Strategies for Information Retrieval and Virtual Teaming to Mitigate Risk on NASA's Missions

Daria Topousis
Jet Propulsion Laboratory
1942 Sycamore Court
McKinleyville, CA 95519, USA
1-707-496-8555
daria.topousis@jpl.nasa.gov

Gregory Williams
Jet Propulsion Laboratory
4800 Oak Grove Drive, M/S 303-300
Pasadena, CA 91109, USA
1-818-393-0748
gregory.williams@jpl.nasa.gov

Keri Murphy
Jet Propulsion Laboratory
4800 Oak Grove Drive, M/S 600-140
Pasadena, CA 91109, USA
1-818-393-1674
keri.murphy@jpl.nasa.gov

ABSTRACT

Following the loss of NASA's Space Shuttle Columbia in 2003, it was determined that problems in the agency's organization created an environment that led to the accident. One component of the proposed solution resulted in the formation of the NASA Engineering Network (NEN), a suite of information retrieval and knowledge sharing tools. This paper describes the implementation of this set of search, portal, content management, and semantic technologies, including a unique metasearch capability for data from distributed engineering resources. NEN's communities of practice are formed along engineering disciplines where users leverage their knowledge and best practices to collaborate and take informal learning back to their personal jobs and embed it into the procedures of the agency. These results offer insight into using traditional engineering disciplines for virtual teaming and problem solving.

General Terms

Management, Performance, Design

Keywords

Knowledge management, portal, search, communities of practice

1. INTRODUCTION

Following the loss of NASA's Shuttle Columbia in 2003, a board was convened to identify underlying causes of the accident. After extensive review, the board determined that "NASA's organizational culture and structure had as much to do with this accident as the External Tank foam." [2] This uncovered an endemic problem across the Agency: that the full body of NASA's existing knowledge and resources are not sufficiently accessible or utilized to solve engineering problems. [Strategic Plan] The Columbia Accident Investigation Board also discovered that while most NASA Centers capture lessons learned, they tend to keep knowledge of problems contained within their Center [2]. In the end, the board determined that "NASA has not demonstrated the characteristics of a learning organization." [2]

In addition, big IT investments in the 1990s led to fragmented information spaces that inhibited engineers from finding and sharing the accumulated knowledge of the Agency. Engineers "have had to guess where a particular piece of information might reside in order to query the system and retrieve it." [1]

One component of the multi-faceted solution resulted in the NASA Engineering Network (NEN), a suite of retrieval and knowledge sharing tools specifically aimed to facilitate information retrieval and sharing among engineers at all of the NASA Centers and affiliated contractors. NEN includes a metasearch capability to find data among distributed engineering resources, communities of practice formed along engineering disciplines, and a portal to integrate these components.

2. NEN IMPLEMENTATION

2.1 Metasearch

The NEN search utilizes capabilities already in place at NASA, and integrates, without changing, underlying data resources. A Verity search engine was implemented in 2003, and by 2006 the search had more than 2,000,000 searches per day. Studies of search trends conducted between 2003 and 2006 also showed that the average NASA user uses 2-3 words in a search engine to find data and 75% will only look at the first page of results. [5]

The metasearch uses a service oriented architecture. When a repository has been identified for inclusion in the NEN search, a service level agreement is invoked that determines whether NEN will use a GET command via a web interface for indexing, an SQL query into JDBC or ODBC compliant database, or exporting in XML using a tab-delimited file.

To the extent possible, systems across NASA are encouraged to use metadata and NASA's taxonomy [link to taxonomy]. Specific metadata fields in use, based on the Dublin Core standard, are author, date, description, competency, subject, organization, content type, mission/project, and instrument.
The system currently searches over a dozen, in addition to the Lessons Learned Information System (LLIS), which contains over 1600 lessons dating back to 1972, the Problem Reporting and Corrective Action (PRACA), NASA's Online Directives Information System (NODIS), the NASA Image Exchange, and self-contained lessons learned systems such as those from the Clementine Mission and Orbital Space Plane. The system also searches external resources such as reports from Best

Manufacturing Practices Center of Excellence and U.S. Department of Energy resources on the environment, health, and safety.

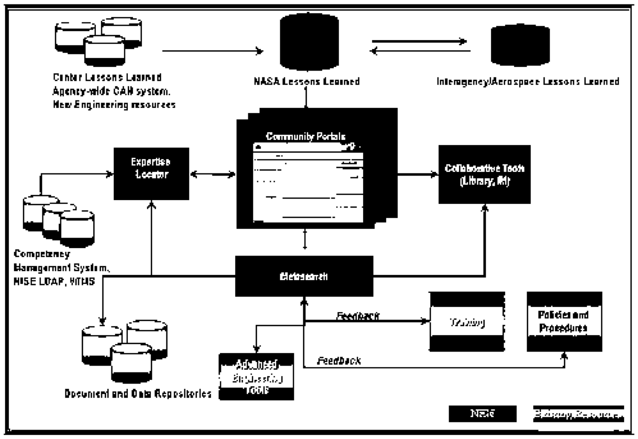

Figure 1: NEN Functional Architecture

The Lessons Learned Information System was moved into NEN's content management system to make updating and accessing files easier. The system is overseen by the Lessons Learned Steering Committee, which includes members from 10 NASA Centers.

2.2 Communities of Practice

According to Shenhar's study of project management, an "organization should learn to learn from its previous project experiences. Sharing information across projects and summarizing project lessons should become a common norm." [9] One way to enable such learning is through a community of practice.

Within NEN, communities of practice are defined as "groups that form to share what they know, and to learn from one another regarding some aspects of their work." [10] Communities of Practice connect individuals with expert peers and promote collaboration, information exchange, and the sharing of best practices across boundaries of time, distance, and organizational silos. NEN rolls out its communities of practice with the goal expressed by Hoadley and Kilner that purposeful conversation might occur around content in context. [4]

All communities are implemented using Vignette Portal software, version 7.2. This tool was selected because it adhered to IT security policies within NASA and provided a range of tools useful to engineers, including discussion boards, calendars, resources links, and web connectors into electronic library collections.

On NEN, communities are formed according to engineering disciplines identified by the NASA Engineering Safety Center (NESC). At the core of the NESC is an established knowledge base of technical specialists pulled from the ten NASA Centers and from a group of partner and organizations external to the Agency. This ready group of engineering experts is organized into 15 disciplines areas. Similar to Orr's study of photocopier repair technicians, where the "construction of their identity…occurs both in doing the work and in their stories" [8], engineers at NASA are accustomed to identifying themselves by their discipline.

Communities of Practice in NEN have a technical leader, who is aligned with the NESC Technical Fellows program; a facilitator, who is part of the NEN team; and a core group of members identified by the leader. The NESC Technical Fellows are chosen as leaders because they have already been acknowledged by NASA to be leaders in their respective field, thus eliminating the need for the NEN team to evaluate and select leaders. Engineer trust in the appointed leader is key to the success of the communities; as Nahapiet & Ghoshal wrote, "where relationships are high in trust, people are more willing to engage in…cooperative interaction." [7]. The leader's role is to align the community with strategic and operational goals, energize the community, and organize meetings and events. [3]

Each community is given one or two facilitators from the NEN team, who provide expertise on layout and communication approaches and who conduct behind-the-scenes maintenance on Vignette and any associated electronic libraries.

The NEN communities are being rolled out in three phases. The first phase included Software Engineering, Nondestructive Evaluation, Systems Engineering, Structures, the NASA Engineering Management Board, and Program/Project Management. The second phase of communities, which are currently being developed and prepared for launch include Materials & Processes, Aerodynamics, Guidance & Navigation, Thermal Engineering, Propulsion, and Electrical Power. Following the successful rollout of this second phase, 11 more communities will be implemented.

The approach to community rollout is shown in Figure 2. Each community leader is presented with a set composition and layout of portlets, which he or she may then modify based on his or her discipline's needs. Once the community is designed according to the leader's approval, a core group of community members are identified and given review of the site. As community interest widens and builds, the leader determines when the community should be made live. Once a community is live, it is by default available to anyone with access to NASA's intranet. This includes NASA personnel and badged contractors.

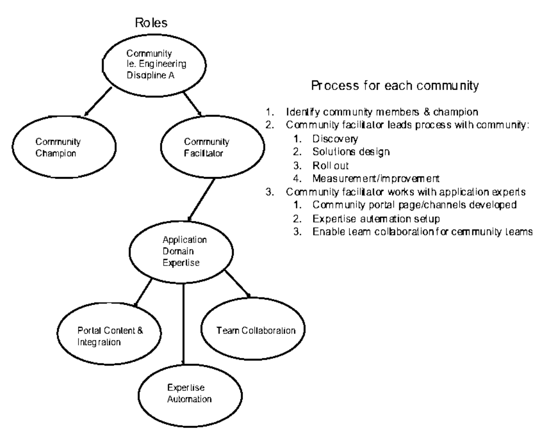

Figure 2: Community of Practice Rollout

2.3 Portal

NASA has an agency-wide portal called InsideNASA, and some Centers have their own internal portals, but these all cater to all personnel regardless of their role within the organization. In order to bring together the metasearch, communities of practice, and general resources of interest to engineers, the NASA Engineering Network also includes a NEN portal. From the site, also built using Vignette, engineers can access live communities of practice, see relevant announcements, learn about training opportunities, and access popular resources such as specific NASA requirements and external resources such as professional associations and societies. The portal also includes a section dedicated to Lessons Learned, where engineers can browse by several categories such as year and topic, and can upload new Lessons Learned.

3. METHODOLOGY

3.1 Search

In February and March of 2007, a study was conducted to determine search engine research methodologies and derive a baseline performance. The study included surveys and direct observations of how they interacted with NEN.

The surveys were completed by 29 engineers at 9 NASA Centers using Survey Monkey. Questions focused on the participant's current job, most common searches, most frequently accessed documents, and awareness of NEN.

One week following the survey, remote observations were arranged with the five individuals who expressed an interest in participating further in the study. Observations were conducted using WebEx with telephone support for the audio component. Participants were asked to conduct a known-item retrieval and an unknown-item retrieval. Data gathered during this process included query modification count, click count, average ranked position of relevant documents, and satisfaction score.

3.2 Communities of Practice

As of May 2007, seven communities have rolled out with several more under development. Only those communities rolled out were analyzed. Activity in the communities was measured, and input from community leaders and facilitators was gathered to create success criteria. As a measure of success, communities are expected to roll out within 6 months of initiation and course correction and establishment should happen 6 months after completion. In the second year and beyond, the facilitator and leader would be involved in ongoing maintenance to keep the community active and engaged.

System logs were used to generate metrics for the communities. These logs analyzed how many hits the specific communities received in general as well as popularity of individual portlets within the community.

3.3 Portal

The NEN portal is currently being redesigned and integrated with the InsideNASA portal to have a common look and feel and to increase visibility for NEN. Once the portal has been in place for several months and rollout campaign communication initiatives are complete, the NEN portal will be analyzed using surveys and web statistics.

4. FINDINGS

4.1 Search

The first finding from the study was that users who don't use NEN either use a different search engine (46%) or were unfamiliar with NEN (38%). Figure 3 shows the most common reasons for not using the site. A communication plan with clear target audiences and messages is being created to resolve this issue. The rollout activities mentioned in 3.3, above, will also improve NEN familiarity by users across NASA Centers. In addition, a simple search is being added to the InsideNASA interface which would improve the NEN search visibility.

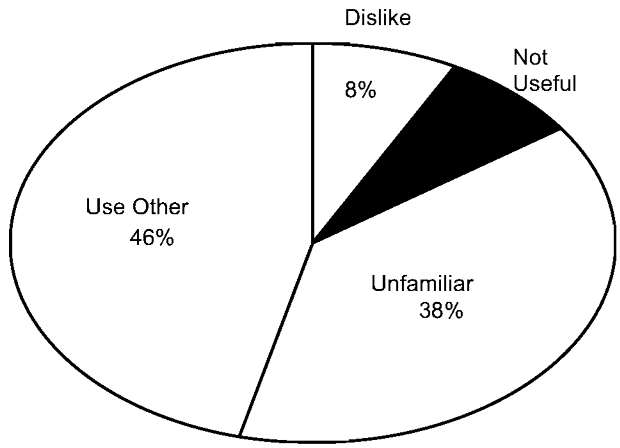

Figure 3: Reasons for not using NEN [6]

During the observation phase, it was found that users were most satisfied with the query modifications mode part of the observation. Query modification mode measured how many changes a user made to their initial search query in order to find what they were looking for. Of all the metrics gathered during this phase of the study, overall satisfaction was quite low. However, during the study it was observed that formatting issues resulted in participants thinking the search did not return the desired results when, in fact, it had. This issue can be easily resolved with improved formatting. For example, if query results showed where the search term was found within any given document, users might be more satisfied.

According to Paul King, who conducted the study, "The ARP [Average Ranked Position] for known-item searches was very high, which indicates poor performance. A mean ARP of .85 indicates that the average relevant document was found on the bottom of the set of search results (between the 8th and 9th positions.) This may be addressed by changing the weights of the indexed fields of cataloged items in the Verity search engine." [6]

See Table 1 for observation results. Note that known-item mean count is skewed upward due to a single exercise in which the participant clicked on many items.

Table 1: NEN Search Observation Results [6]

Metrics	Known-Item			Unknown-Item		
	Mean	Median	Mode	Mean	Median	Mode
Query Modifications	0.8	1	1	0.6	1	1
Click Count	4.4	3	2	1.89	2	2
Relevant Hits	N/A	3	3	2.89	2	1
ARP	0.85	0.8	N/A (nulls)	0.31	0.3	N/A (nulls)
Relevancy Score (RS)	N/A	N/A	N/A (nulls)	2	2	N/A (nulls)
Satisfaction Score	4.6	5	5	3.6	4	4

Satisfaction Scale
1. Very Satisfied – *would use it regularly*
2. Somewhat Satisfied – *would use it periodically*
3. Adequate – *may use it sometimes*
4. Somewhat Unsatisfied - *probably would not use it*
5. Not Satisfied – *would never use it*

Relevancy Scale
1. Relevant
2. Somewhat Relevant
3. Irrelevant

Following the study, it was decided that regular log metrics should be automatically gathered on a monthly or quarterly basis to enable an ongoing analysis of search efficiency. IP addresses of NEN team members would be filtered out of this analysis. In addition, the noise of full-text searching could be reduced by adjusting weights.

4.2 Communities of Practice

In combining the responsibilities for each of the facilitators, it is recognized that in practice, and from the human factors standpoint, most individuals have unique skill sets that typically favor either a business background or a technical background. It is our experience that each community of practice requires expertise for both, so the NEN team selects and aligns each facilitator with a co-facilitator based on the combined skill sets and strengths of each. For example, if the facilitator's skill set has a stronger background in business and communications, then he/she is aligned with a co-facilitator that has a greater strength in technical skills. The reverse is the case when the facilitator has a stronger technical background, whereas he/she is aligned with a co-facilitator with a strong business and communications background.

Through the concept of aligning the facilitator and co-facilitator based on a "synergy" and pairing of skill-sets, this creates a natural synergy between the two positions, one supporting the other. This also provides an opportunity for each to learn from and grow their individual abilities in the areas where they are weaker and less experienced. A comprehensive list of skills and competencies was created and once those competencies were identified, a survey instrument was created using a Likert Scale to identify and quantify the existing strengths and weaknesses of each facilitator for each competency. Two different survey assessments were done, one to determine technical competencies, and one to determine the business and communications competencies. The results of the two surveys were used to conduct a gap analysis to determine facilitator skill levels at both the team level as well as at the individual level for both the technical and business competencies. Training is then targeted on both an individual level as well as a Team level, starting with those areas identified as most crucial and where they are lacking the most.

Once a community rolls out, announcements are placed on the InsideNASA portal and on the NEN portal. The metrics shown in Table 2 are from April 2007 and indicate hits to a specific page; i.e., how many times a page or community has been accessed at the top level.

Table 2: Page Level Metrics

Page Name	No. of Hits
InsideNASA Home	69,229
Knowledge Management Community of Practice	64,379
Structures Community of Practice	47,368
Nondestructive Evaluation Community of Practice	18,571
Systems Engineering Community of Practice	14,784
Software Engineering Community of Practice	5,709

The team then used April 2007 metrics to analyze which portlets within the active communities were most accessed. According to these metrics, portlets listing popular documents from an electronic library were by far the most popular portlets across the communities. The key documents portlet in the Structures community had 32,106 hits and the Systems Engineering community had 11,826 hits.

While one month provides a snapshot of metrics, the NEN team plans on analyzing monthly metrics to see how hits to pages and communities change over time and whether or not these changes are precipitated by parallel community activities such as conferences, meetings, etc.

Personal experiences among facilitators have shown that spending time upfront with a community leader to plan the community is essential in a successful rollout. This is in line with the findings of Greenes [3]. In addition, communities with a clear charter that can be expressed to potential members helps clarify and streamline the message.

When rolling out a community, attending conferences has proved to be one useful means of garnering interest from potential members.

5. ACKNOWLEDGMENTS

Our thanks to Jeanne Holm for providing expert input on this paper, and to Jayne Dutra, Lisa Smith, and Paul King for conducting search analyses for the NEN Team. NEN is sponsored by the NASA Office of the Chief Engineer.

6. REFERENCES

[1] Dutra, Jayne, and Smith, Lisa. Strategy and Implementation Plan: Providing High Precision Search to NASA Employees Using the NASA Engineering Network. JPL internal document, January 2006.

[2] Gehman, Harold W. (Chairman). Columbia Accident Investigation Board Report, Volume 1. NASA and GAO, August 2003.

[3] Greenes, Kent. Communities of Practice Facilitation. Presentation to the NASA Engineering Team, 2006.

[4] Hoadley, Christopher M., and Kilner, Peter G. Using Technology to Transform Communities of Practice into Knowledge-Building Communities. SIGGROUP Bulletin, 25(1), 31-40.

[5] Holm, Jeanne, and Yew, Manson. Overview on Engineering Search at NASA. NASA internal document, June 2006.

[6] King, Paul. NEN Search Study: Report Out for the NASA Engineering Network. JPL internal document, March 2007.

[7] Nahapiet, J & Ghoshal, S, 'Social Capital, Intellectual Capital and the Organizational Advantage', Academy of Management Review, 23(2), (1998).

[8] Orr, Julian E. Sharing Knowledge, Celebrating Identity: Community memory in a service culture. in Middleton, D. and Edwards, D. eds. Collective remembering, Sage Publications, Newbury Park, CA, 1990, 169-189.

[9] Shenhar, Aaron J. Strategic Project Management: The New Framework. Proceedings of PICMET, Portland International Conference on the Management of Engineering and Technology, 1999.

[10] Wenger, E., and W. Snyder. 2000. "Communities of Practice: The Organizational Frontier." Harvard Business Review. 78(1): 139–145.

End of File