Tribal Knowledge: Engineers Share Knowledge

Presented to

KMWorld and Intranets Conference 2006

San Jose, CA

Dr. Paulette DeGard
Knowledge Strategist
Flight Deck Engineering/Boeing Commercial Airplanes
The Boeing Company
P O Box 3707
MS 07-21
Seattle, WA 98124
425-717-9238 (voice)
Paulette.h.degard@boeing.com

There is a tendency among many companies, government agencies and non-profit organizations conducting process capturing and knowledge management activities to believe "we're special, we're different, we don't do it that way!" Typically, in working with these organizations, I try to disprove this by showing that 80 percent of what they do is repeatable and only 20 percent of their work is "special." I have even said that about my current employer, The Boeing Company, on more than one occasion. However, I have discovered one thing about Boeing that is different – our product life cycle is 40-plus years. There are very few products that have that kind of shelf life and service requirements. At Boeing we call the 777 our new airplane – its first flight was almost 15 years ago. I still think of the 757 and 767 as new airplanes and we have already discontinued production of the 757 aircraft. We have, however, not stopped servicing because the fleet will be flying for many years to come.

Because of this long-term relationship with our customers, we are forced into a longer view of time than most companies. Our knowledge needs to have a lifecycle that exceeds that of our airliners. Just yesterday, Boeing received a technical service query for an A-20 – that's an airplane that served in World War II. Interestingly, the knowledge requested was available and provided to the customer.

With the advent of technology, Boeing has had to re-address how we manage knowledge and capture processes. The speed of business and the changing demographic of the technical workforce (from long-term "Masters" to nearer-term "Padawans") requires quicker, more efficient and accurate access to knowledge. Suddenly the issue of making tribal knowledge explicit becomes grossly obvious and quite difficult due to the long time that has elapsed since much of the knowledge was generated.

In 1999, Gartner Group did a study on the use of knowledge workers' time. They concluded that knowledge workers spend 60 percent of their time looking for information to do their work.[1] When you couple that with the fact that 70 percent of the working population will be eligible for retirement within the next 4 years[2] a huge issue is posed for large corporations. So how does a

[1] McCampbell, Atefeh Sadri, Linda Moorhead Clare, Scott Howard Gitters, "Knowledge management: the new challenge for the 21st century"; Journal of Knowledge Management; Sep 1999, Vol. 3, Iss. 3, Page: 172 - 179

[2] Holtz-EakinBaby, Douglas; Boomers' Retirement Prospects: An Overview; November 2003

company like Boeing address a situation of this magnitude? We eat the elephant one bite at a time!

Over a year ago, the Flight Deck Engineering organization determined that they needed to make a concerted effort to address the changing demographic of the group; an infusion of younger technical workers and the “bluebeards” retiring. This requires the organization to be more disciplined about how technical and project knowledge is captured and managed. After a number of tries to achieve this goal, it was recognized that engineers, whose job is to analyze things to the very detailed level, were not really suited to create a knowledge management system. So, I was hired into the company to help the Flight Deck Engineering organization get organized.

As you can imagine, the Flight Deck organization’s role is to solve issues with the human aspects of flying the machine and the systems that pilots touch. So, we are very concerned about communicating with the appropriate stakeholders, getting the right answer, and making certain that the answers are documented. We have an important safety and operational responsibility for the tens of thousands of airplanes that are in service. That’s a very large task and one we are very proud to be responsible for at Boeing.

The issue of capturing, retaining, and retrieving information is critical to the long-term success of our organization. The first mandate I was given to design our system was that a new Boeing engineer ten years from now could go to this system and find the information they needed and know that they know it’s the right information. The second mandate I was given was that this new employee would be able to bet their paycheck on the accuracy of the information they found.

So that’s where we started. In identifying the underlying problem, it became clear that a technology phenomena had led us to this situation. I call it “Individualism Diffusion; the antithesis of knowledge management”. What this means is that when Mr. Gates of Microsoft made as his goal that everyone would have a personal computer on their desk, they diffused out into the workplace. However, most organizations did not make agreements on how information would be stored and shared so each person created their own information organization structure. There is typically no agreed upon hierarchy of folders or agreed upon naming convention for files. There is typically inconsistent use of templates and forms. As a result, groups have a difficult time finding and sharing knowledge and information not only within their groups but across the different groups they work with daily.

The solution that was offered up to Flight Deck was that the group would identify a taxonomy for their folder structure and there would be an agreed upon naming convention for documents. As a group, everyone would have to discipline themselves to use the agreed upon file structure

and naming convention or else it would not work and we'd continue to dumpster dive servers just like today.

Several repositories (stand alone servers) would be created to capture previous design decisions, graphics, working knowledge and presentations. These repositories would have varying levels of confidence. For the previous design decisions, an approval board made up of senior engineers and managers would be used to validate the correctness of the information and therefore be read only access. The graphics and presentations repositories would be medium confidence and would be read and write access allowing individuals to place information into the folder structure as they see appropriate. The working knowledge would also be medium confidence and likely be developed in a database allowing individuals to input small amounts of knowledge and information that took less than 4 hours to create. The information would be a good start for someone looking to address the same issue in the future.

So a pilot project group was identified and it was determined that Flight Crew Operations Integration (FCOI) (a subgroup of the overall Flight Deck Engineering group) would be the test case. When the pilot was launched, we determined that FCOI **must** become a learning organization. So what does that mean? The **focus** of the group had to shift from individual learning/organizing to organizational learning/organizing. We would have to establish a reliable means of sharing knowledge now and on into the future. The group agreed that reliable knowledge management could **reduce** mistakes related to inconsistent design decisions. Finally, we acknowledged that knowledge management would **require** reliable and repeatable processes; therefore, we must capture, deploy, and monitor our processes to ensure accuracy.

The methodology I used to begin the design of this process was to establish a taxonomy for the folder structure. Over the years I have learned that without that agreed upon taxonomy or language, it is difficult to get data and information organized. The first technique I used was to interview some of the engineers in FCOI about how they organized. I quickly learned that they all organized differently. What I did discover was that many of them felt they were organized best in the Microsoft Outlook folder structures. So, I asked every engineer in FCOI to take a screen capture of their Outlook folders and email them to me. I then conducted a review of the terms used in the Outlook folders. I was trying to identify commonalities in the group's organization of data. What I discovered from my analysis was that they in fact did organize in a similar way at an aggregate level.

In the aviation industry there is a standard used for categorizing the many various parts. This standard is called the Airline Transport Association (ATA) chapters. The ATA chapters work particularly well for hardware and systems on the airplane. At Boeing, we have numerous products and derivatives of those products. So the majority of the engineers organized either by ATA Chapters or Boeing models. This was a bit of a revelation for the group since most people had never even discussed how they organize individually.

The next step was to determine whether the model or the ATA chapter should be the top level of the taxonomy. I discovered that more people put model first, followed by chapter or function. So, I developed a mock-up folder structure that would be used for the repository we were establishing to store our past decision documents and the set about gaining approval by the technical people.

The next issue was to determine how to assure accuracy of what would be retained in the repository. At Boeing, like many large corporations, people believe what they read because they believe the author was qualified to write the document. This personalization strategy works as long as you know the person who authored the document or if you know you have the original document. To be able to “bet your paycheck” on the validity of the information meant we had to move from a personalization strategy to a process that validated the accuracy of the documents being saved. In order to address this issue, it was determined that an approval board of some kind would be needed before documents were migrated into the repository. This task was a bit daunting for the volume of documents to be migrated into the repository; however, it would assure the validity of the information and the placement within the folder structure.

Another issue that arose around the approval process was whether every single document or issue resolved needed to go through these types of rigors? It was determined that in fact no, not all documents needed that kind of scrutiny, just those that had design decisions and would have major impact on the future. So, we established a confidence level for each of the repositories. The design decision repository would have an approval board and therefore have high confidence in the accuracy of the information. The “working knowledge” repository would be a place where information would be stored that was of value, but not necessarily validated for accuracy.

The third issue that was identified was how to assure that these repositories would be able to stand the test of time. As like many major corporations, Boeing experiences long-term business cycles in production and workforce. So how do you assure that when the downturn happens that the repositories would in fact survive budget cuts? The solution we determined was that we would make certain that the repositories had an economic value that would validate the need for maintenance. When documents are identified for placement into the repository, then a dollar value is associated with them. We chose to use a formula of 8 hours per page times $130 an hour burdened rate for engineering. These values are then aggregated up and a total asset value is established. The spreadsheet for logging the documents to be migrated into the repository is presented to the approval board each month. They then identify if the document is the appropriate one to be kept, agree on the value of the document and identify where the document is to be stored. After the approval process is completed, the administrator will then migrate the documents to the identified places and alert the group that new documents have been placed into the repository.

In testing the approval process, we encountered several issues. First, many of the documents could be placed in numerous folders because there was more than just a single issue documented. Two, the engineers had a difficult time agreeing on where the document should be placed. Three, we did not want multiple copies of the same document in the repository (single source of data only). To resolve these issues we went back to the drawing board on how we accomplish approvals. First we determined that we would work with an individual subject matter expert on placement. Second we would not allow debate of the placement decision unless there was a compelling reason to do so. Last, we agreed to hyperlink within the repository to the original document if multiple locations were required.

The next big issue was one of how to find documents out in the larger world of servers, emails and filing cabinets. As we talked with people in the group about what they wanted to see migrated into the repositories it became quite evident that a lot of the information from the past ten years was resident in emails – both conversations (threads) and documents attached to emails. I initiated an Internet search for tools that search and organize emails. The result of that search was the discovery of the Email Intelligence Platform™ by Clearwell Systems. This system allows you to feed Outlook files into the system, which then quickly crawls and indexes the emails. The benefits of the system are to be able to rapidly capture, and accurately categorize Flight Deck knowledge found in emails and their attachments, provide simple and intuitive access to Flight Deck Knowledge for existing and new employees, accomplish KM capture in a cost effective and efficient manner, and establish repeatable process for future KM capture needs. Bottom line – Clearwell discovers, organizes, and analyzes the information captured in email.

Our method for using the Clearwell System was first to copy Outlook folders from individuals' computers and feed them into the system. Once we had people's PST files captured then we went about having them tell us issues that they wanted researched to see if we could actually find the information. One test we did was with a new employee who had little technical knowledge was to pose a highly technical question to see if he could find the answer in the captured email. In less than 45 minutes, he was able to identify the information and answer the question. Additionally, he located documents valued at over $500,000, which were actual test flight documents validating the answer. These documents, although they could probably be found somewhere else, were imbedded in email and we are now able to migrate them into the repository for future use.

What we are finding is that much of the previous formal decision-making process – use of formal decision making documents and processes – has become more informal through the use of email. In the past, there were meetings and formal events that constituted a design decision or some other important milestone. Today, many of these decisions are agreed upon in emails and not formally documented in the same way. As like many other companies, we have moved from formal structured to informal unstructured over time. This shift in decision-making has been good

because it speeds up the process; however it is more difficult to locate specific decisions and ensure appropriate stakeholders are included unless you are actually involved in the email chain. The Clearwell System is allowing us to capture and store these informal unstructured decisions, which will allow employees today and in the future to have quick, accurate access to vital information for accomplishing their jobs.

Other benefits we have found in organizing our email through the Clearwell System include the ability to use our taxonomy in organizing email and therefore using a consistent classification technique. Probably the best feature we have found is that when the system indexes the emails it removes all duplication of content. This is extremely valuable for being able to read through threaded conversations and reduces the amount of “noise” you typically have to sort through when reviewing an email thread. Two other features we have found helpful are the relevance ranking and hit highlighting. Both of these features make it easier to identify if an email or thread of emails are relevant to the topic being researched. This also speeds up the process for capturing and storing relevant information.

Boeing is a very large company and has major concerns about privacy, retention of information, and security of information. Over the course of the pilot we have had to look closely at things like International Trade and Arms Regulations (ITAR), which Boeing is significantly impacted by because of having so many global partners and employees. Additionally, there are issues of proprietary information that must be protected and restricted to authorized viewers only. These issues have to be resolved if we want to move from an offline pilot to a more accessible system. To date, we have identified issues we need to address from varied groups such as: Information Protection, IT, Legal, Procurement, Records and Information Management, and Email IT architects. These issues have to be resolved in order to move forward and replicate the toolkit and systems we are putting into place for Flight Crew Operations Integration.

There are several major company-wide initiatives happening at this same time. In addition to capturing, storing and retrieving knowledge and information we are also being asked to adopt Lean Engineering as a corporation. This is a huge effort and one that impacts everyone in the company. For us in Flight Deck this means that we need to look closely at our processes and deliverables and identify ways of eliminating waste. In the engineering world, waste is defined as re-work or queue time waiting for someone to make a decision or input. Again, Flight Crew Operations Integration is our pilot group for capturing and re-working processes. Interestingly many people do not think they own any processes nor that there are processes in place. We spend some time educating people on what a process is and why we need to capture them. In the future, new engineers will need to know how to repeat processes to assure that waste is reduced but also that they are following best practices. When we have completed our pilot project, we will have captured in both Visio and in work instructions, using info mapping, over 25 processes that impact not just FCOI but the entire Flight Deck Engineering community. I believe

those 25 processes will be integrated into 5 major processes – issue identification, requirements identification, requirements testing, approval and acceptance, and deployment. Once these 5 major processes are refined and documented, they will become the basis for other sub-processes.

Finally, what are the benefits to Flight Deck for their investment in capturing and managing data and knowledge? First and foremost, the need to assure that organizational memory, in the form of the different repositories, is stored and maintained for future use. Next, knowledge management represents an effort to avoid mistakes and is an insurance policy against the loss of that organizational memory in the future. What it requires is a new mentality – use/get data (current state) versus learn/contribute to knowledge (future state). Third, by addressing legal, security and privacy concerns and working with different stakeholders within the company to assure compliance with each set of requirements, we are assuring that Flight Deck is in line with corporate mandates.

In order to realize these benefits we are instituting some learning organization metrics. The underlying theme for identifying these metrics is to provide processes for seamless and timely transfer of and access to pertinent information. These actions will be realized by reduction of errors in issue resolution, defined commonality of folder structures and adherence to that commonality, increase in reuse and efficiency of knowledge, increase in accuracy of information captured, and usability and retrieval of information.

The planned date for conclusion of our pilot is December 1, 2006. Our goal is to have the tools/toolkit identified and tested then move into implementation of continuously capturing and storing knowledge and information. This will be an ongoing process for quite a while. We will then test the tools and toolkit on the other groups within Flight Deck – moving from a subset of about 25 people to the entire group of 230 people. This leap will help us understand if our tools are robust enough to capture, store and retrieve information as well as capture, document and deploy processes. When we have completed our work with Flight Deck, I believe we will have a replicable system for knowledge and process management that can be deployed in other segments of the corporation.