have to be mapped across the three components in order to provide a foundation for the integration of such KM enterprise projects as knowledge retrieval, for both document based knowledge and the tacit knowledge located within the minds of the companies experts. It is also the foundation for collaboration, both for capturing the knowledge that is generated in those collaborative communities, and for providing the framework within which knowledge facilitators or knowledge managers will operate as they provide services for those collaborative communities.

Perhaps the best way to describe a knowledge map is to describe how we are building ours.

The place to start is with a meta data standard. We started with the Dublin core and then added some additional tags like ContentType and audience. For both of these tags we are developing controlled vocabularies. The audience tag will be the basis for our personal and community taxonomies.

We are looking to expand our Meta Data standards in two ways. First by expanding the controlled vocabularies for audience and subject or keywords. Once you have a rich enough set of values in your

controlled vocabularies, the more difficult, but much more rewarding task of creating semantic networks to capture the rich mosaic of relationships among your values can begin.

The second way we are looking to expand is by looking at the new meta data approach offered by RDF, or Resource Description Framework. This technique can be described as an XML layer added to standard meta data. It is still an open question how much effort to put into implementing this standard. The alternative is to rely on the tacit knowledge of a KM team to enrich the current meta data standard.

With the meta data standard in place, the first application was to build a browsable taxonomy of content on the corporate Intranet. We call ours, the Yellow Pages, and it not only will be a project with an immediate payback in saving users and web developers time and effort, it is also the first step toward building a knowledge map.

To grow our yellow Pages into a full scale knowledge map, we will need a variety of approaches. First, we need deeper and more dynamic categorization of all our content. And then, we have to develop the audience meta tags into more complete personalization and community descriptions.

To do this, we are planning on a three prong approach. The first is the hard method track and it will rely on expanded search technology. There has been a recent flurry of offering in the area of automatic categorization tools offered by search vendors and smaller specialty vendors. We are currently doing an evaluation of this product space with the idea that what is really needed is not automatic categorization (still too low quality) nor human categorization (too costly), but rather what we are calling Cyborg-Categorization.

Cyborg-Categorization is the fusion of automatic and human categorization. It means that the automatic tools are going to be used by human categorizers to speed up their efforts and allow them to systematically and creatively explore the rich content and varied communities with our company.

The second method is to send knowledge engineers or analysts out into the field to interview knowledge workers and incorporate their findings into the knowledge map. These KE's will also begin the process of exploring how best to use the knowledge map within the various target groups.

The third method is to incorporate the results of the first two tracks into a enterprise content management and document management system. A content management platform allows a dynamic and distributed publishing procedure and provides for the transition from a web site or department publishing model to a taxonomic publishing model. This will mean that we can work with content providers to continue to build the knowledge map that supports that content, by creating work flows that integrate meta data capture and routing of content to both Subject matter Experts and Knowledge analysts.

The role of content management and a team of Knowledge analysts and/or managers points out the last and most important part of a knowledge map and why it needs to be viewed as the third infrastructure leg rather than, as many authors do, a project that is done at the beginning of KM and then is largely done or slips into an easy maintenance mode.

It's a shark! Like a shark (or a relationship in a Woody Allen movie), a knowledge map has to keep moving or it dies. It is an evergreen project that not only is never finished, but its continuing development and use is perhaps as significant as the cultural readiness of a corporation as a deciding factor for the success of KM. It is not just that there is new information, people, products constantly being developed and added to the system, it forms the framework for how knowledge is incorporated into all employees work.

Workers and content providers are ultimately the experts when it comes to the subject matter of their areas, but they are not skilled knowledge analysts and do not normally think of their content in terms of the category of knowledge it represents, the relationships between knowledge chunks or objects, or other KM components. Therefore, in order to keep the system moving and alive, the knowledge map and its utilization in content management, collaboration, and knowledge retrieval must be an essential part of your KM system.

Finally, the constant development and utilization of your knowledge map by both employees and knowledge facilitators is an important part of obtaining the feedback loop that any good knowledge management system needs. For example knowledge managers need to know the map and use it and since KM is a self-referential initiative, knowledge managers need to use it to capture the knowledge in their projects to refine the knowledge map based on actual experience.

In summary, for industries in the Financial sector as well as most large companies outside the consultant sector, I would recommend, based on our experience and analysis, that they not choose between codification or personalization as 80% of their KM direction. What they should do instead is devote resources to creating the third infrastructure leg of KM, the intellectual infrastructure which will become the means for an integrated balance of codification and personalization components. This intellectual infrastructure rests primarily on creating a knowledge map, a living, breathing, evolving knowledge map that consists of dynamically categorized repositories of content, people, and communities and the knowledge managers and facilitators that use the knowledge map and contribute to its continual evolution.