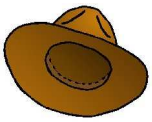

Dr. Tom's Method of Multiples

A Case History of Concrete Taxonomy Development Method

Revised 02 June 2006

Thomas D. Wason, Ph.D.

tom@twason.com
www.twason.com

2006

Summary

The development of metadata and taxonomies is a human endeavor. It is people who will use this information, either directly or indirectly. This is true of many human activities. The method for developing a system of metadata and taxonomies described here is a specific implementation of the general principles that can be applied to many endeavors. An effective metadata system can be established with the participation of **multiple** teams each with a different perspective, the subject matter expert (SME) teams. Each SME team is comprised of **multiple** members. The SME teams are given a carefully chosen concrete task that spans their different perspectives. As they work on the task in facilitated joint meetings, a taxonomy team records the comments of SME teams. The taxonomy team is comprised of **multiple**, independently tasked recorders. The intent is to define and capture metadata and taxonomy definitions from each of several different vantage points. Each recorder provides separate reports that are consolidated into a single report with resulting recommendations for metadata and taxonomies. These recommendations are then validated by an independent set of SME participants. A case study using this method is presented. The results are compliant with SCORM, IEEE-LOM and IMS-MD specifications. The method is moderately ethnographic.

If knowledge is power, new knowledge is command.

When new knowledge is generated,
one of the first places it shows up is in a taxonomy.

Contents

1. Introduction

The development of metadata and taxonomies is a human endeavor. It is people who will use this information, either directly or indirectly. This is true of many human activities. The method for developing a system of metadata and taxonomies described here is a specific implementation of the general principles. Let us first proceed with those principles, framed relative to metadata and taxonomies, and then look at a specific case applying those principles.

2. Principles

The principles for metadata and taxonomy management have several components: opportunities, purpose, people, task focus, body and validation. There is an underlying fact.

> A rule of metadata and taxonomy development is:
>
> **Whatever you do will be wrong**.
>
> Accept that from the start.
>
> Your process should come up with a "good" solution and a method for continually improving the results.

The method described here does not discuss the use of tools. Tools are important, however. They can search out new terms in the resources that can be tested against the existing taxonomies and thesaurus. A principle is to manage taxonomies based on definitions. New terms can be defined and the terms within the definition can be tested against the taxonomies and thesauri. Truly new terms may reveal new knowledge within the enterprise. The method described here can be augmented with tools for both terminology discovery and taxonomy management. Indexing is a different activity that will not be covered here. People are more apt to know what the do than what they need. Often the needed components are so automatic as to be invisible while performing the task. This method has its emphasis on "do," although participants are encouraged to express needs.

2.1. Opportunities

A primary objective of this method is to capture knowledge. Taxonomies are of great value to an organization. A taxonomy is useful for organizing and retrieving resources and information. It can be the basis for knowledge (KM) management. The management of a taxonomy is a source of great value to an enterprise.

If knowledge is power,
new knowledge is command.

When new knowledge is generated, one of the first places it shows up is in a taxonomy.

I used to do quite a bit of patent work for clients and for myself. This often involved searching for "prior art." Before the current computerization, this would sometimes involve going to the patent office to search; a most enjoyable activity. The patents were in boxes—called "shoes"—racked in shelves. After narrowing down the classification with a computer directory I pulled a shoe off the shelf, took it to a desk, and leafed through the cover pages. I soon learned that I was getting into recent knowledge when the material was not available because it was being reclassified. Think about the first transistor patent: there is no category of "semiconductor devices." When a few of these patents have accumulated, reclassification is required. Typically one examiner is in charge of a particular area, and he or she maintains a personal library for reference, including patent cover sheets, during the reclassification process. This person cannot tell you about patent applications in process, nor should he or she. But the fact that a reclassification is underway tells you that you have hit a "hot" area. The new classification(s) tells you something about a dynamic area of work. This can be the basis for new knowledge.

As a general rule, new taxonomic development occurs at the edges of the existing taxonomies. The terms represent new knowledge. People are debating what to call the new knowledge or what external knowledge is to be imported into the enterprise. The core taxonomies, although useful, represent old knowledge that has already been integrated into the enterprise. I am omitting obvious knowledge management and data mining activities as we are discussing new concepts developed within an organization. For an organization, a main purpose for managing the evolution of its taxonomies is the recognition of the primary new areas of interest to the organization. A core change in knowledge will usually be reflected in a change in the metadata structure. Some aspect of the facets may be subdivided into multiple separate facets. New facets may be required. Old facets may no longer be useful. Relationships among facets may change. All of these present legacy issues of course. The key is to capture the knowledge. That is critical to the management of the enterprise. We should always keep in mind the purpose of the metadata and taxonomies. The role of taxonomists in this particular method is to help extract the information from the people who use it. A general principle people are more apt to accurately perceive what they do rather than what they need. Our objective is to let people reflect on what they do in a way that reveals what they need and carefully record their activities. For an enterprise, management of its taxonomies provides a window on important new knowledge, and should be conducted by its own staff using methods that expose this knowledge. An external consultant may be useful in setting up an effective process. As the method described uses a taxonomy team,

setting up the process results in the training of internal staff. This approach has some ethnographic elements although it is artifically constrained by the task and the subjects know they are being observed. The occasional interaction of the recorders also violates pute ethnographic principles.

2.2. The trust problem

A major problem is trust in the quality of metadata. When IMS (www.imsglobal.org) set out to develop specifications for online education one of its first specification was for metadata. As I led that effort one of the first questions we had to address was the reliability of the metadata. During the requirements meetings people said they could not evaluate the metadata itself for its reliability; instead they would look at the source of the metadata. That information is carried in the IMS (and the IEEE LOM and SCORM) metadata structure[1]. Within an enterprise the quality of metadata can be controlled best by incorporating its generation within the workflow. We'll look at that more below.

On the open web metatags (little metadata-like bits in the HTML code) are of little value due to their unreliability. Within an enterprise, this need not be the case. Ensuring the quality of metadata is important. Its value approaches that of the resources they are attached to. Expend effort on it. Is the metadata any good? Are the terms used properly? How do you know? These questions are largely out of scope for this paper. If you don't address them the exercise of making metadata is not very useful. In fact, untrustworthy metadata is worse than none. The process of creating meaningful metadata should be part of the metadata system. The inclusion of users in the metadata structure and taxonomy management process provides an opportunity to develop the framework for a solution. It cannot be stressed too much: **ADDRESS THE TRUST PROBLEM**.

2.3. Summary of Method of Multiples

An effective metadata system can be established with the participation of multiple teams with different perspective, the subject matter expert (SME) teams. Each SME team is comprised of multiple members. The SME teams are given a carefully chosen concrete task that spans their different perspectives. As they work on the task in facilitated joint meetings, the comments of SME teams are recorded by a taxonomy team. The taxonomy team is comprised of multiple, independently tasked recorders. The intent is to define and capture metadata and taxonomy definitions from each of several different vantage points. Each recorder provides separate reports that are consolidated into a single report with resulting recommendations for metadata and taxonomies.

The method illustrated here for developing a metadata and taxonomy has a few basic steps:

1. Design a task for people with different viewpoints.
2. Have multiple people with different viewpoints work on the task together.
3. Record what they need and do from multiple viewpoints.
4. Make a metadata structure and taxonomies.

[1] That information would need to be digitally signed if not under enterprise control.

5. Validate the metadata structure and taxonomies.

The method is moderately **ethnographic**, recording the behaviors of the SMEs and subsequently analyzing them for metadata, taxonomic elements and new knowledge (Figure 1).

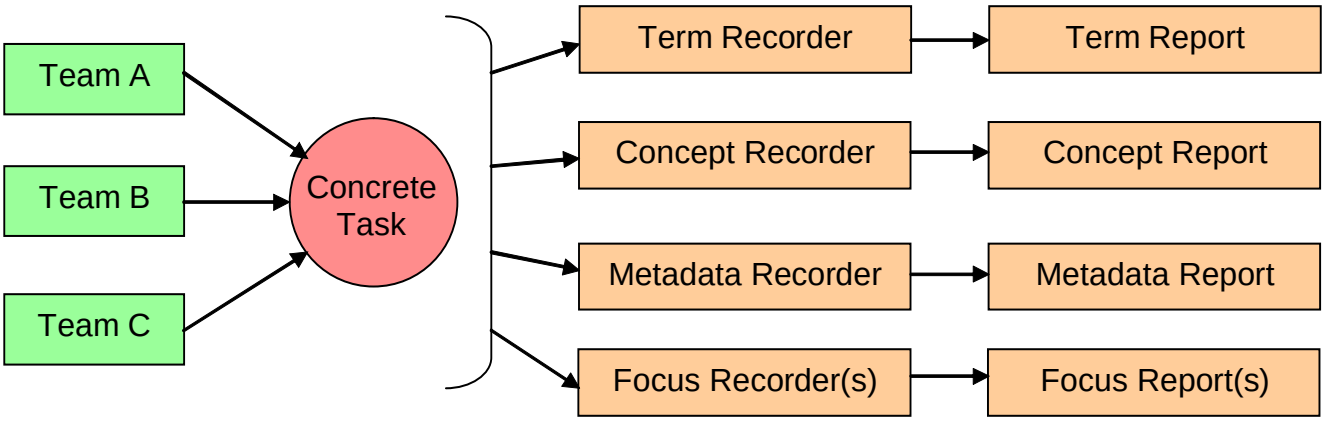

Figure 1. Method of multiples with multiple members of multiple teams collaborating on a conrete task while being observed by multiple, independently tasked recorders.

Of course nothing is ever as simple as this. So let us turn to an actual case history to see how this plays out.

3. Case history

The director of the US Air Force Institute for Advanced Distributed Learning (AFIADL), working with her team, was setting up a system to convert a large corpus of essentially paper-based (PDF) instruction in to online instruction. These courses are widely used in the Air Force, delivering some 140,000 course per year. This same system would be used for the creation of new courses. The AFIADL team recognized that good online instructional materials are expensive to develop. In the interests in providing a good return on investment, and creating some uniformity across the Service, the team recognized the need to reuse materials. In order to do so, it would be necessary to provide metadata on the resources that would allow course developers to discover them without being overwhelmed by a large number of hits.

AFIADL instituted a project with Boeing Training Systems & Services and Teleologic Learning Company. I was brought in by Teleologic to lead the project after it had started. The first thing done was to redesign the project. The original objectives of creation of a metadata and taxonomy system to support the reuse of instructional resources were maintained. The original ethnographic method origina by was significantly modified from being a series of field observations to being a directed colloaborative activity with independently tasked recorders. AFIADL, Boeing and Teleologic all accepted the reformulation of the project plan. The project centered on a collaborative meeting. A considerable amount of work was done before the meeting.

3.1. Task design

The project plan was centered around a group of subject matter experts (SMEs) from different areas collaborating on a common task. Our objective was to define a task that spanned several important domains but had some commonality. Note that we designed the task before we attempted to recruit participants. This is a **key** part of this method. Working closely with several AFIADL personnel we decided that a task of locating a reusable instructional module (RIM) on **Basic Electronics** to be used in several courses would span several areas (Figure 2).

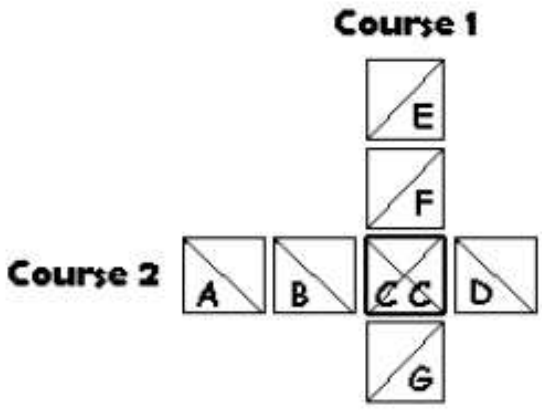

Figure 2. A module (C) common to two courses.
From AFIADL Final Report 4TFG57030729

The AFIADL staff searched the corpus of career development courses (CDCs), locating a number that could or did have a section on basic electronics. From these we selected three courses from quite different areas:

CDC 2A351, Avionic Systems
CDC 4A251A, Biomedical Equipment
CDC 2E151, Satellite/Wideband

All of these courses existed only in PDF format and had been developed independently. The reusable module was defined as that in the avionic systems course. The task for the group was to discover and fit that module into the appropriate courses. It was understood that the fit would not necessarily be ideal, and the adjustments necessary would also be of interest, as the overall objective was the reuse of resources.

3.2. Team collaboration

The participants were recruited by AFIADL based on the task needs; specifically, they were subject matter experts who developed courses who had direct experience either with the courses or with the specific subject area. The participants were recruited from three Air Force bases (AFBs). The participants from different bases did not know each other in advance. The participants were told that there would be a multiday meeting in Montgomery, AL at Maxwell AFB. The pre-meeting materials instructed the participants on the task. They were provided with a summary of metadata concepts. The complete courses were made available online and as email attachments. The AFIADL coordinated the meeting dates—no small feat.

The team was given the basics of metadata and taxonomies both in advance of the meeting and at the beginning of the meeting. This instruction is summarized in the Appendix (Section 5.). The concept of reuse was also discussed, using the model described above (Figure 2).

3.3. Record

A project team recorded the collaborative activities in the meeting. The project team consisted of a facilitator and two or three recorders. Each recorder was tasked to focus on a particular area:

- The *Terminology Recorder* will focus on the terms used, indicating context or brief definition. The recorder will have pre-made forms to work in.
- The *Concept Recorder* will focus on concepts and content. This will include the definitions that will underlie the taxonomies, course content, and the concepts to be represented by metadata fields.
- A *Reuse Recorder* may be present. This recorder will focus on factors relating to reusability of a module and the nature of the contexts. This is the tarrget focus. This recorder is optional. If not present, the Concept Recorder will perform this task.
- The *metadata facet recorder* will focus on the facets of the resources. In this case, the facilitator also served as the facet recorder. I performed this role.

The concept and terminology recorders had each prepared worksheets for capturing information. The concept recorder's worksheet had a simplification of the SCORM metadata structure with spaces to record the associated concept, team and time. The terminology recorder's worksheet had spaces for the terms, teams, equivalent field of use, definition and time. Time was captured to assist in coordinating the results of the terminology and concept recorders. The recorders were minimally intrusive, only asking for clarification when needed. AFIADL also supplied instructional systems specialists (ISSs) and administrative personnel. Three of the four recorders were expert in online course development. The facilitator (me) had less experience. Two of the four recorders were expert in metadata systems; two had only limited knowledge. These two were educated in the principles of metadata.

The meeting was held for two days. Most of the participants arrived the evening before the meeting. The meeting was carefully constructed and managed. The meeting room was arranged in a "U" of tables with a video projection screen at the open end. There was a row of chairs behind the chairs for the "U." Typically AFIADL staff used the rear chairs to be unobtrusive. A blank flip chart was provided. A computer was connected to the video projector. An open notes format was used. The facilitator could make working notes that were displayed. Each participant had a name card on the table.

The meeting started with introductions. It was soon apparent to me that use of first names of military personnel resulted in unease. Use of rank, e.g., "Sergeant Smith," relaxed the group. As most of the AFIADL personnel were civilian employees, the converse was true for them. We worked in a manner to make the military personnel most comfortable. After the director of the AFIADL made introductory comments, I made a brief presentation outlining the task and, using the dreaded PowerPoint (PP), summarized the concepts of metadata and taxonomies. Although the results needed to be conformant to the SCORM standard, no mention was made of this standard. To a degree, the process was a test of the adequacy of the specification. It passed. A minimum number of slides were used, as the intent was not to be a lecture, but to help frame the task. The PP slides were mainly illustrations similar to those in the Introduction (**Error! Reference source not found.**). The

reuse recorder (Boeing) made a presentation on the principles of reuse of instructional resources. Again, limited use was made of PowerPoint slides.

The work process required significant initial facilitation. As it was explained to me, the enlisted personnel were accustomed to following orders. A free-form process was not a normal work model. For much of the initial part of the meeting I stood at the front of the room. We knew we had "lift-off" when the team members started explaining the process to each other. It was a good group that really got into the spirit of the exercise. One member was having some difficulties with one of the concepts; the other participants all joined together to explain it to him until he understood. This was a good time for me to sit down and be quiet.

The team was encouraged to focus on definitions: what concept did they intend to convey? I have found this to be a very effective mode for people to achieve agreement on metadata and taxonomic terms. People consistently agree rather quickly on what they mean to say. The definition is recorded and then the group works on the specific terminology. Interestingly, the spirit of agreement that ensues from agreeing on meaning carries over in a spirit of collaboration in deciding on the appropriate terms. I have seen this happen a number of times, and am always amazed at how well it works. I guess it shows how we are bound to the definitions we have for our personal or local terms, but concepts, which are more abstract, are more accessible for negotiation. I have used this process in the harmonization of international specifications.

The team was led with open-ended questions of the form: "How would you describe the course module you want to put into the course?" "How would you find it?" "What would you want to know about it?" This revealed a great deal of detailed information. For example, the skill levels to be achieved were "proficiency codes" in a document (Career Field Education and Training Program, CFETP) that precedes the development of each course. The CFETP document is developed at a meeting called the Use and Training Workshop (U&TW). A key part of the document is the Specialty Training Standards (STS) or Course Training Standards. This was found to also include a fine-grained hierarchical table of the topics and skill levels required to be taught. The STS thus provided the basis for two taxonomies: Skill levels (a two level taxonomy) and topics (which could be up to 3 levels). The team members brought out a number of resources that were pertinent. In the end, all metadata field taxonomies were found to be pre-existing, one in particular was extremely rich. It is used widely across the military. When this major taxonomy was mentioned, all of the participants immediately agreed that it was the right one to use. This was gratifying to all. A comment form one of the participants was "We knew the answer, you just we're asking the right question." This seemed be a consistent theme: once they figured out the right question, they could find an answer. Not once did an answer go in search of a question. When I mentioned that the major taxonomy (called the Technical Order Numbering System) was extremely well organized, the AFIADL director commented, "It's the military." She was kind enough not to add the obvious "stupid." More detail on metadata development is beyond the scope of this paper.

Interspersed with the discussion of the fields and taxonomies needed to discover resources were discussions of the process of developing a course and how this would be accomplished to enable reuse. The group proposed a model of course development that surprised the

instructional designers. Sketches were made. At several times the staff revisited the model; always the response was that the SMEs could use the model. Everyone recognized that this would make resource reuse much more feasible. The group proposed some metadata fields that would support this, and promptly provided vocabularies. It was realized that the use of metadata needed to be **integrated into the workflow** for both resource development and course development, including resource discovery. This was new, and valuable knowledge.

Each recorder captured information according to his or her focus area, including the team that was the source of the information. This was not always possible due to the rapidity of the interchanges once the meeting got up to speed. This division of focus resulted in four different perspectives on the discussions. The recorders were able to ask questions to clarify points as needed. Although this happened infrequently, it was an advantage over making a videotape of the meeting for subsequent analysis. Videotaping is intrusive in its own way. The individual recorders each prepared reports without consulting with each other.

I consolidated the individual reports a single report and submitted it to the other recorders for comments and edits. This final meeting report was then submitted to AFIADL. The individual reports were included in the final report to have an "open" process.

3.4. Develop metadata and taxonomies

From the notes and final report a set of metadata fields was developed with proposed taxonomies. Most of these had been developed during the meeting. These fields were then mapped to the Sharable Content Object Resource Model (SCORM) specification metadata mandated across the US Department of Defense (DoD). There were no difficulties in accomplishing these mappings. Similarly, the reuse model developed in the course of the meeting mapped nicely into the SCORM specification. As the SCORM metadata is closely related to the IEEE-LOM (Learning Object Metadata) and the IMS MD specifications, it is also compliant with those specifications. A simple spreadsheet was created for test purposes. Metadata was created for the Basic Electronics module. Boeing developed a prototype metadata tool that interfaced to the AFIADL metadata repository. This was configurable via external files. Thus local names for fields could be presented to the user, but mapped to the SCORM specification when communicating with the repository. Configuration files also allowed specific fields to be exposed during different parts of the workflow. Proposed workflow integration is beyond the scope of this document.

3.5. Validation

The metadata and taxonomies needed to be validated. Due to the size of the project, an in-depth validation was not possible. This was accomplished in two ways. A Teleologic course developer working on the conversion of an AFIADL course to an online form was provided with the table of fields and taxonomies. She was asked to create the metadata for several modules she was developing. Following some instruction, she accomplished the task rapidly—well within an hour. The project team and AFIADL reviewed the results. All were satisfied with the results. This validated the ability of a tagger to apply the metadata system.

Boeing's prototype tool and a small database were used for testing the metadata and taxonomies. The AFIADL leadership team, the Teleologic team and the Boeing team

developed metadata for sham objects in addition to the real Basic Electronics module. The metadata was constructed to be quite similar to the target metadata, but differing in some key aspect: user level, instructional content and so forth. The development team (not the participants) created the sham metadata. This metadata was placed in the database using the prototype tool. A new set of participants was recruited to test the metadata system. Their task was to find the test object using the tool. A detailed description of the need was provided without using any of the metadata or taxonomy terms. They were instructed to conduct searches until the search returned only a good test object. Their activities were tracked and recorded for subsequent analysis. As it was not possible to recruit many new participants, the original team members were also invited to attempt the searches.

The following is extracted from the final report (Wason, T., & Perrin, B. (2004) AFIADL Taxonomy Project Final Report GSA Task Order ID: 4TFG57030729 Dec. 2004 http://www.maxwell.af.mil/au/afiadl/adl/TaxonomyFinalReportDec2004.pdf):

> Six users attempted to find a target resource using a prototype online search tool developed by Boeing. The principal objective was to determine which fields the users favored. Of the 26 unique fields, 5 were found to result in 50% of the use. 80% of the use was restricted to 12 fields. The success rate was not high, owing to multiple factors associated with the prototype status of the tool and the limited data set size. Since the objective of this evaluation was to explore the use of available fields, this did not constitute a major problem. In fact, the difficulty in achieving successful results produced an abundance of instances for analysis.
>
> The difficulties with the metadata tool and the small object data set resulted in a large number of attempts from participants. This produced a large data set for analysis from a small population. A small number of searchers (6) created a large number of searches (257), averaging 42.8 searches per searcher.
>
> The number of the six participants that actually searched is not known. Of the 257 searches performed, 145 produced some results (56%). Of these, 21 included the target learning object (#38) in the result (14% success in searches that returned any results, 8% of all searches). 4 searches achieved the goal of returning only one return, the target (1.5%). Three (3) of the participants reported never getting any hits on any searches. Five (5) of the six participants responded with emails. One responded with a simple "done!" The others had comments. The comments generally expressed a dissatisfaction with the experience. Optimization of the tool user interface should be a subject of future work.

This validation test was indicative but not conclusive of the validity of the metadata set. It was concluded that with further study only a subset of the original metadata fields would be adequate.

4. Discussion

To summarize, here are the components of the method:

- Multiple teams with different perspectives
- Multiple members of each team
- Multiple independently tasked observers.

These multiples focus on a common unifying task.

The case study illustrates an application of the method, which evolves with experience. Anyone who has worked on such a project recognizes that:

> The project MUST have a champion in the organization.

AFIADL had a strong champion who built enthusiasm in her skilled team. The AFIADL project exceeded everyone's expectations. Even mine. We developed results that had far broader utility than expected. We surfaced strong taxonomies that we didn't even know existed. The participants created a model for course development that was unexpected. The independently targeted recorders each provided viewpoints that we could consolidate into a stronger overall picture. The team members built bridges across domains, working together on a seemingly straightforward concrete task. Everyone learned something from the exercise.

It is important that the objectives first be defined; it is too easy to let these things to either be open-ended or to be limited. Obviously the first things to be determined are the purposes of the metadata and taxonomy systems. Typically that beast called "Upper Management" declares what these shall be in very broad terms, for example:

- Let us provide our customers with information they need.
- Let us provide the right resources to the right person or system at the right time
- Help our suppliers provide what we need.
- Let us manage our resources effectively.
- Let us reuse resources we have already created.
- Let us manage information effectively.
- Determine what our competitive advantages are.
- Discover commonalities across divisions within the organization.
- Let us discover each other when appropriate.

These purposes are good things. You will think of others. These purposes provide the vision for the endeavor. You can't pick them all; rank them. Pick the top **3**. They should support each other. Now the purposes have to become a reality. In the process the enterprise must embrace the process and the results. If Upper Management does not support this effort, it is purely an academic exercise.

Back in the days of the nice librarian behind the desk there were two models relative to resources: keep them neatly organized on the shelves or keep them circulating. The use of technology has removed this split. One must ensure that it doesn't creep in again. An overriding purpose of metadata is to keep knowledge and resources in circulation. Neat digital "shelves" are not the objective. It is easy for taxonomists to think in terms of neat organizations. It is the nature of taxonomists to be "neat." None of us wants to intentionally impose our sense of neatness over the users needs. Remember the Law Librarians' Discovery (Section **Error! Reference source not found.**). My general orientation is toward smaller taxonomies with a richer thesaurus. Both must be searchable. This approach focuses on the definitions of the elements in both the taxonomies and the metadata structures (i.e., fields). The method and subsequent **maintenance effort** must include a plan to capture new material that represents new knowledge within the enterprise. A plan for consolidating this information and sending to the appropriate users should be in place.

Potential boundaries of use need to be established from the start. The Upper Management should define an intended **scope** of use of the metadata system. It should also define or discover how it will interact with pre-existing systems. Typically the objective is to start locally but to create such a compelling system that everyone wants to use it. This will not happen. People want ownership, particularly of such a core function. Develop an initial **focus**. Over what scope of the enterprise's activities can one hope to achieve use of the results? It will always be more limited than you hoped. Think in terms of what should be a clear win, and then expand the scope later. That is not a new idea, but we all forget it, don't we? Something about the "buffet of life" comes into play here, as we load up our plates with lots of different tasty things, far more than we can eat (Figure 3).

http://www.metroweekly.com/articles/attachments/2003-09-18_arts_and_entertainment_647_792.jpg

Figure 3. A loaded plate.

It takes discipline to restrict the initial scope to what can be reliably accomplished. It is far better to shoot small and be successful, as you are attempting to recruit the organization, which will be watching to see if the effort falls on its face. There should be an **expansion plan**. How will this project extend into other domains within the enterprise? What will be the obstacles? One would like to use the same initial methods to expand the system within the enterprise. You will need champions elsewhere. This will mean building bridges from the original participants to new participants, creating new teams and going through the exercise again. The original participants will serve as proselytizers.

There is a tendency to seek out those people who will willingly participate. This is preaching to the choir; gratifying, but not very productive. What you want to do is engage people from several different domains around a task that will unite those domains. For example, if sales, marketing and production need to work together, it is important to select a task that each domain can relate to. Each will relate to it differently. This is good. It is core to the method. In the case study the three domains were represented by different courses. Clearly we could have sliced the pie in different ways (I must be hungry).

The method focuses first on developing a **concrete task**, and then recruiting participants. Once the decision makers have determined which domains within the enterprise will have use for the taxonomies, a selection among those domains is made. Initially limit the participating domains to some manageable number of 3-5. More that this will be like herding cats, particularly on your first attempt. In the case study (see above) we had three participating groups; this was quite effective. Fewer than three does not include useful differences of need. Selection of the initial participating groups is not a trivial exercise. How do they represent the core activities of the enterprise? Do they have enough differences to provide a rich interchange? Are the differences so great that there will be no common ground? I would not say the selection is an art, but it certainly requires good judgment. Having selected the primary participating domains, create a task that spans the domains. This is not trivial.

I have found many times that focusing on definitions rather than terms is extremely useful in getting people to work together.

Terms in a taxonomy have only **one** definition.

Once the definition is agreed upon, then you can work on the terminology. There is one "master" term for each definition. There may be a thesaurus that provides synonyms. Don't try to slice the bologna too thinly. Terms (and definitions) should be distinctly different for the user community, not just for the taggers. Remember, metadata is a communications system between tagger and user.

In the process of the case study it was clear to all that the creation and use of metadata must be integrated into the workflow. Although tools are nice, you can get by without them if you have to. I worked on a project with Indiana University in which we could see the need to create metadata during the creation of the resources. There was no budget to do this, and it was unreasonable to expect the authors to learn to code metadata in XML. We devised a Word template to collect metadata during the workflow; an XML specialist would then convert this into a standard format using an XML Schema. In the AFIADL case we recognized the need for a powerful tool; Boeing's prototype was built on this principle. It is clear that a good workflow metadata tool would be valuable. A tool could be configured for both resource creation and resource use. I have felt this to be true for sometime; I think XML provides a good underlying data form.

A metadata system management plan must include a validation plan. How do you know you have made the right stuff? A good method might use groups:

1. a group creates metadata for a group of objects
2. a group creates metadata for a group of objects
3. a group attempts to use the metadata for the objects

The members of groups 1 and 2 have the same task with the same objects. In the interest of creating metadata for a large number of objects, each member of each group does not attempt to create metadata for all objects. Study inter-tagger reliability. The objective is to create a significant number of objects for different uses. Group 3 attempts to discover or use the objects via the metadata. Group 3 is given some narrow task and must find or use the appropriate objects. Determine the number of attempts made by the members of group 3 to accomplish a narrowly defined task.

I've not discussed a maintenance plan. Clearly you need one. It should continue to extract new knowledge. It can be built on the same concepts described here to set up a metadata system.

Ultimately, the success of the project relies on people and preparation. It takes effort to prepare for the critical meeting(s). In the AFIADL project the project team met with the client well in advance of the meeting to discuss objectives, the plan, the timeline and other logistic elements. You have to do it. The meeting itself, once planned for, uses many methods from furniture arrangement, open notes and simple framing presentations, to an understanding of the social dynamics with continual adjustments. It is important to

recognize significant issues—such as the method of course development fostering reuse—and to pursue them in a positive manner. It is also important to recognize resistance and to address it carefully without alienating anyone. You do not want a polarized group.

Ultimately, with an organizational champion, reasonable objectives, good preparation and alert participants, this is fun. Really.

Tom Wason
tom@twason.com
www.twason.com
+1 919.839.8187
+1 919.602.6370 (cell)
1421 Park Drive
Raleigh, North Carolina 27605 USA

5. Appendix

Let me first introduce some basic metadata and taxonomy concepts. Imagine going into a grocery store in which none of the objects have labels (Figure 4). You would have to open all of the cans and boxes to see what to buy. The grocer might not appreciate this. It's messy, too[2].

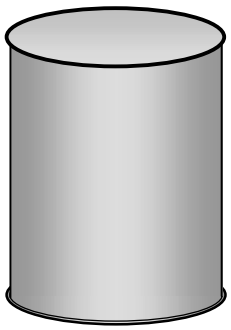

Figure 4. An unlabeled object.

Labels, particularly with some standardization as to listing contents, make the commercial grocery process possible. The packager can tell the end user what is in the can even though it goes through a complex distribution chain. Clearly, it is useful to label objects (Figure 5).

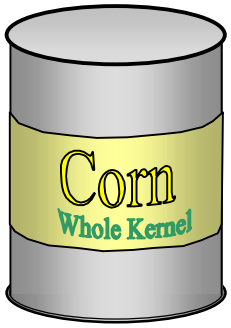

Figure 5. A labeled object.

Discovering appropriate digital objects can be a daunting task given the large number of resources available. Many resources are labeled to support discovery; these labels are generally referred to as "metadata." Metadata, or "data about data," provides a number of bins or fields into which values are placed. The labels are sometimes referred to as "tags." More properly they are referred to as "facets." A collection of related objects will all use the same kind of bins. Each item has its own set of bins, but all of the items will usually use the same kinds of bins. For example, every object may have a "subject" bin. A defined set of terms may be associated with each kind of bin. Labeling an object then entails selecting a

[2] These have data management analogies.

term from the set of possible terms for a bin and placing the term in the appropriate bin. You can think of these as cards drawn from a deck, dropped into bins, Figure 6.

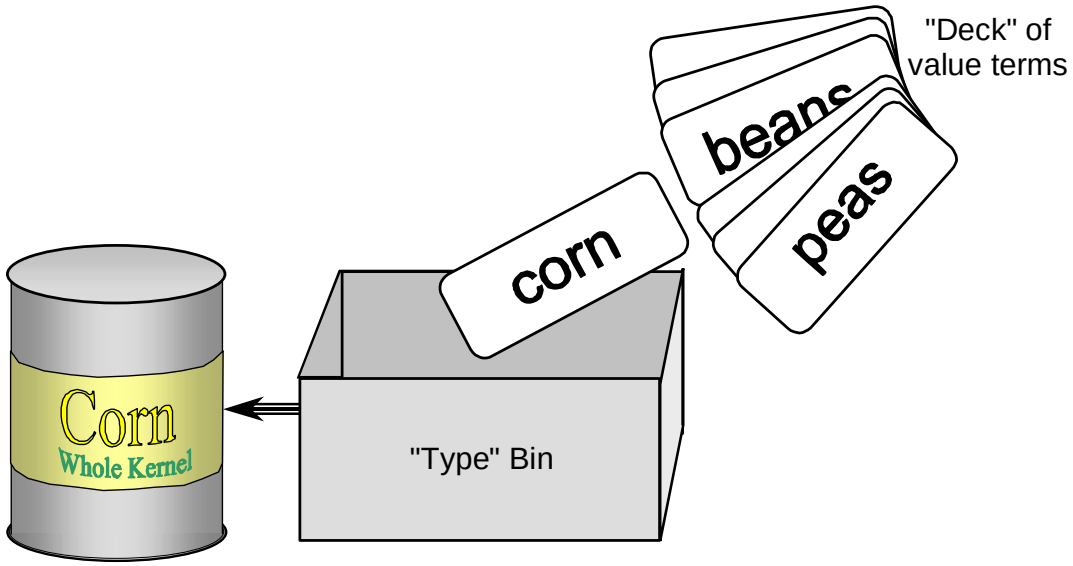

Figure 6. Metadata bin and values describing an object.

When trying to find items of interest, you want a computer to search through particular bins looking for particular cards, for example, the term "corn" in a "type" bin. Search systems can issue queries specifying fields (or bins) and the values (cards) of interest in each of those fields.

Often the values for a field are drawn from a taxonomy—a controlled system of values. A taxonomy may be a simple list of terms or terms arranged in a structure. The most common type of taxonomy is a hierarchy with terms of greater specificity under more general terms. For example:

food stuffs
packaged food
canned
vegetable
corn
whole kernel

These values are drawn from six levels in a taxonomy from the top level of the general class of items you might find in a store down to the specifics of the state of the vegetable. Other metadata might include weight—in English or metric units—nutritional content, calories, producer, country of origin, expiration date and so forth. This example illustrates tagging things with metadata using standard taxonomies or value sets, ranges or formats. It is also possible to tag business relationships and actions; in fact the limitations are those of imagination. Tags, if they are reliable and valuable, can be used not only to discover things,

but also to discover relationships for knowledge management or data mining. The description above is a bit pedantic; my intent is to provide a basic summary the concepts of metadata and taxonomies in a manner that might be useful in talking with subject matter experts about metadata.

> A core principle of metadata is that it is a **communications system**.
>
> It is method for the tagger (the person who makes the labels) to communicate to the user sufficient information for the user to find a resource and use it properly. Both the tagger and the user must speak the same language[3].

The tagger attempts to communicate to the searcher (Figure 7). The searcher's response to the tagger is less direct, being a function of usage. If the user doesn't find or use the resource the metadata describes, an astute tagger will recognize that either the resource is of little use or the tags are not adequate. This analysis is needed for healthy two-way communication. The tagging community must "listen actively."

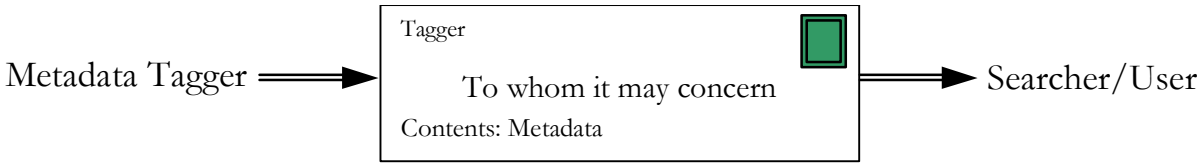

Figure 7. Metadata is communication.

This paper will deal with the tagging of objects rather than with the tagging of processes or people, although some of the same principles probably apply. If I write too much more text on metadata structures and taxonomies your nose will be broken as your head hits your keyboard. If you want to more about metadata structures and taxonomies, take a look at the Dr. Tom Guides (www.twason.com/DrTomGuides.htm). Wear a nose guard.

In reality, the line between metadata and taxonomies is blurry. The people using a metadata system generally try to define a structure (metadata) with well-defined bins into which values (that may be from taxonomies) are placed. The term "field" is often used for the "bin." More technically, these are facets, as each field describes some characteristic of the resource. These fields define the structure of the information. These categories and terms have meaning to a grocer stocking shelves—in this case canned vegetables. Metadata describes an information architecture. A controlled taxonomy is one type of authority list or file, a system of names of

[3] Accurate translations are allowed.

allowed values to be used in that structure. The definitions of terms may be contained in a glossary. These distinctions are abstract; the domain of taxonomists. Keeping these concepts separate can become confusing, missing the point that we are putting labels on objects to find and use them. Let us keep it simple for this document:

- **Metadata** is a set of fields that will have values that can be selected from the taxonomies, if applicable, or implemented with defined formats (e.g., dates).
- **Fields** may have interrelationships.
- A **taxonomy** is a value set that can be applied to a field.
- A **thesaurus** maps among terms either through equivalence, opposition or proximity.
- An **index** is a set of terms that point to specific locations within a resource. We don't consider indexes here.

See also the Glossary at: www.twason.com/Glossary.htm.

When we talk about the metadata on an object, we are talking about the conceptual structure and the actual or potential values within it. In this context metadata is intended to mean the tagging on an item such as the metadata on a picture of a can or an activity such as a purchasing cycle.

I shall use the term "taxonomies" in the plural, although the case can be made that for an enterprise, there is only one taxonomy. The plural is used to reflect that fact that metadata is the structure of information reflecting the facets, and taxonomic entries are the value set from which terms are drawn to populate the metadata facets. It is a useful, albeit abstractly artificial, distinction. Let us not argue about it, but move forward to see what this difference may yield. The objective of metadata is to communicate a description of a resource with categories and terms that have meaning to the searcher. The development of a metadata system of fields (used as general term) and taxonomies should reflect the users' needs, not the tagger's assumptions, no matter how well intended they might be.

The Law Librarians' Discovery. Librarians from the Harvard and Stanford law schools asked for my advice in developing common metadata and taxonomies so that resources could be shared. After some discussion, I asked whom the resources were for. There were three potential users types: students, professors and alumni—e.g., practicing lawyers. The latter were also potential sources of donations, so important for double reasons. I suggested that the metadata system should be derived from the needs of these users; therefore a working group of the types should be assembled. This was done. The results were quite different from the categorization that the librarians had originally planned. The librarians were intelligent people, but they were focused on the manner in which they might "shelve"

the resources; the users focused on how they might use the resources. The designation of "user" is not without meaning.

It is clear that it is essential to have a metadata system that meets the needs of the users. It is also clear that a method for extracting this information from the users would be valuable. A process of developing metadata and taxonomies is realized in a useful method based on practice in directed team meetings. The results of the method are tested and improved in a quality assurance cycle. In the process, the valuable new knowledge can be revealed.

