Introduction

There has been an ongoing interest in story telling as a component in Knowledge Management over the last few years, but it has never really become a major focus. Story telling has been touted as the best way to make the leap from information to knowledge, and as the best way to capture and transfer tacit knowledge.

However, the idea causes some real disquiet among practical businessmen, information specialists, and even many in Knowledge Management. The image of a group of business people squatting around a campfire swapping stories is scary on a number of levels: Ripped and dirty suits and dresses, setting off fire alarms, poking people with sticks, and finally, huge amounts of wasted time as really bad story telling executives unload their fears and anxieties.

In addition to this type of comedic disquiet, there is another problem with the use of stories in knowledge management, which is, the knowledge embedded in stories is very difficult to codify in such a way as to capture the richness and multiplicity of stories without losing the immediacy and power of the story telling experience.

There also seems to be a huge disjunction between theory and practice when it comes to story telling in a corporate environment. On the one hand, the use of stories is flourishing. It is practical and direct. However the theory of the use of stories in corporations seems to be languishing between a deep academic theory that is unanchored in the practical reality of business and a strategic view of KM that looks askance at any attempt to capture and/or codify the living, breathing reality of actual story telling in day to day business. The net result is that story telling is divorced from any systematic foundation that would give it both a dimension of rigor and practicality.

I would like to suggest that one answer to both the fears of semi-mystical story telling rituals taking over the board room and the difficulty of truly capturing and representing the deep knowledge within stories is to create a rich and powerful knowledge architecture. This knowledge architecture must be organizationally powerful enough to overcome the flaky image of story telling circles and, at the same time, rich and flexible enough to represent the multi-dimensional nature of stories, allowing the knowledge in stories to be captured and indexed and made re-usable across multiple contexts,

This article will look at the issues and some approaches to creating such a knowledge architecture for stories. We will start with looking at why knowledge management should incorporate story telling and how stories are being used in corporations. We will then look at what knowledge architecture is and how it can be applied to story telling, what benefits can be expected, and how do you create meta-stories. Finally, we will look at some possible future directions that a good knowledge architecture can enable.

Story Telling and Knowledge Management

Why should Knowledge Management incorporate story telling?

Humans have been telling stories as not only a form of entertainment, but as a way to make sense of the world for a very long time – probably almost as long as they have had language. So it is not a surprise that we continue to use this powerful medium in the corporate environment. What is a surprise is how little we have incorporated story telling into knowledge management, although with the efforts of such projects as the IBM Story Project, that is changing somewhat.

We know that story telling is going on throughout the enterprise, but what is it about stories that we not only don't have to launch a program to get people to tell stories, we can't get them to stop telling stories even if we wanted to

Stories are a fundamental form of knowledge

First, stories are a fundamental means that humans use to structure the world. Our brains seem to be wired to easily and almost automatically organize information into stories. Listen to small children play and you hear the most wonderful stories being created, all without the benefit of major skills acquisition programs. Story telling seems to develop along with language skills and perhaps even before a sense of causality fully develops.

However, stories are not just for entertainment or children. Stories select events on the basis of importance and fit with other events and we then combine those events into an ordered and at least partially causal chain in order to explain the world around us and ourselves and our place in that world.

Stories also support chunking of facts and events in ways that correspond to how our brains are designed both for paying attention and for remembering.

Since story telling represents a significant form of knowledge, knowledge management needs to come to grips with the nature of story telling and how story telling is being used within the enterprise.

Story Telling is used throughout all businesses

A second reason that knowledge management needs to deal with story telling is that there is one certainty when it comes to business, whether you like it or not, story telling is going on in every business, every department, every team. Story telling is not only natural, it is being used right now throughout your enterprise and it is being used heavily, probably more heavily than any other information or knowledge sharing channel you have.

If we don't understand the activity and nature of story telling and if we don't understand what stories are being told in our enterprise, then we run the risk of creating a corporate environment that not only doesn't support knowledge transmission through stories, but could stifle or distort the use of stories in our environment.

True, people will tell stories given any chance at all, but the question is, what kind of stories will they tell? Will those stories get to the right people? Will those stories be used in ways that improve your culture or create an undercurrent of negativity? Will those stories teach new employees what they need to know or mire them down in the practices of the past?

Stories are particularly suited to knowledge management instead of information management.

While there are numerous variations in story type and use, they do share some common characteristics. One of the most important characteristic is that stories exist in the realm of knowledge, not information.

First, stories convey not information, but meaning and knowledge. The information they contain is seamlessly incorporated into the story through the use of context. And since stories create clusters or chunks of information, they are easier to pay attention to and to remember. It may be harder to codify knowledge than information, but it is easier for humans to remember knowledge rather than strings of unrelated bits of information.

Second, stories are particularly good at transmitting tacit knowledge. Indeed, given the difficulty in capturing and making explicit the tacit knowledge residing in your internal experts, stories seem to be the one way that we not only can, but easily do, capture and transmit tacit knowledge.

Third, listeners react to stories differently (and better) than to charts and logical arguments. Stories provide their own context which makes them more believable. In addition, people tend to hear stories in a more receptive mode according to some research. It has been suggested that the cognitive processes underlying story hearing are different. Stories are told within their own context, but the interaction between teller and listener is not one of transmission of information, but rather, the story becomes the means through which listeners create their own context, interpreting, and filling in blanks and links of the story.

Knowledge management can improve and be improved by story telling.

Knowledge management offers a number of advantages for story telling. First, it places story telling within the context of knowledge management thereby providing a framework of legitimacy. Second, research into the rich store of corporate stories can lead to better story schemas than many found in the current research. Adding the corporate context with its practical orientation can focus the organization of story elements in ways that often seem lacking in the academic literature.

Lastly, knowledge management can be improved by incorporating stories that are great examples of knowledge and the transfer of knowledge. Currently, there is a lot of

information management in knowledge management and a focus on the embodied knowledge within stories is a good way to enrich the architectures in knowledge management to include elements that go way beyond simple reference library systems.

What is Knowledge Architecture?

Before we look at what a knowledge architecture needs to be to come to grips with something as hard to codify as stories, I should say a little bit about what knowledge architecture is. For a more complete discussion of knowledge architecture see my article, From Information Architecture to Knowledge Architecture in Intranet Professional, September/October 2001.

However, the simple answer is that knowledge is information plus various kinds of contexts and so knowledge architecture starts with information architecture (organization, navigating, labeling, and retrieval of information) and adds different types of intellectual, personal, and social contexts.

A context is something that gives meaning and depth to information. Rather than try to define context further, let me tell you a story. Recently, I was at a doctor's office and the doctor came out and told a young woman that since she had changed her appointment from the following day to today that he wasn't able to get her charts. He repeated that information and waited for a reaction from the young woman who, since she didn't have the context that would give meaning to that piece of information, remained silent. She didn't know if the doctor's information meant the doctor wouldn't see her, if the appointment could be held but wouldn't be as productive, or if she needed to do something. The doctor knew the context which was twofold – Not having a chart meant that the doctor would be less effective and that he would have to work harder to elicit information from her.

The doctor transmitted the necessary information but not the necessary knowledge and the result was a complete lack of understanding. I can only hope that the doctor was better at communicating context in the actual interview. Perhaps if the doctor had a button that she could click on (Explain or More Info or What does that mean?). In other words, if there had only been a knowledge architecture supporting the interaction.

Knowledge architecture then is the attempt to create an intellectual infrastructure that can support the organization and retrieval of not just information but sets of related contexts around information, contexts that change over time and with different dimensions of applications. Knowledge architecture deals with a richer, more multi-dimensional intellectual universe of discourse and through that richer universe, must deal with the shifting chaotic world of applied information, i.e., contexts of actions.

Knowledge Architecture for Stories

In the case of stories then, knowledge architecture has two primary tasks. One is to create the intellectual infrastructure for deconstructing, capturing, indexing, organizing, and retrieving stories and elements of stories in a variety of applications and in a variety of communities within the enterprise.

We will look at some of the specifics of how to create this infrastructure in a later section, but in order to create this infrastructure reasonably well we need to understand how stories are being used in a corporate environment. And the answer is, of course, in a wide variety of ways.

One of the most basic ways that stories are used is in informal education and training. Stories are typically used to present the finer points of an area. They are not particularly good at transmitting the details and low level procedures that new employees need to learn, but once that basic context is learned, stories provide guidance and lessons in the advanced or more sophisticated application of those basics.

In other words, standard training tends to impart information, while the training embedded in stories tend to impart knowledge. For example, knowing which resources a call support person should use in answering simple questions is part of basic training, while knowing when to not look in those resources, but instead ask the real expert who sits in the other room, is the kind of knowledge that is often taught by a story.

Second, stories are the foundation for many formal and informal communities that form within an enterprise. The act of sharing stories creates the knowledge flow that makes a community alive and valuable. And the store of stories that are created become the foundation or context within which and from which, the community looks at the world.

Stories represent and/or contain the values and informal rules by which a community is organized. These community stories create an impact immeasurably greater and richer than a corporate or departmental newsletter. They use emotion and can engage listeners, individually and communally, in ways that mission statements will never equal.

A third use of story telling can be particularly valuable and it is what Richard Denning of the World Bank calls springboard stories in his book of the same name. Springboard stories are stories used to create a new paradigm or to not only introduce new ideas, but get people on board and actively promote the new idea. Denning argues that one reason stories are particularly suited to getting people to accept new ideas is the different way people react to stories versus charts and logical arguments.

There are certainly many more uses of stories in a corporate environment and, indeed, one task in the construction of a knowledge architecture for stories is exploring, identifying, and categorizing all the different kinds of uses of stories.

In addition, there are also a lot of different kinds of stories. There are anecdotes, myths, fables, and metaphors. There are cautionary tales (horror stories), success stories, lessons learned and hero stories, puzzle or detective stories, bonding stories and attack stories or

propaganda. You can find examples of all of them in a corporate environment, so if you are going to try to create a knowledge architecture for stories, it will need to be very flexible and rich.

However, a good knowledge architecture should also create a powerful infrastructure for supporting face to face story telling. This can include rewards for tellers of good stories. Is it really too much of a stretch to imagine "A very good story teller who functions as an important informal source of education and training within their department" as part of a performance review?

Another way that knowledge architecture can support all kinds of story telling, face to face and virtual, is to work with education and training to develop story telling skills and story understanding skills. In addition, capturing stories is not simply a matter of recording them. The skills of story crafting or story creation are also valuable and can be rewarded and they can also be used by a central KM group to create vivid and artistic and effective stories.

This means, not only knowledge architects categorizing and structuring the information contained in stories, but also performing such roles as knowledge facilitators or knowledge managers, and knowledge engineers. These additional roles would be used to support face to face story telling in communities and to facilitate the capture of those stories.

Should you create an architecture for capturing and retrieving stories?

Given the dynamic nature of stories and story telling in the corporate environment, it is clear that the more traditional reference library approach of information architecture is not sufficient to do it justice. So whether we call it knowledge architecture or information architecture on steroids, it needs new and innovative thinking to come to grips with the universe of stories.

However, there are a number of authors and speakers who argue that even if we can create an architecture to support capturing and retrieving stories that we shouldn't. They argue that the whole attempt is misguided and will kill the very thing it tries to support, that the attempt to codify stories will kill their impact, their magic, the very things that make story telling so powerful. Captured stories become static and lifeless.

However, on the other side of the question, I offer another story.

Recently I watched the movie, The African Queen, and just as I did the first time, the second time, and who knows how many other times, I was thrilled, uplifted, and damn if I didn't get all teary eyed one more time. Static and lifeless? I don't think so.

It is true that once you record a story or make a movie of a story, one component of the story becomes static, but story telling is not simply or even primarily a transmission from one person to another. One reason stories are so powerful is that they are common

creations of teller and listener and the components that listeners bring to story telling they can and will bring to re-constructed story telling.

One final point has to do with the relationship of face to face and virtual stories. Some authors suggest that as you move away from face to face, you lose the power and impact of stories and gain nothing in return. However, that is not necessarily the case.

For example, deconstructed stories may even be easier to remember and have as great or even greater an impact. Some experiments with stories as well as earlier cognitive studies suggest that sometimes abstract qualities are more powerful, much like studies on birds that found that chicks responded more strongly to cartoon beaks than to their own mother's beak.

Or as Peter Orton of the IBM Story Project put it;

"One of the most important yet least appreciated facts about story is that perceivers tend to remember a story in terms of categories of information states as propositions, interpretations, and summaries rather than remembering the way the story is actually presented or its surface features."

So I would argue that stories can be deconstructed, captured, indexed, analyzed, and retrieved and that the sum total of all this activity, if done well, would be to enhance, not kill the magic and power of story telling in the corporate world. The only real danger would be from an attempt to substitute the indexed, deconstructed story for the living, breathing, evolving story.

The outcome of all this analytical activity would be a library of stories structured and indexed to facilitate retrieval. This library can be used to compare stories or extract general common elements. You can create an abstract or summary of the story that will enable people to decide if that is a story they want to read-view-hear.

For example, someone searching for help on how to deal with customers with certain needs might launch a search within a knowledge retrieval system. A search result could return a list of links to documents and web sites, a set of related documents (and how they are related), and some links to some stories that illustrate how to deal with different kinds of customers.

Another use of a library of stories can be seen from the example of case law. Stories, especially stories that contain best practices and other lessons, can be used both formally and informally, as part of a real life education and training initiative. Just as lawyers study specific cases to gain both general knowledge and specific real world knowledge, so stories can be used to train sales reps or customer service reps or whatever.

There is, however, a unique aspect to a library of stories that information libraries don't have. Stories are fun and interesting to human brains and a library of stories becomes an educational tool just by its very existence. If it is set up right, people will browse and

read-view-hear just for the fun of it – and they will be learning lessons and culture as they have fun.

Another reason for capturing and deconstructing stories and creating libraries is the growing virtualization of businesses. As communities become more virtual, they will need stored and indexed stories to overcome their separation in space and time. Adding a collection of captured stories to a virtual community can enhance those communities and make them more powerful and effective.

Libraries of community stories not only enhance those communities directly, the act of creating a community library can reveal the rules and values of your formal and informal communities. You are then in position to use that knowledge to support those communities better and to better understand how to maximize the value of those communities within your organization.

How do we capture stories?

If we accept that we should capture and deconstruct stories, the question then becomes, how to do it – how should we represent stories, both in terms of capturing active stories in their native habitat, communities, and how go about creating effective stories that can function as a resource for all communities.

There are many ways to capture stories, some better than others. For example, these stories are often referred to as water cooler stories, so one method might be to install audio and video bugs at every water cooler to capture everything that is said. I can only think of a few dozen reasons why that might be a bad idea – nobody has water coolers any more, the job of winnowing out the interesting stories from the chaff of self puffery would be roughly comparable to a job monitoring the infinite number of monkeys typing on an infinite number of typewriters to produce the works of Shakespeare, all your employees would either stop talking or quit, you would be liable for the most interesting law suits, etc., etc.

OK, so that is one way not to do it. Some ways to do it are:

First, you need a central group to administer, meta tag, and facilitate story capture. This group could be an intranet team, a corporate communication team, or a dedicated knowledge management team. It could include members from all three plus people from your library and/or training and education. Some of the activities of this group could include adding meta data, maintaining a library of stories, teaching story skills, publicizing the role of the story library, and facilitating the story capture process.

Another essential component is to create a reward system for submitting stories to a central repository of stories mapped to the communities within which they are generated and publish the whole process on your intranet. In this case, the whole process includes authorship and rewards as well as the stories themselves. The process should take place

within a highly visible story telling medium – corporate magazines, recognition and awards, monetary and otherwise.

Capturing stories also means representing the content of stories in a variety of ways. For example, stories could be submitted in text format, but it would work much better if they could be done verbally. People will tell stories for hours, but ask them to write them down and 15 minutes seems very long. Until we have speaker independent voice recognition, that would mean having a person and/or camera on the listening end.

In addition, there are at least two very different ways of using multimedia to capture the richness of stories. The first is to create a movie that captures the story and/or exemplifies the story in a way that goes beyond simply a talking head telling a story.

These movies could be anything from edited actual stories (talking heads with music, visuals, and animation) to mini-movies that present the story. A variation might be filmed scenarios or case studies.

You might want to use human actors for some stories or you could generate a stable of animated characters that could be used and reused. Another option would be to offer the chance to star in a movie as a reward for submitting the story of the month with, of course, the option to bow out for the shy.

This might sound extravagant but with the development of web technology, web cams, relatively cheap and easy to use equipment and software, and advances in animation, a company could put together an internal team or hire external contractors for not much more than existing corporate communications teams.

A second way of using multimedia to represent stories is to create a multimedia representation of the elements of the story and their relationship.

One reason stories are so powerful is that they contain so much in such a small amount of packaging. To try to capture all the multidimensionality of stories in simple text would expand the story reading experience to mind numbing size, not to mention academic jargon overload. One way around this is to represent the multidimensionality with a multidimensional multimedia. For example, a hyperbolic tree representation of the relationships between elements of the story combined with sound and video effects as the reader/user explores the story.

Other than as a way to find employment for out of work video artists, why would you go to the trouble of adding all these elements to a story that can be told face to face in three minutes? Two reasons come to mind: virtual communities and the design of the human brain. Of the two, I find the latter reason the more compelling. Quite simply, our brains are designed to work better when more than one sensory channel is activated by incoming stimuli.

So, if this multimedia presentation activates more parts of the brain, two things happen, people can pay attention better and they can remember better. Those seem like good reasons to me – as long as we remember, that not all stories can be so treated – and the economic pay back spot will vary from company to company and over time.

Knowledge Architecture – Meta Data for Stories

Regardless of how stories are captured, stored, and represented, in order for anyone to find them again we will need to develop a reference or indexing system based on meta data of some sort. Just as with the case of representing stories, an index or meta data schema for stories will need to be much more than a simple library reference system. The meta data schema will have to be as rich and multidimensional as its subject matter.

However, a good place to start for any meta data system is still with the basics, in this case, the Dublin Core. While the Dublin Core is a good starting place, it is ultimately, both too much and too little for a meta data schema for stories.

Many of the Dublin core elements are geared toward the necessary meta data you need to create a well indexed reference library, but it doesn't support retrieval as well as it should except perhaps for professional librarians and searchers.

But for informal communities, and real time searching, that is, searching in support of job activities and procedures, it is not the right set. The Dublin Core is too much for authors to easily add the necessary values for such fields as Rights, and too much like book keeping to fill out values for fields like Format, Contributor, and the like.

The Dublin Core is too little, however, when it comes to actually finding items. Outside the reference library, the most important Dublin Core elements are usually title, description (or short abstract), and subject (usually referred to as keywords). For good searchers, date and author can sometimes help. However, as has been often noted, search using simple keywords is not very powerful.

In general, what is needed is additional categorization or taxonomies of content and/or a controlled vocabulary or other organized conceptual set of keywords. These are in the process of being developed either within enterprise intranet environments or by vendor companies that are offering automatic and semi-automatic categorization with built-in world knowledge expressed in predefined ontologies.

This is a good and necessary trend, but the situation with meta data for stories is even less developed and requires some different thinking. One reason for this need is the different structural and temporal elements of stories.

It is beyond the scope of this article to propose complete meta data schema's for stories, but I would like to look at three areas that could fruitfully be developed:

Meta Data extensions – to capture the sort of narrative schema's that researchers have developed
XML within stories – particularly to capture the temporal relationships of stories
RDF for representing the rich, multidimensional knowledge contained in stories.

There have been several proposed schemas that all stories follow. One describes stories as consisting of 5 stages revolving around equilibrium: equilibrium, disruption, recognition, effort to restore, results of effort to restore. Another schema uses different labels and subdivides two of the stages and comes up with the following schema:

Introduction of setting and characters
Explanation of state of affairs
Initiating event
Emotional response or statement of a goal by the protagonist
Complicating actions
Outcome
Reaction to the outcome

If we look at these stages as meta data fields, we might fill them with the following types of values:

Introduction – Description, Names of characters, keywords or concepts, subject matter or category
Explanation – description
Initiating event – Short Description (between description and entity)
Emotional Response – description
Complicating actions – Short description, keywords or concepts
Outcome – description
Reaction - description, keywords as part of an ontology of lessons or morals of stories

Having a set of meta data fields such as these could form the basis of an organizational scheme, but still would not capture the complex relationships found in stories.

XML
Another approach to the narrative schema above could be using XML to delineate the stages.
<Introduction></Introduction>
<Initial Explanation></Initial Explanation>
<Initiating Event></Initialing Event>
and so on.

XML could also be used with similar story elements or relationships. For example, <follows x></follows x> and <precedes y></precedes y> tags could be used to relate the timing of two events. A similar relationship though more dependent might be the <must include> relationship. If one element exists in the story, then the preceding one must exist as well.

Other relationships are conflict<>resolution and <supports> and <opposes>.

RDF
Meta Data is good for capturing certain aspects of the content of stories and XML is particularly good at structuring the sequential nature of stories, but RDF might supply a missing element which is a way of capturing the rich conceptual relationships within stories. These relationships include both subject matter hierarchies within which a concept or keyword is located and a typology of world relations.

A story that deals with lying, for example, with the right RDF description would also be related to broader concepts such as dishonesty and ethics. In other words, lying (is a member of> dishonesty and dishonesty <is a member of> ethics.

Adding a typology of world relations also helps provide a layer of structure that would greatly improve a user's ability to find what they were looking for. In our lying example, if "lying" was a keyword that described one of the events in a story, the types of keywords that lying was related to might include things like the results of lying, the morality of lying, lying about subject matter X, and so on.

Another important element of meta data for stories, whether expressed in RDF or some other format is distinguishing between concepts or keywords that are merely mentioned and those that are elaborated upon. This is an area that many of the automatic categorization companies are dealing with also. In that arena, it is referred to as the "aboutness" of a document and there are different approaches to characterizing the aboutness of a document ranging from a weighted set of keywords to bar codes to a hierarchical representation of the concepts within a document.

There is no clear answer for how to create a meta data framework for stories, but one thing is clear is that just the attempt to create one will have enormous benefits for knowledge architecture and knowledge retrieval. Stories provide a good test bed for dealing with the additional complexities of knowledge in that they are relatively well structured and there is a rich body of existing and new material being generated all the time.

Future Directions

It should be pretty clear from the above that we still have a long way to go to create a knowledge architecture that will support the full value of story telling in business. I offer the following, not as a specific goal or approach, but as an image of one way this might work in the future.

The first point is that a knowledge architecture for stories must be much richer than traditional library or information architectures. This richness permeates the entire process, from the multiple contexts of stories themselves to the variety of categorization

schemas that will have to be developed. In addition, both the representations of stories and the story categorization need to be much richer than currently available.

In addition, the representation of stories will have to be much richer than simple text. We've already seen how multimedia could be used to represent story components as well as directly present the story itself. And the representation of the retrieval process of stories would also be much more powerful if it went beyond the simple text listing of story titles, dates, and a brief description.

In fact, story telling retrieval systems could be good candidates for new computer interfaces such as those envisioned by David Gelernter and others as they look beyond the standard desktop metaphor of files and folders.

Imagine a browse interface into a library of stories. It might consist of a graphical map of a world. The world could include countries (communities), cities (stories clustered around communities), natural resources (stories clustered around categories of various types), roads or connections between cities that could be customized by the user or the community that the user belonged to. In addition, there might be actors moving around the map performing actions – such as mining some resources, exploring the connections between stories or even representing a particular story teller.

Double clicking on a city would open up another display that might show a structured set of resources, charts and spreadsheets, and a decision panel that was personified as the governor or librarian of the place, an avatar who could offer help, information, and act as an agent, launching searches out into other information spaces and digesting the results according to rules the user or community could set up.

Of course, one response to this image might very well be, why on earth would you go to all that trouble, setting up complex graphics, etc., when simple interfaces can do the job? As we saw in earlier discussion, the answer lies in the way the brain works – we remember things better when more than one sense is involved. In fact, we do most sensory tasks better if more than one sense is involved. This is particularly true when it comes to stories rather than simple informational documents.

Another future direction might be to use stories or narrative schema to improve learning in a variety of situations even including one that are not really stories.

As one IBM researcher on the IBM Story Project put it:

"Narrative schema may be applied in many situations that are not story related….It seems that narrative schema is an option in processing data even when there are no human characters or the events are essentially nonsense. …we may see stories in random material, information and data where stories do not otherwise exist."

One example where story elements are already in use is with the use of persona's to organize information about certain key common elements. Rather than ask a user to

wade through a series of specific questions about what information they need (or typically need), simply ask them to find the closest matching persona. Persona's are already in use in areas like financial services and computer sales sites.

There are many directions that story telling in a Knowledge Management environment could take. However, underlying any new direction should be a rich and powerful knowledge architecture. Without this architecture, story telling will likely continue to languish either in abstract academic research white papers or hidden in the undiscovered byways of personal interactions within corporate communities and knowledge management will miss the opportunity to extend its scope and depth by incorporating one of the most heavily used knowledge transmission mechanism with corporations today.