

From Information Architecture to Knowledge Architecture

Borrowing a page from Animal Farm,
"IA good, KA better, IA good, KA better"

Or, in less sheepish words, information architecture has added a dimension to corporate intranets that was lacking once those intranets grew beyond the vanity press home page sites: the ability to find information at a reasonable cost in time and labor. As needed as information architecture is, however, knowledge architecture adds a dimension that will dwarf the cost savings and productivity gains achievable by Information Architecture.

Unlike the Animal Farm 4 legs good, 2 legs better, KA doesn't replace IA, nor is it a political coup. Rather, KA rests squarely on a foundation of IA and the better the IA, the better the KA that can be built upon it.

Before we begin to discuss the evolution from IA to KA, we need to spend a little time on what the relationship of information and knowledge is. I recently gleaned a number of definitions of knowledge and came up with the following incomplete list:

Knowledge is information + meaning
Knowledge is information + action
Knowledge is organized information
Knowledge is applied information
Knowledge is understanding patterns
Knowledge is information + a person reading, understanding, interpreting, applying
Knowledge is information + more complete patterns
Knowledge is something that we digest rather than just hold
Knowledge is justified true belief.

I'm not going to try to add to the list, but if you examine the list, you quickly realize that despite the differences, they all say pretty much the same thing. There is an intuitive sense that knowledge is broader, deeper, and richer than data or information. In other words, Knowledge is information + something. Again, I'm not going to try to define too precisely what that something is except to call the something extra, contexts. So knowledge become information + contexts. Which means that knowledge architecture is information architecture plus contexts.

What kind of contexts? I would like to discuss three major types of contexts: intellectual, personal, and social or community contexts. Each of the three types has different implications for the shift from IA to KA.

Intellectual contexts.

This is where you find traditional IA at its most developed. Traditional IA deals with organization of information in taxonomies with alphabetical or hierarchical schemas, navigation structure like a Yahoo-style browse tree, and search.

Information, however, is relatively static and one dimensional compared with knowledge and so we need more dynamic taxonomies and multiple taxonomies to support knowledge activities. For example, take the piece of information, that "Our policy is to do X." This becomes knowledge when you add context. Some sample contexts:

History. We had policy Y, but we had to change it because of problems with B.

Applicability. This policy only applies to the United States to people with job ranks above VP.

Personal. This policy has never been read by person A.

Practical. This policy should be referred to when dealing with customers in situation Z.

Valuation. This policy is important because it protects us from certain liabilities.

It is relatively easy to abstract information from its context and store it. However, knowledge cannot be so easily separated from context, indeed, it is its context that gives information the added meaning to become knowledge. Knowledge is not something that we try to store. Instead we store information and ask people to convert it into knowledge. What this means for KA is that it is not only concerned with the structure of the storage system, but with embedding KM into the procedures of people interacting with their information systems and converting information into knowledge, thus making it easier, faster, and more productive..

While I'm not sure I agree with notion that in 2 years E-learning will be a subset of KM, or KM will be a subset of e-learning, nevertheless the two are as tightly intertwined as two fields can get. In fact, within the realm of intellectual contexts and storage of knowledge, it is true that knowledge architecture is an architecture for learning, an architecture that enables people to convert stored information into knowledge.

Another example of the evolution from IA to KA based on a foundation of information architecture is a meta data initiative. The first step was to create a meta data standard based on the Dublin Core. A good start, but then we had to design an implementation path for getting meta data into our documents. We have done that. But now comes the really hard part, developing controlled vocabularies to fill those meta data fields with. And as we develop those controlled vocabularies, we will be describing the intellectual or knowledge space of the corporation.

Creating these controlled vocabularies (and there are more than one), will require a significant knowledge engineering effort that will involve not just compiling lists of terms, but more knowledge engineering efforts including interviewing content owners and their customers. It will involve not just an alphabetical or subject matter index, but will need a semantic network of terms.

I want to close this section with an idea that is still in the theory stage, but which potentially could be the most valuable, the most interesting, and perhaps, even the most fun. It is Knowledge Architecture for stories.

Stories are one of the most powerful and effective means of transmitting knowledge that we have, but there has been little done to incorporate stories into an effective knowledge management program. Of course, one very good reason for this might be that it is a task on the difficulty scale approximately on a par with walking on water or finding the lost city of Atlantis.

However, there are some approaches that could begin the process of putting in place a sophisticated KA for stories. Let's start with a quick look at stories in their natural environment, a community of practice. Story telling is essential to the creation and maintenance of a community of practice. It builds and sustains the community, supports learning by individuals and the community as a whole, and establishes the place of the story teller within the community.

Story telling as a knowledge transmission mode teaches in a multi-dimensional fashion which enables such added benefits as participants learning the vocabulary and values of the community virtually effortlessly.

The first thing a KA for story telling must do is to establish a framework to support the activity of story telling. Technologically, this means incorporating email, bulletin boards, voice, and digital media into the way people work, and to build smart search technology on the back end.

It means putting in place reward systems for adding to the story hoard. It also means teaching people how to tell stories as part of their interaction with their community of practice. A related technology might be to capture more of the tacit knowledge involved in daily activities and implicit in the stories themselves.

But what about the actual intellectual infrastructure needed for stories as a knowledge transfer mechanism? There is no one answer, but there is a single approach that itself uses multiple means of understanding human interactions, its the field of history., particularly the history of ideas and critical historiography.

What it needed is the intersection of history, cognitive science, artificial intelligence, and knowledge management. Of all those fields, history is the real stranger, and potentially it could create the most value of them all.

An information architect looking for a place to start might consider the following quote by a German philosopher and historian of ideas, Ernst Cassirer,

“But even in historical thinking the particular fact is significant only by virtue of the relationships into which it enters. Although it cannot be regarded as an instance of a general law, yet in order to be historically conceived, to appear sub specie the mode of history, it must take its place as a member of a course of events or belong to some teleological nexus.”
Language and Myth, p. 27

In this case, the course of events are the story elements and the teleological nexus could be considered the moral or point or lesson of the story.

Finally, any knowledge architecture that grapples with establishing a rich architecture for intellectual contexts, needs to also take into account the need to model the differing knowledge state of the individuals that will access this knowledge. Which bring us to our next set of contexts, personal contexts.

Personal contexts

IA and KA for various personal contexts deal with the infrastructure for personalization, one of the oft-cited advantages of portal technology, but which often fails to deliver its promises. One reason for this failure is in the lack of a knowledge architecture or information architecture to support the various ways of defining a person and relating that definition to a set of tasks that a person can do.

For example, one portal element that is often touted is the ability to create a personal newspaper. However, a personal newspaper is useless in so far as it is personal. People don't work as isolated individuals, they work in communities. In other words, selecting only those stories that interest me might save some time, but having an additional dimension of stories that others in my field deem important enriches me.

Another example of the interplay of personal and community can be seen in the new search technologies that include an adaptive component. Having a search engine that remembers your search behavior and which adapts its ranking of results based on that history of behavior is a very powerful tool that can save time. Having a search engine that adapts based on the history of community behavior can enrich as well as save time.

However, creating an infrastructure that supports multiple membership in a varied range of communities and that supports an individual in their work does require a good personalization foundation. Two critical elements of this foundation are categorization schema for people and for tasks. Like the schema created for intellectual contexts, these personal and task-based categorizations need to be dynamic and multiple.

Individual can be categorized by role and by function or job title. One individual can be categorized in many ways, but even more important from the perspective of KA, you must be aware of the historical and temporal dimension and build that into your schema. For example, having a group category of client service representative is good, but distinguishing between an inexperienced and an experienced rep and providing a systematic way to track individuals as they grow from one category to another better.

Another dimension of a KA for personalization is setting up rich profiles to capture the tacit knowledge that individuals carry around in their heads. Connecting people with people is often much more important and rewarding. As has often been stated, especially in knowledge management literature, it's not what you know but who you know that knows what. An overstatement, but one with a useful lesson.

An exciting new area in the creation of personal contexts is connecting people with agents and agents with agents. See the recent Scientific American article, The Semantic Web. While the notion of using RDF to embed meaning into documents is very powerful, it will require even more work than is needed in the arena of turning meta data standards into a controlled vocabulary of a corporations knowledge universe.

A necessary adjunct to personalization categorization is creating a catalog of tasks. Creating these catalogs is an evergreen project and is not something that can be done by a central group of catalogers or even information architects, but requires the kind of field research techniques of anthropology and knowledge engineering to observe people in their actual work environment.

Social contexts

As noted above, social contexts are much more important than personal contexts despite what The marketing literature of portal vendors might say. In fact, socialization is a much better term than personalization for what should be the focus of a portal initiative. While it is important how you organize and label individuals, the most important elements of those categories and labels are the ones designed to combine and merge sets of individuals into a rich and varied set of communities.

There are a variety of communities and dimensions within which to define communities. There are a variety of activities that define communities such as communication communities or collaboration communities. Some communities have well developed formal structures and are designed for long life spans, Other communities are informal and come together to support a single project limited in time and organizational space.

Communities can be formed around a variety of interest types, from the specific (Java Programming) to the general (Stock market), and from the social (Scuba club) to the organizational (Intranet Content Owners Network).

In addition, communities use a variety of channels to communicate and collaborate. One channels is a formal project or document collaboration using highly structured, well supported software platforms, such as content management or document management software. Another channel is online discussion groups with their balance of broad interest areas and an ever changing set of discussion topics. Another virtual community channel that is becoming more important is real time video conferencing. Finally, there is a whole set of very informal channels from real time, virtual chat (instant messaging) to real time real chat (water cooler meetings). Underlying and impacting all these channels is the one truly successful communication and collaboration tool, email.

In addition to developing organizational and labeling support for communities, Knowledge architecture also must deal with navigation and search. In this case, that means not only connecting people with information, it also means navigating through the people space of the corporation, connecting people with people. Again, as in the case of personalized newspapers, the crucial element of connecting people with people has to do with connecting communities of people, not just individuals. This means, for example, supporting navigation in and from within a community of practice.

Here again the close connection between learning and knowledge architecture is evident. Learning has been described as becoming an insider in a community of practice. What this means for KA is that it must support becoming a practitioner, not just learning about the practices of a community. It also means supported the process whereby an individual learns to view the world from the model of the community, the world according to client reps or java programmers or administrative assistants.

Knowledge architecture must support all these social activities. KA must model communities and support the communication and collaboration activities of those communities. KA must model, capture, and retrieve the tacit and explicit knowledge expressed inside those communities and expressed through their activities. KA must model individuals in a flexible organizational schema that will support modeling all the various types of communities.

One last point about architecture for communities. Communication based communities are relatively simple to model and support. The real challenge is creating an architecture that supports collaboration. A number of writers have noted that most collaboration efforts fail to produce the desired outcome. I've seen some figures as high as 75% fail. I believe that the fundamental reason behind this failure rate has to do with the lack of a good knowledge architecture. An architecture that not only models the informational activities of a collaborative community, but also its personal, political, technical, and intellectual infrastructures. How to create such an infrastructure and what components go into it, is the subject of the last section.

Knowledge Map

One of the more valuable products of a good information architecture is a content map of the information within a corporation. Whether its called a content map, a directory, or the corporate Yellow Pages, it rests on a powerful and flexible set of taxonomies representing a balance of global and local labeling practices and provides the means to navigate by browsing or searching through that content map.

Knowledge architecture begins with a content map and then goes beyond it to produce a knowledge map, a richer, more complex, more multi-dimensional, more human, more valuable, more – well you get the idea - more.

A knowledge map not only creates an organization and navigation platform for content repositories, it adds in repositories of tasks and people or communities. A knowledge map must provide a complete integration of the intellectual, persona, and social contexts discussed above.

A knowledge map can be thought of as a knowledge network mapped to a network of actors mapped it a network of tasks. A knowledge map is designed to answer that simple little question: Who needs to know what when and why do they need to know it? To answer that question, knowledge architecture must identify and model the who, structure the what and the when, and provide a platform for the members of the knowledge community to continually provide answers to the why.

How do you create a knowledge map. First, it is an evergreen project, it is never done and you can't design it from the beginning. The best you can hope for is to not have to have too many generations of the map produced. In other words, the growth in content is incremental and doesn't require significant new features too often.

A knowledge map is based on multiple taxonomies (and multiple types of taxonomies). It involves categorization, technology, and active field work with subject matter experts. It should use soft methods like a day in the life of a field rep and it should use hard methods like auto-categorization.

It should be a self referential project in that you need a knowledge map of the knowledge map project.

Finally, what can you do with a knowledge map. The short answer is:
Anything you want.

A slightly longer answer is to use it as an infrastructure that will support virtually all knowledge activities in the corporation and since there is a knowledge component to virtually all that humans do, it becomes a corporate infrastructure. Some sample activities might be a personalized and adaptive learning technical support. Another area of major impact would be in the areas of e-learning and best practices. a knowledge map can provide the framework for expertise locators and other community support tools.

There is no end to what it takes to create and maintain a knowledge map, nor an end to what you can do with one. I would add more and go into more detail, but I have to get back to work on our knowledge map.