Hubbub - An innovative customer support forum

Duong Nguyen, Simon Thompson, and Cefn Hoile

British Telecommunications Plc
Adastral Park, Martlesham Heath
Ipswich, IP5 3RE, United Kingdom
{duong.nguyen,simon.2.thompson,cefn.hoile}@bt.com

Abstract. Internet user forums have been proven to be effective not just as a community meeting place but also as a supporting tool for various business products. Traditional forums are designed with "browse and read" journey in which users have to select the right sub forum to get into and select topics to read from within. However, we have identified a new trend in forum design toward community question answering systems with an "ask questions first" user journey, a topic-less organisation, search based information retrieval and social network inspired alerting. Here, we report on the implementation and trial of such a forum, Hubbub, that epitomizes the aforementioned trend. It is designed to eliminate key issues found in current forum technologies and has been fielded as a support channel for a BT Softphone product, resulting in a significant reduction in support costs. We then report on the performance of this forum in practice and speculate on the reasons that forum design is taking this direction. Finally, we conclude the paper with some thought about future direction of forums in the Web 2.0 era.

1 Introduction

Modern converged communication products pose considerable support challenges. For example, a VOIP service is dependent on network connectivity, proprietary client software, a VOIP server, PC hardware, a microphone and drivers, a modem and an operating system. Any failure of any part of the solution will leave the product inoperable and the user with a problem.

This is also a problem for the service provider, because users are not able to identify the source of failure for their service. Even articulating the issue in the expected manner is often beyond a non technical user. In addition converged services are often supplied at a very low cost, while they may be profitable there is often insufficient funding from these non traditional services for a large scale traditional support infrastructure. However, in order to make the wider economics of converged service provision work it is critical that these services are provided as part of a complex converged services ecosystem (i.e. the Internet or voice services), which implies all the big company brand risk and customer loyalty dynamic which makes traditional support groups necessary.

Internet forums offer a potential way for companies to handle this support challenge [2]. Community derived support (customers helping each other) offers

zero cost support, one-to-many support (a support agent helping a customer on the forum and many other customers then accessing and reusing the support information) offers a much more efficient support model than single point interactions. Companies working at the top tier of the internet services ecosystem are exposed to a number of risks by internet forum activity.

1. Internet forums can be a vehicle for disruptive customers to damage brand and distort the perception of a product via anecdote and (sometimes) malice. When surveying 25 different forums (selected from the top results of Google searches over forum provision software) we found that on average each post is viewed 119 times (see figure 1, note that the graph is plotted on a log scale). It seems from this observation that the minority of users are posters and any views on the forum may or may not be representative of the experience of the majority. Companies may honestly stand by the view that their customers have a right to publicly state their issues with a product or service, but at the same time it is legitimate to try and ensure that these views are given an appropriate context and appropriate weight [3]. The usage characteristics of internet forums lend themselves to polemic and propaganda and are therefore unsuitable as vehicles for customer advocacy development.

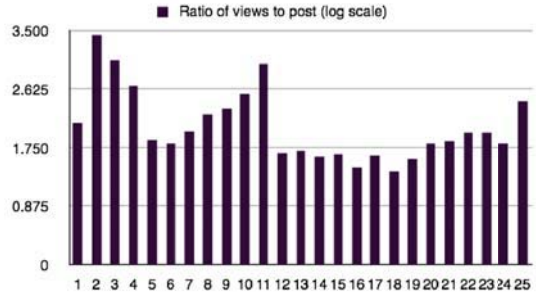

Fig. 1. *Ratio of views to posts on 25 different internet forums. On average each post is read 119 times.*

2. By extension, the use case for internet forum use is overwhelmingly that of customers seeking to discover and read a post that contains the information they want. The rise of internet search engines has demonstrated that users far prefer using search as a mechanism for information retrieval via a web-browser to the experience of information retrieval via web-surfing or browsing.

3. Multiple threads and topics can result in a fragmentation of activity in a forum which militates against the development of a community of expertise that is the commercial objective. Topic structures are initially imposed on forums arbitrarily by their owners, in best-practice cases these structures are frequently revisited in response to the use that is actually being made of the forum by its users. Unfortunately, as soon as a topic structure is created it changes users perception of the forum, and changes their pattern of use.

These three issues provided the non-technical motivations for the implementation of the Hubbub forum [7] which is described in this paper. In addition the development of lightweight Web 2.0 engineering techniques in recent years was a key inspiration to the team. In particular blogging aggregation and tagging engines such as del.icio.us provided models of implementation that changed our view of what could be done with an internet forum.

In the rest of this paper we describe the features of Hubbub which provide enhanced support to the information retrieval use case and reduce the supporting cost for BT business. We provide a qualitative and quantitative evaluation of the forum in use and conclude by pointing to future developments that we hope to make of the technology.

2 Hubbub

In this section we present our community forum solution: Hubbub (see figure 2). It offers a simple means for users to interact with each other, and exploits those interactions to provide functionality, rather than attempting to provide fully automated knowledge management. The knowledge in the system should primarily come from the human contributions, and the system should aim to help users in navigating and coordinating these contributions.

The Hubbub user journey begins with an invitation to the customer to submit their query in their own language and with their own title. This is then parsed and the keywords from the query are extracted and matched to other posts in the forum. When a post is made to the forum the keywords in it are added to the customers profile and an interest in other posts containing that keyword set is registered with the system. Each discussion in Hubbub are indexed via keywords, each of which is a word that has been selected by at least two users of the system.

The posts retrieved via the keyword indexing system are then dynamically compiled into a list for the user creating a "virtual topic" which contains all the relevant posts for the customer to browse. This is in contrast to the static topic structure that is presented by most internet forums. Because the virtual topic is built dynamically for the user it is less likely to be irrelevant. In addition because posts are retrieved with a search function posts which focus on negative issues are unlikely to be retrieved, not because Hubbub filters or removes them, but instead because they are less likely to contain keywords that match the users current post unless the user is looking for negative comments.

If the customer does not find an answer amongst the posts in their virtual topic they can modify their query by interacting with Hubbubs tagging system.

Users are offered tags which are relevant to the post, but which have a relevance below the threshold used for retrieval (see figure 3). If the posts retrieved are inadequate the user is able to select these less relevant tags and will be rewarded with a modified search result. Alternatively the user is also able to add keywords of their own which may not appear in the text of their query. This fea-

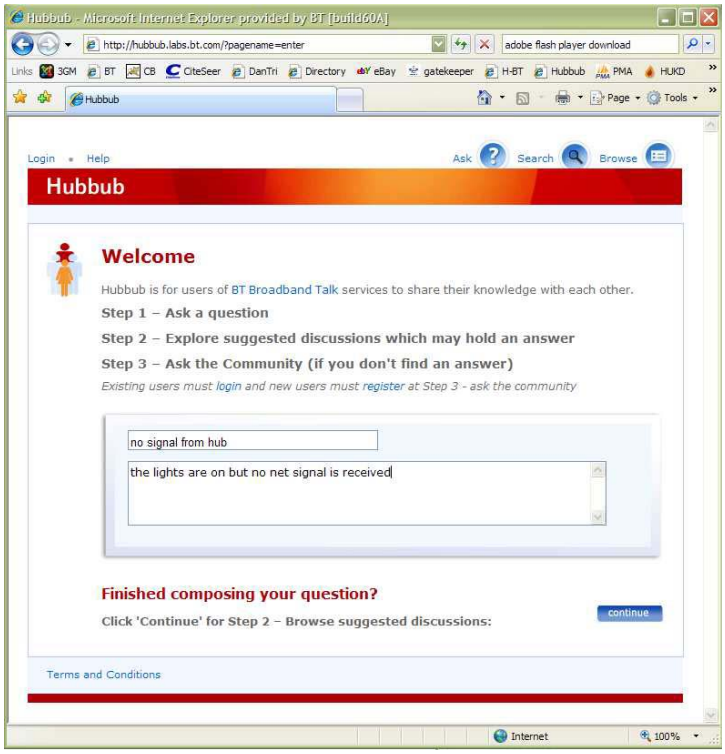

Fig. 2. *Hubbub landing page.*

ture enables tags to be introduced that are separate from the users text enabling the system to link groups of posts in its knowledge base.

If the tag modification process fails to generate an appropriate result they can post their query onto the forum. At this point they must either register or login in our current implementations, this is in order to limit the opportunity for spammers and forum abuse in general and to control the quality of the discussions on the forum.

A key feature of Hubbub is the email alert facility. As mentioned above, author will be automatically notified of any changes related to his/her discussion. Hubbub however provides a mean for end user to subscribe to other thread(s), particular keyword(s) or user(s) so that he or she will not miss any other potential source of information that could provide a solution to his/her problem. Of course, all users have total control over this email alert feature and can decide how the information will be sent to them.

Apart from being able to see relevant discussions, users can also browse the list of discussions on site and view the tag cloud composed of most frequently

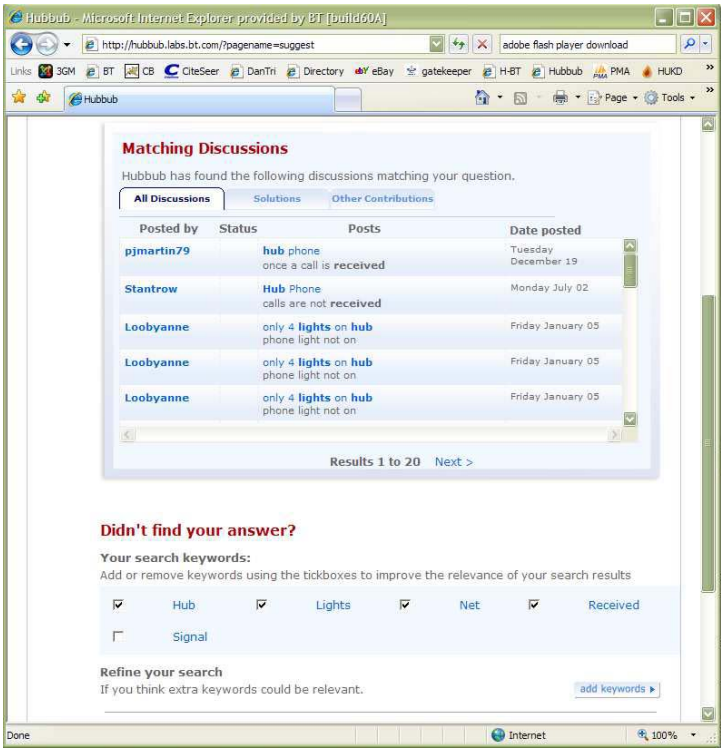

Fig. 3. *Suggested discussions and keywords.*

used keywords[1] (see figure 4). This feature provides a similar facility to that of traditional forum solutions. However, due to the category-less design nature of Hubbub, users will only see a limited subset of all the discussions on site namely latest, unsolved, most read posts. Viewing a keyword allows users to see the list of posts containing that keyword. This list can be either solved or all discussions contain that keyword.

As Hubbub is a community forum, its contents are provided by the members of the community and they do have some features to control this content themselves. For example, if a discussion is considered offensive to an user (i.e. containing an offensive word), he/she can report this discussion as offensive and the system will immediately ban this post from displaying to other users. It can only be displayed again after a customer agent has reviewed and decided that the content is appropriate. Importantly, once the customer support agent or other user with the appropriate administrative rights has asserted that a post is not offensive users will no longer be offered the option of banning it because they find it offensive. This means that it is not possible to use the community

[1] We have not identified a particularly successful mechanism for producing an informative keyword cloud on our data. Nonetheless, we have retained it as a useful navigational feature but its development is beyond the scope of this paper.

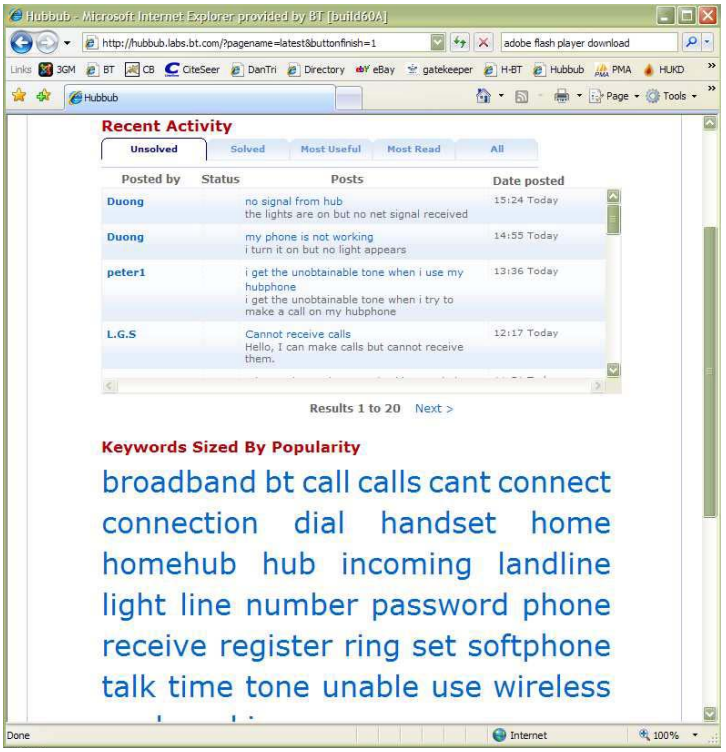

Fig. 4. *Browsing Hubbub.*

moderation feature to remove all the posts from the forum permanently as an act of vandalism.

Administrative users are able to edit the list of keywords to provide automatic deletion of posts that contain certain offensive terms. If any user tries to post a query that contain a banned word, both the query and his/her account will be automatically banned.

Another feature of Hubbub is that it provides customer support agents with the log of activities, including the list of discussions that have not been answered, list of posts/users banned from the site, etc. The agents can then decide to provide answers to such discussions or review the banned posts/users to ensure their validity.

2.1 Technical Design

Figure 5 shows the main components of Hubbub. Hubbub has been built using standard open source solutions including Apache/PHP as the web server and the scripting language, MySQL as the database design and Smarty as the template for the UI. The use of open source software was critical for the Hubbub project because it enabled implementations of the system to be created and run at very

low cost reducing the barrier of investment that potential internal customers were required to commit to.

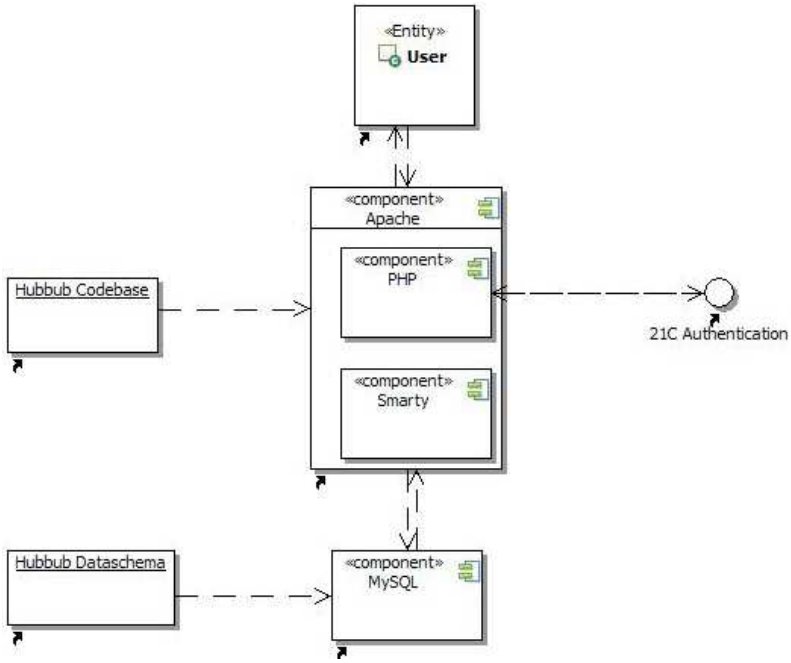

Fig. 5. *Hubbub components.*

A discussion in Hubbub consists of an opening query and is followed by a number of reply posts if they exists. Here a post will be stored with a title and a body. Keywords are automatically extracted from both the title and the body with the rule that any word that is not a stop word or common word will be considered as a keyword. The more times a keyword is ticked by an user, the higher its importance will be.

In terms of database design, there are three nodes in Hubbub which are user, keyword and post. There are number of edges that link these nodes together such as contain_edge links post with keywords (i.e. a post contains few keywords), author_edge links post with user (i.e. authorship). The nodes are indexed for faster data retrieval.

When a query is posted, the keywords are then searched (using simple text matching technique) in the edge tables to find posts sorted by the number of relevant keywords. If the user alters the list of selected keywords the search is restarted. This process is crucial to the functioning of Hubbub since it is the bridge between user queries and potential related solutions. The optimality of this search method cannot be claimed but it provides an adequate, scalable and easily engineered solution.

3 Results

Initially Hubbub has been set up as a trial for 3 months to used as a supporting channel for BT Softphone product on July 2006. After its successful trial, it has been used as a main supporting channel since October 2006. At the time of writing it has more than 5500 registered users and has served more than 4.1 million page requests of which nearly 1.5 million requests come from legitimate users (the other 2.6 million requests are from search bots - which is common for any web site that is live on the internet).

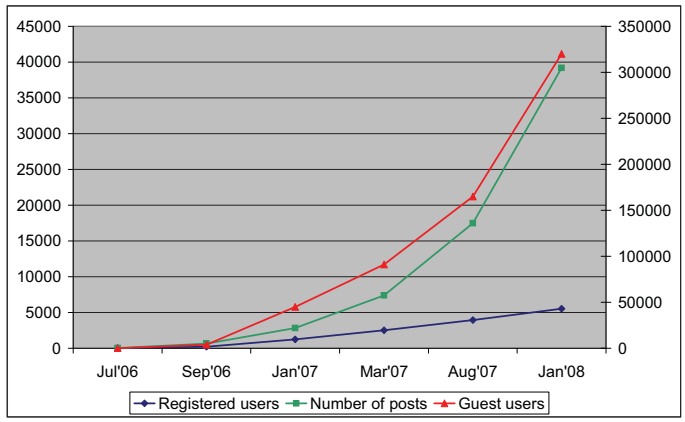

Fig. 6. *Hubbub performance. 2nd axis is for guest users only.*

Figure 6 visualize the statistics of Hubbub since launch. As can be seen, the number of posts and registered users are doubled every 5 months since launch time. In overall, the community uptake of Hubbub is rising steadily.

Table 1 shows the latest average visitor information at this point of writing (Jan 2008). With more than 5000 visitors per day, Hubbub is a significant engine for deflection from BT's contact centers.

5083	visitors per day (excluding bots)
13662	page views per day
190	visitors per day referred from BT.com
510	downloads of robots.txt per day.

Table 1. *Latest Hubbub statistics.*

It is not possible to determine why a particular user decides not to call a contact center, but if we were to measure the value of an interaction via Hubbub as equivalent to a contact center interaction on the basis of published contact

center costs we can quantify something of the business value of this approach. A survey of large organisations in the USA found that the fully loaded cost of a contact center call is between $2.70 and $5.60 [1] on this basis amount of saving per week by Hubbub is approximately £69,000. or approximately £300,000 per month. This is a significant reduction of customer support cost for a relatively low maintenance cost since Hubbub only requires a minimal monitoring (the main server has not been restarted since Oct 2007). Compared to the cost of having a real person behind the phone to give the answer, it is a substantial saving. Furthermore, by going online, the customers will be able to control the amount of information that they can see and have more flexibilities in searching for the right answers. With Hubbub, they are allowed to alter their queries and keywords to search for different posts that might be relevant to them. They might not be able to have such flexibility when dealing with real person behind the phone. In our opinion, it might be beneficial for the supporting agents to actually make use of Hubbub as well but we are not able to try out this idea yet.

Based on the results we gathered during the period of 18 months, it can be seen that Hubbub is quite a success. Users has made 39,000 queries in total and only 12,000 have been registered as a post. Furthermore, of these 12,000 posts, 6000 have at least one answer from other community members and 950 have been marked as solved. Thus, the optimistic reading of the success of Hubbub can be projected at 71.7%[2].

4 A Taxonomy of Internet Forums

There are a number of successful forum solutions available at the time of writing. In this section we will describe some decision making factors that should be considered before selecting the appropriate solution and then outline the main types of forum that are in the market currently.

4.1 Discussion Forums

A number of companies and opensource projects offer complete forum solutions for adoption. Some widely used examples are Jive Forums, mvnforum and Lithium/Rightnow. These solutions can either be installed on a local infrastructure or, in many cases rented on a software as a service basis. The all provide a way for users to read publicly posted messages and to post messages of their own for other peoples consumption, but beyond that basic functionality it is hard to identify their differentiating characteristics.

Wikipedia [5] offers a comparison of 84 forums on the basis of the following features:

[2] Here we assume that if an user enters his or her problem and found a solution based on Hubbub suggestions, that query will not be registered as a forum post. Thus, the success percentage can be measured by (number of unregistered queries + number of solved posts)/total number of queries.

– Flat or Threaded: does the forum support threads or not
– User-selectable themes: is it possible for users to customize the look and feel of the forum
– Calendar functionality: can events or other calendar related information be integrated onto the forum
– Image attachment: can users put images on the forum
– Unread message tracking: can unread messages be detected
– WYSIWYG editing : is there an advanced editor for message formatting on the forum

While the list provided is extensive, these features are not really useful for understanding the usefulness or otherwise of the forums described as channels for customer support. For example, none of the requirements gathering exercises that we undertook in BT identified Calendar, or WYSIWYG functionality as needed. In fact it seems that there is almost complete uniformity in the core features of these systems, which is interesting given the issues that we have identified in this paper.

4.2 Community Question Answering

Yahoo Answers [12] has become one of the most successful sites on the internet by providing a place where users can ask question and receive answers from a community of experts.

In spirit and execution Yahoo Answers is superficially similar to Hubbub, but differs in that Hubbub is specifically designed to serve information rapidly to the user and deflect further enquires, whereas the Yahoo design seems intended to draw users into further interactions. For example, there is no step to present users with similar questions when they post an enquiry on Yahoo answers, and no interactive search mechanism (see Section 3 of this paper for information on how Hubbub does this).

Microsoft have adopted a similar approach with the Microsoft Office question answering system [10] where user can ask a brief question and then be presented with relevant discussion topics. If their question cannot be answered by these topics, they have the option of posting it in the community for other members to reply. It is also similar to Hubbub but lacks the ability for the users to fine tune their query. Here the users are not able to interact with the search process and if the presented solutions are not relevant, users will have to ask another question. On the other hand, Hubbub allows the user to fine tuning their search by adding or removing keywords which in turn will directly affect the list of suggested posts that users can read (see section 2).

Another example of this question answering system is GetSatisfaction [6]. It is a pay-per-support commercial platform that various companies can register to host their support and provide their agents to work on site. GetSatisfaction provides such companies with the forum infrastructure so that end users can search for solutions to their problems on site and get answered by real people if

no solution could be found. At this time of writing, there are about 200 companies with around 40000 posts in total. However, their system has the same disadvantage which is similar to the Microsoft Online where users are not able to interact with the search process.

The Start system from MIT [9] is another question answering system that takes an alternative tack. Whereas Yahoo Answers and Hubbub rely on a human community of experts Start uses advanced AI to attempt to automatically answer user questions from a formal knowledge base. At this time, while interesting, the performance of this kind of system in open domains such as customer support has not been proven.

These examples have clearly demonstrated the new trend in supporting forums in which community question answering systems play a pivotal role. Getting away from the traditional browse and search journey provides user with much greater flexibility and efficiency and allows them to access the right information in much lesser time than previously needed.

4.3 Product Support Forums

Ksamba [8] and Fixya [4] are two well known examples of websites that offer support for products and services from human experts.

They differ somewhat from the Yahoo model in that the objective of the site is to steer the customer into a one-to-one interaction with one of the experts on the site, enabling the expert to provide consultancy in return for remuneration.

This kind of model presupposes that the customer for a product or service does not feel that their relationship with the service provider places the provider under an obligation to support them for free. If there is no chargeable interaction then the site cannot fund itself, and there is no efficiency gain from publishing the expertise elicited from the expert by the customer in the interaction on the site because it is not in the interest of the expert to allow this to happen.

Table 2 summarizes the list of supporting forums that we have discussed in this section.

Features	Lithium Rightnow	Hubbub	Y! Answer	Microsoft GetSatisfaction	MIT - Start	Fixya Kasamba
Query-specific search	*No*	*Yes*	*Yes*	*Yes*	*Yes*	*No*
Category-based browsing	*Yes*	*No*	*No*	*No*	*No*	*No*
Direct human contact	*No*	*No*	*No*	*No*	*No*	*Yes*
Email notification						
- posts	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*	*No*
- keyword/users	*No*	*Yes*	*No*	*No*	*No*	*No*
Solution marking	*No*	*Yes*	*No*	*No*	*No*	*No*
Customize search process	*No*	*Yes*	*No*	*No*	*No*	*Yes*
Administrative functions	*No*	*Yes*	*No*	*Yes*	*No*	*Yes*

Table 2. *Comparison between different forum solutions.*

5 Conclusion and Future Work

Hubbub has been successfully deployed as an innovative customer service solution which combines insights into the fundamental structure of customer forums with web 2.0 technology for tagging and social networking. It has been proven to be popular with customers and BT has plans to use it as a support channel for several new products and services.

Strategically Hubbub will need to be integrated with the BT information architecture under the group wide rule of one for solutions. We are engaging with BT's suppliers and systems architects to attempt to ensure that this happens as efficiently as possible. In parallel we are exploring how to make our solution available for other companies to adopt.

The social and interaction aspects of Hubbub also require further analysis. The motivation of users to participate in Hubbub is complicated by the business objective of the system which is not simply to encourage discussion or community problem solving but also to be an efficient call deflection knowledge base. Therefore the map of corporate ownership and community territory differs for Hubbub when compared to other forums.

A number of research projects are currently underway to develop advanced technology to improve Hubbub's functionality. These include investigations into using the knowledge base as a recommender system for customers and visualization techniques for the management of the forum. We hope to be able to report on these projects in subsequent publications.

In the not too far future, we envisage user forums to have more human like processing power such as the ability to integrate a natural language parser to have a better understanding of user queries, to apply learning techniques in improving both the accuracy and the efficiency of search results for the users. In addition, we also expect internet forum to become a standardized component of an open web application in which end users can fully customize in order to find the right information for themselves with minimal effort.

References

1. C. Barry. *Managing your Cost Per Call.* MultichannelMerchant.com, http://multichannelmerchant.com/opsandfulfillment/contact_center_advisor/cost_percall/, January, 2007.
2. Z. Cui, G. Ducatel, M. Thint, B. Assadian, and B. Azvine. Towards automated customer self-help. *BT Technology Journal*, 24(1):96–106, 2006.
3. J. Dyche. *The CRM Handbook: A Business Guide to Customer Relationship Management.* Addison Wesley, 2001.
4. Fixya - Technical Support, User Guides and Repair Service. http://www.fixya.com.
5. Forum software comparison in Wikipedia. http://en.wikipedia.org/wiki/comparison_of_internet_forum_software.
6. GetSatisfaction - People Powered Customer Service. http://getsatisfaction.com/.
7. Hubbub Community Forum. http://hubbub.labs.bt.com.
8. Kasamba - Live expert advice. http://www.kasamba.com.

9. B. Katz, S. Felshin, D. Yuret, A. Ibrahim, J. Lin, G. Marton, A. J. McFarland, and B. Temelkuran. Omnibase: Uniform access to heterogeneous data for question answering. In *Proceedings of the 7th International Workshop on Applications of Natural Language to Information Systems (NLDB 2002)*, June 2002.
10. Microsoft Office Only Community. http://www.microsoft.com/office/community/en-us/wizard.mspx.
11. G. Urban. *Don't just relate - Advocate*. Wharton School Publishing, Upper Saddle River, New Jersey 07458, 2005.
12. Yahoo answers. http://www.yahoo.com/answers.