The Role of the Community in a Technical Support Community: A Case Study

Donald M. Allen[1] and Thomas Schneider[2]

[1] Cisco Systems, Inc., Technical Services, Colorado Springs, USA
donallen@cisco.com
[2] Cisco Systems, Inc., Technical Services, Zurich, Switzerland
thschnei@cisco.com

Abstract. Resource tagging has become an integral and important feature in enabling community users to easily access relevant content in a timely manner. Various methods have been proposed and implemented to optimize the identification of and access to tags used to characterize resources across different types of social web-based communities. While these user-focused tagging methods have shown promise in their limited application, they do not transfer well to internal business applications where the cost, time, tagged content, and user resources needed to implement them is prohibitive. This paper provides a case study of the process, tools, and methods used to engage users in the development and management of a tag taxonomy (folksontology) used to characterize content in an internal technical support community in the Cisco Global Technology Center.

Keywords: component; tag, taxonomy, folksonomy, social community, folksontology, case study.

1 Introduction

Cisco launched an internal technical support community to enhance support and troubleshooting collaboration among 3500 support engineers across eight primary technical centers, as well as numerous other smaller technology groups around the globe. The focus of the internal support community (called TechZone) is to replace the historical standard of practice using internal email mailing lists to post questions and request assistance from experts. The use of the mailing lists relied heavily on the submitter's knowledge of the mailers associated with specific technologies and the willingness of the email list members to respond to those requests. This type of interaction led to poor predictability of when and whether a response would be received. In addition, different support engineers would repost the same requests when they encountered the same problem at a later time.

One goal of the support community is to provide a better structure for posting requests for assistance; to make submission targets more obvious; and to prove a mental

A.A. Ozok and P. Zaphiris (Eds.): OCSC/HCII 2013, LNCS 8029, pp. 335–344, 2013.

model of the technologies and the semantics of the technology hierarchy. Cisco has thousands of products and services that are supported by the Cisco Technical Services group, each with its own vocabulary to describe the product, its features, and the types of problems encountered by customers using it. This issue is similar to large retail and community support web sites where web site structure and navigation is important to guiding users to the correct location and obtain product support. User enabled support communities are common on the internet. Companies and end user communities establish them to facilitate the sharing of configuration and troubleshooting lessons learned among community members. The business impact on a company when a user fails to post to the correct discussion board or does not receive a correct and timely response via this type of free service is minimal since the user is getting what he paid for. When a company, such as Cisco, uses the same user support community to assist support engineers who are in turn assisting paying customers the impact to the company could be considerable. Cisco considered it paramount to address the navigability of the over 290 support communities supported in TechZone in order to ensure proper posting and facilitate timely response. A number of product hierarchies exist within Cisco to support various business activities associated with designing, developing, manufacturing, and selling products. None were deemed applicable since our support organizations were not consistently organized around products. This paper describes how Cisco employed collaborative tools to enable target users of the community support system to construct and manage the navigational structure of the internal support site.

A second goal of the development of the support community is to capture the knowledge developed during discussion posts and convert it into content available to other engineers to avoid future repetitive posts. Various enhancements were made to the community application to improve access to the available content. The enhancements required a quality of content characterization (metadata) which was not available from the content in the system. Tagging is common among content rich web sites (e.g. Flickr, Delicious) to improve resource (e.g. pictures, websites) access. In order to maximize content "findability" and to move the organization to a common support-focused vocabulary, the development team worked with the internal support community using social collaboration tools to develop an augmented folksonomy[1] to improve content access. The logic, tools, and processes that were employed in this effort will be described below. Three efforts within the overall implementation effort are discussed in the paper:

1. Definition of the discussion board navigational model,
2. Identification of the structured vocabulary (folksonomy) used to characterize the discussion posts and knowledge content in support of timely and accurate access.
3. The design and development of a tool to enable members of the community to self-manage the structured vocabulary.

[1] A folksonomy is a system of classification based on the activity of sorting information into categories based on the consensus of the users of the information.

2 Tags, Taxonomies, Folksonomies and Labels

Web sites such as Delicious (formerly del.icio.us) and Flickr and the research associated with their use demonstrates the importance of augmenting user supplied content with tags as a method to improve user access [1-2]. The emergence of folksonomies [3], including user-created and managed tag sets, demonstrates the efficacy of utilizing tag taxonomies to further enhance user content accesses [4] and to provide additional structure to the deployed tag set. The majority of the research around collaborative tag assignment and tag taxonomy focuses on two areas. The first focus is on the assignment of tags involving the collaboration of a group of content reviewers. Their goal was to assign the appropriate tags to a piece of content [1]. The second area of focus is on the application of computational logic to large, pre-tagged content data set with the goal of surfacing tag-relatedness in support of identifying tag navigational hierarchies [5,9] or tag presentations in the user interface (e.g. tag clouds). Both methodologies require specific resource availability. Collaborative tagging requires the availability of users who can review content and assign the appropriate tags, while algorithmic analysis requires the availability of large data sets, in terms of tag content, in order to characterize tag usage.

Note: Cisco utilized *Social Support Online Community* software licensed from Lithium to implement the TechZone support community. The Lithium software provides both a tag and label infrastructure for annotating content (discussions and articles). The tag system is consistent with the community-based unmanaged, informal functionality common to Web 2.0. The label infrastructure utilizes an administrator controlled folksonomy system [1]. In order to avoid confusion between the two terms in this paper and to be consistent with prior research, the term tag is be used in place of labels.

The application of existing research methods to the TechZone support community is problematic. As noted earlier there are 290 separate communities within TechZone: in essence a community of communities. Each community targets a specific technology and contains discussions and articles focused on that technology. The TechZone underlying infrastructure supports a document library of articles developed to address specific product problems and solution. An analysis of the available articles was conducted during the development of the new tag taxonomy. At the time of the analysis, the one year old site contained over 6,000 articles. The number articles per community ranged between 4 and 233, with an average of 21 articles per community. A detailed analysis of the tag utilization revealed that on average only 2.7 tags were assigned per article with a range between 2 and 22. Considering that the system required that at least one tag be assigned to an article prior to publication and that required tag was automatically assigned based on the content type of the article (e.g., How-to, Troubleshooting, etc.), users were assigning fewer than 2 tags to characterize an article. Applying established tag refinement methods to the available data set was considered impractical. Collaborative tagging would require the time and coordination of support engineers across the globe whose job responsibilities preclude such an effort; moreover, there was no management buy-in to the process. The existing data

set is not sufficient within each communities to apply the techniques used in prior research since the since tag sets would need to be community-specific (a tag used in the Security Firewall community would not necessarily be applicable in the TelePresence Call Control community).

A hybrid, targeted folksontology [6] approach was taken to construct a tag taxonomy that identified the common tags across communities as well as those that would be specific to a single or subset of communities. A folksontology combines the distributed, community driven tag set with the disambiguated, hierarchical and formality of a taxonomy. Our approach involved technical leaders from each of the different technologies (e.g. IP Telephony, LAN Switching, etc.) who worked together using collaboration tools to construct the set of communities and associated list of tags targeted to their technology. Technical leaders were paired with a taxonomist to augment that tag set with structure and additional metadata to improve overall utility of the resulting tag set [4]. This approach accommodated the rollout plan of new technology to the different groups (which included training, communications, and content migration) while at the same time afforded the construction of a tag taxonomy that included tags common across technologies (e.g. install, configure, troubleshoot) as well as those specific to the technology (e.g. AAL2 Trunking, ASR5000).

2.1 Tag-Zation with Representation

Kiu and Tsui [7] identified two problems associated with constructing and using taxonomies traditionally constructed by taxonomists and/or domain experts: 1) the disconnection between the terms identified by the "experts" and those employed by the vocabulary users and 2) the difficulties involved and costs associated with keeping the taxonomy up-to-date. Kiu and Tsui's propose the application of computational algorithms and data mining techniques to address these issues. We chose a more user-centric design approach, one that actively engaged representatives of the user communities and gave them ownership of the tag taxonomies (to their taxonomies) in order to address the disconnection and update problems.

In order to provide direct control of the tag taxonomy to members of the community, we designed and implemented an application to enable the technology community representatives to manage the tag set (additional information about the Label Manager can be found in the corresponding section later in this document). Previous research on the use of controlled vocabularies versus social tagging demonstrates some advantages to using controlled vocabularies in terms of content retrieval [8].

In order to address the issues of tag differentiation, group think tendencies, and relevant tag availability, we decided to provide richer information about a tag to enhance the end user's access to and understanding of a tag. In addition, there was an interest in expanding the application of the tag taxonomy into other areas of the business. One area of interest is in the streamlining of the flow between the creation of a customer support case, the content in TechZone available to address the case and the support engineers with the knowledge and expertise to provide assistance. A single, common vocabulary to characterize a support case, content, and support engineer knowledge provides a foundation for automation.

The need to share and exchange the tag taxonomy within Cisco and with Lithium and a desire to base our development on industry standards led us to select the Simple Knowledge Organization System (SKOS). SKOS provides a standard method to represent knowledge organization using the Resource Description Framework (RDF) [11]. The SKOS schema provided the foundation for many of the metadata elements needed to support tagging. The table below identifies the tag metadata used:

Table 1. Tag metadata used

Attribute	Description
ID	System generated unique ID
Preferred Tag	Tag that will be used by the system when presenting the tag to the user
Definition	Text describing the tag – could include an expanded acronym
Scope Note	Text describing the context (technology or sub technology) where the tag is intended to be used
Alternate Tags	Set of tags that are equivalent to the preferred tag such as synonyms and expanded acronyms
Hidden Tag	Set of tags that are equivalent to the preferred tag but are not displayed to users (e.g. project code names)
Category	Specialize grouping of tags based on their similarity, used to make user tag access easier – Activity, Content Type, Environment, Feature, Miscellaneous, Problem, Cisco Product, Third-Party Product, and Protocols & Standards
Communities	Set of support communities where the tag has a strong affinity
Broader/Parent	Set of tags that have a more general definition than the tag (the parent in a parent-child relationship)
Narrower/Child	Set of tags that have a more specific definition than the tag (the child in a parent-child relationship)

Tag metadata is used by the search, filter, and tag assignment components in the UI to improve the overall usability of tags in the following ways:

- Tool tips presented when a user moved the mouse pointer over a tag displayed the tag description and the broader and narrow information if present.
- The preferred, alternate and hidden tags are used during search and when assigning tags to content.
- Tag category and the broader/narrower attributes are used display the tags in a hierarchy component used for filtering and assignment.
- The community attribute is reserved for future use (see the Future Plans section).

3 Support Community Identification

3.1 The Process

A participatory design approach [10] was utilized during the construction of both the community structure and the tag set. A design team was established for each of the major technology areas. Each design team consisted of experts from the technology group, a program manager, a developer, and human-computer interaction expert

(who also played the role of the taxonomist). Each technology design team completed the following process:

1. Held a kickoff meeting in which the overall process was discussed along with deliverables and timelines.
2. Trained technology experts on the social community application which included the different capabilities they were expected to configure (community hierarchy, tags, etc.).
3. Seeded a pad in Etherpad [13] with a template which included instructions and an example structure to follow.
4. Established a schedule with checkpoints for group review.
5. Mapped the community to existing problem code taxonomy to facilitate system-to-system data exchange. This step was purposefully sequenced after the group constructs the community hierarchy in order to prevent the existing structure from influencing how team organizes the community.
6. Identified the set of tags to characterize the content for each of the identified communities (details of this process are provided in the Tags and Taxonomy Process section of this document).

3.2 Lessons Learned

Using Etherpad to collect the community structure proved to be a good method for collaboration among the team members. The ubiquitous access to the pad enabled users across different geographic locations and time zones to collaborate on the building and refinement of the community structure. A signification shortcoming however was the inability to provide comment or the reasoning behind a change or suggested change. Typically these kinds of discussions migrated to email and became unavailable to some of the team members.

4 Tag and Taxonomy Identification

4.1 The Process

The support community structure developed during the initial phases of a deployment was used as the framework onto which the tags were identified and organized. The following process, similar to the one used for the community structure, was used: a kickoff meeting, followed by training, interim checkpoint meetings, and a finalization meeting. The goal of the training was to educate the design group on how the users in the support community would leverage the tags to search for and filter content. The training also provided some general guidelines and heuristics for identifying relevant tags. We wanted to avoid being too prescriptive in terms of the number of tags we were expecting and avoided providing hard and fast rules that needed to be followed. Our desire was to have the design teams identify the tags most used for folks in their organization to characterize their questions (discussion posts) or to find the solutions to their problem (articles).

4.2 Tag Identification Guidance

The following guidance was provided as part of the technologist training:

When to add a new tag

- Tag does not exist or new tag is not an alternate to an existing tag
- Existing tags do not provide the differentiation needed

When to use acronyms versus words

- Acronyms should be used when they are the common method of identification
- Do not create new acronyms, only use established ones

Using single word versus multiple word tags

- A tag should not combine object and action
- Multi-word tags should be composed of words that are unable to stand alone (Active Directory) and identify the intended target
- Avoid using existing tags in new multi-word tags (categories and product names are the exception)

When to modify an existing tag

- Changes to shared tags (used by other dictionaries) should be negotiated where it makes sense

Use the lemma/headword as a tag

- Other forms should be alternate tag (configure – configuring)

Differentiate noun and verb version of a headword

- Example - Install (verb) and Installation (noun)

How to identify product model numbers

- Avoid the use of model numbers without qualifier (e.g., “7921 IP Phone” instead of “7921”)

4.3 Tag Identification

The technology representatives collaborated in the identification of the tags using the Etherpad created to construct the community structure. When the community structure was complete, reviewed and agreed upon, the Etherpad was versioned and all subsequent changes to the finalized tag set were either communicated to the Taxonomist for update or were created using the Tag Manager application (see below for more details). We found that some groups migrated to other tools to identify their tags (such as Microsoft Excel) in order to better track common versus community-specific tags.

4.4 Role of the Taxonomist

During the early phases of the technology rollouts to Tech Zone the Taxonomist played an active role to ensure that the proposed tags followed the guidelines presented during training to preserve the integrity of the tag metadata. There was a propensity demonstrated by some groups to combine product and activities together into a single tag (e.g. "Configure ASR500") and for acronyms to be added without the corresponding expanded version. Additional tag refinement was required in the Activity category in order to differentiate the task being performed versus situations where the task had already been performed (e.g. installing a router versus configuring an installed router).

5 Tag Infrastructure

As noted earlier many social communities have moved toward tag hierarchies primarily driven by computational analysis of tags or user collaboration to create tag taxonomies. Cisco lacked the body of tagged content and the business justification for user collaboration (i.e. the cost associated with engaging support engineers by taking them off the work queue). Our decision was to develop a tag taxonomy infrastructure that could support our tag metadata, hierarchical needs and would integrate into an existing user interface of the deployed support community web application. We worked with Lithium for approximately a year to get the infrastructure in place in Social Support Community (Lithium) and to design and implement an application to manage tags (Cisco - see the next section).

5.1 Tag Management

A tag management application (representative self-managed tag taxonomy) was developed to provide the technology representatives a tool to create, edit, and update the shared tag set used within the support communities. The application included the ability to search for tags as well as to filter the tag set based on community and tag category. In order to address the need for tags that are specific to certain technologies, a user is able to add tags to communities they have permission. To support common tags users are able to associate exiting tags to their communities. The application supports user entry of all of the metadata described above. A number of usability enhancements were added to prevent users from entering duplicate tags either as single or multiple word tags.

The Tag Management application is a web-based application deployed via Apache HTTP server and Apache Tomcat. The internal data representation of the tag taxonomy is based on the Resource Description Framework (RDF) which is based on the W3C standard to describe metadata [11], provide an interoperable taxonomy exchange format, and can easily be merged with other Cisco ontologies. The Sesame [12] infrastructure was used for backend storage of the tags, user access/authorization, and the storage of the community discussion board hierarch.

5.2 End User Request for Tag Changes

A technical community within TechZone is used to engage end users in the tag management process. Users can post new tag requests to the discussion board to facilitate discussion with their peers and to interact with the taxonomist and technology tag representatives. Once a change has been properly vetted the Tag Manager application is used to update the tag set and push the update to TechZone.

6 Future Plans

6.1 Automatic Tagging

Ideally technology could be used to automatically assign the tags to content or to suggest tags to the user based on an algorithmic analysis. We are currently evaluating natural language processing capabilities that might be able to enable this. We suspect having a controlled tag vocabulary will make this processing a bit easier; however, some degree of specialized logic (such as considering context) will be required to ensure the user is not overburden with spuriously relevant tags. We hope to report on this work in the future.

6.2 Existing System Integration

As with any large, established company there are a number of existing systems that might benefit from understanding the mapping of the TechZone tags into their workflow. A couple of efforts currently underway include: 1) Automating the association of tags in a legacy document publishing system used in the external publication of content to Cisco's public site to streamline moving TechZone articles to that site, and 2) Integrating the use of the TechZone tag set into our customer case reporting and management tool to enable the cross communication between the two systems more efficient.

6.3 Tag Lifecycle Management

As other researchers have noted, in order to maintain the freshness and relevancy of a tag set it is important to analyze the use of tags within the system. To efficiently perform this analysis tools are needed that can map historical tag utilization against product lifecycle for both Cisco and Third Party products. Understanding tag usage as it relates to the evolution of a product (initial introduction, updates, and end-of-life) is critical to effectively manage product tags. It will not be enough to look at simple tag utilization to understand the usefulness of a tag; evaluations must be done in context. Combining the data and the product lifecycle will be a challenge.

The role of a tag taxonomist in the tag management process is one we must explore and resolve. The breadth and depth of product, feature, and problem knowledge required to manage such a large content domain is too large for a single person to undertake. Adopting a hybrid model combining the skills of the technology specialist sand taxonomists appears to provide a viable approach; the details of that relationship have yet to be sorted out.

7 Conclusions

This paper described a process and an application developed that addresses the requirement to develop a tag taxonomy in a business environment where neither the content needed to algorithmically identify neither tags nor the business justification for engaging community users in the tag identification process were present. The material presented proposed a method for engaging technology experts in the tag identification process as well as in the maintenance of the tag taxonomy. Monitoring and analysis of the use of tags in TechZone and the tag maintenance activities of the technology experts should provide insight into the efficacy of this approach.

References

1. Robu, V., Halpin, H., Sheppard, H.: Emergence of Consensus and Shared Vocabularies in Collaborative Tagging Systems. ACM Transactions on the Web (TWEB) 3(4), 14:1–14:34 (2009)
2. Heymann, P., Garcia-Molina, H.: Collaborative Creation of Communal Hierarchical Taxonomies in Social Tagging Systems. Stanford Info Lab Technical Report (2006)
3. Hammond, T., Hannay, T., Lund, B., Scott, J.: Social Bookmarking Tools A General Review. D-Lib Magazine 11(4) (2005)
4. Trattner, C., Körner, C., Helic, D.: Enhancing the Navigability of Social Tagging Systems with Tag Taxonomies. In: Proceedings of the 11th International Conference on Knowledge Management and Knowledge Technologies. ACM, New York (2011)
5. Astrain, J.J., Echarte, F., Córdoba, A., Villadangos, J.: A Tag Clustering Method to Deal with Syntactic Variations on Collaborative Social Networks. In: Gaedke, M., Grossniklaus, M., Díaz, O. (eds.) ICWE 2009. LNCS, vol. 5648, pp. 434–441. Springer, Heidelberg (2009)
6. Van Damme, C., Hepp, M., Siorpaes, K.: FolksOntology: An Integrated Approach for Turning Folksonomies into Ontologies. In: Proceedings of the ESWC Workshop "Bridging the Gap Between Semantic Web and Web 2.0" (2007)
7. Kiu, C., Tsui, E.: TaxoFolk: A hybrid taxonomy–folksonomy structure for knowledge classification and navigation. Expert Systems with Applications 38(5), 6049–6058 (2011)
8. Golub, K., Moon, J., Tudhope, D., Jones, C., Matthews, B., Puzon, B., Nielsen, M.: EnTag: Enhancing Social Tagging for Discovery. In: Joint Conference on Digital Libraries (JCDL), pp. 163–172. ACM Press, New York (2009)
9. Simpson, E.: Clustering tags in enterprise and web folksonomies. HP Labs Technical Reports, HPL-2007-190, pp. 222–223. AAAI Press, Palo Alto (2008)
10. Ehn, P.: Scandinavian Design: On participation and Skill. In: Schuler, D., Namioka, A. (eds.) Participatory Design, Principles and Practices, pp. 41–78. Lawrence Erlbaum Associates (1993)
11. SKOS Simple Knowledge Organization System, `http://www.w3.org/2004/02/skos/`
12. openRDF.org, `http://www.openrdf.org/about.jsp`
13. Etherpad, `http://etherpad.org/`