

IBM Research Report

Closing the Loop with ContactPoint: A Framework for Information Distribution and Collaboration

Maroun Touma, Carolyn Brodie, Catherine Wolf

IBM Research Division
Thomas J. Watson Research Center
P.O. Box 704
Yorktown Heights, NY 10598



Research Division
Almaden - Austin - Beijing - Haifa - India - T. J. Watson - Tokyo - Zurich

Closing the Loop with ContactPoint: A Framework for Information Distribution and Collaboration

Maroun Touma touma@us.ibm.com
Carolyn Brodie brodiec@us.ibm.com
Catherine Wolf [cwf@us.ibm.com](mailto:cwolf@us.ibm.com)
IBM Thomas J. Watson Research Center
30 Saw Mill River Road
Hawthorne, NY 10532

ABSTRACT

This paper describes ContactPoint, a web based application that assists members of large, distributed communities to stay informed about what is happening in that community and to act on that information. ContactPoint consists of a thin java client and a set of server functions that provide users with information alerts and a list of subject experts. Collaboration between the users and these specialists is supported. ContactPoint's goal is to assist communities to stay informed by augmenting more traditional methods such as websites and newsletters.

INTRODUCTION

This paper describes ContactPoint, an application that assists members of large, distributed communities to both stay informed about what is happening in their community and to act on that information. It is very difficult for large, distributed communities to not only keep their members informed about what is happening in their community and industry, but also assist members in finding contacts within the organization who can provide more details about specific activities. ContactPoint assists large organizations by providing functionality that allows members to be alerted to new information quickly and unobtrusively, and to identify and contact subject matter specialists that can provide additional information.

In this paper, we will provide background on some of the problems faced by large, distributed communities. We will then present the ContactPoint architecture and discuss how this architecture allows ContactPoint to address these problems. Finally, we will briefly describe how ContactPoint has been used by one large, distributed community to keep its members informed about what is happening within the community and to facilitate collaborations within the community.

MOTIVATION

As more and more organizations have memberships that are distributed all over the world, it becomes more difficult for members of these organizations to stay informed about what is happening in their organization, their company, and their industry. To keep members

informed, large organizations often send newsletters and/or provide websites that provide highlights on the major activities and projects in which different groups of members are associated. However individuals often miss new information using these techniques because newsletters get lost in a sea of e-mail [7] and individuals do not know when websites have been updated, and may not check often. Furthermore, even if members do see the new information, it is often difficult for them to find contacts to help them obtain additional information if they are interested. Finally, whatever method an organization uses to keep its members informed, the members may not be interested in the information exactly at the time that it becomes available, but sometime later. ContactPoint addresses these problems by providing functionality to: 1) allow an organization to alert its members to new information without adding to the problem of e-mail overload, 2) assist the members in identifying subject matter specialists who can provide them with more detail about the information, 3) assist members in collaborating with these subject matter specialists, and 4) provide an easy to access repository for the information that will remain available to the members over time.

To provide this functionality, ContactPoint has been designed as an "alert system" that has been extended to allow users to respond to the alerts pushed to them. An alert system is a form of push technology [4] that has the goal of providing users with information content with a minimal delay at the user's end. This is generally accomplished by having the information provider push information to the user's workstation, personal computer, or wireless device, and then having the push client display that information directly in the user's workspace. Some examples of alerts systems include BackWeb [5] and InfoGate [6]. Like these systems, ContactPoint can push users many different types of information ranging from HTML documents to videos. However, ContactPoint differs from these systems in that it provides functionality to allow users to respond directly to the information they have been pushed. This is accomplished by providing a list of subject matter specialists along with each article and mechanisms for the user to collaborate with these specialists. The other

Server

The majority of ContactPoint functionality is on the server. Five major functions of ContactPoint that reside on the server will be described in this section. These include the User and Expert Directory, the Alert Distribution Agent, the Content Delivery Proxy, the User Interest Profiling Agent, and the Collaboration and Messaging Services.

User and Expert Directory:

ContactPoint uses the Lightweight Directory Access Protocol (LDAP) to create directories of users and experts. Users are individuals that subscribe to one or more information sources using ContactPoint. Experts are a special class of users who are identified by the information source editors as subject matter specialists. Experts provide additional information or collaborate in other ways with other members of the community regarding various subjects.

The directory stores all information pertaining to the users including *demographics information, password(s) and access controls* to specific websites or sections of the website, *preference settings* and *profiling information*. Users' *demographic information* allows the system to tailor certain services to meet specific requirements such as filtering of certain content based on the users geographical location, routing of the alerts to local data replicates where content customization could be performed or virtual ids could be managed for support personnel located in different time zones. (e.g. "Contact The Editor" function could be routed to different operators around the globe depending on the time of day). *Passwords and access control* information are needed to authenticate the user for different services of the site (Sametime chat, protected areas of the web site, etc.). It could also be used to allow ContactPoint users access, from time to time, to external resources such as a partner's websites and/or other web utilities that require user authentication. The *preference settings* and *profiling information* provide a global store for user defined preferences as they pertain to information "pushed" by the site. The *preference settings* define the way the alerts are presented to the user (level of intrusiveness, frequency, etc.) and are implemented using native LDAP attributes. The *profiling information* define the type of alerts that can be sent to the user and are implemented using native LDAP groups. Profiling will be discussed further in a later section.

Alert Distribution Agent:

The Alert Distribution Agent is implemented as a Sametime chat client. It continuously scans the alert database for new submissions and uses "buddy lists" to register those users with pending alerts. When a registered user comes online, the Agent is notified of his

or her new status via the awareness mechanism built into the Sametime server. At this point, it sends an encoded chat message to the user's java client to initiate the alert download process specifying the URI of the alert descriptor document. The chat session is then terminated by the Alert Distribution Agent once it receives confirmation the message was properly delivered to the client and reflects the delivery status in the alert database. Upon receiving the notification of a new alert, the client initiates an HTTP connection to retrieve the alert descriptor document. As part of the HTTP request, the client sends a copy of the latest descriptor document it had received and the URI for the new document it is about to receive. The Alert Distribution Agent uses this step to compare the two documents, confirm the status of previously received alerts and insure that a new or modified alert is indeed awaiting the user. If a new alert is confirmed the HTTP response from the Alert Distribution Agent would include the actual content of the new alert descriptor document or a redirect to retrieve the document from another server on the net. Once the retrieval process is completed, the client initiates the same HTTP transaction described above thus allowing the Alert Distribution Agent to update its own alert database. This keeps the server synchronized with the alerts being displayed on each of the clients. The Alert Distribution Agent also uses the same mechanism to notify the client when certain documents that relate to previous alerts are no longer needed and can be deleted, freeing up local storage space. These documents will remain available to the user through a database of past alerts. This allows the management of the local cache to be done remotely rather than by the java client running on the user's machine. Therefore, changes in policy can be implemented on the server without requiring changes to all of the clients.

Content Delivery Proxy:

The Content Delivery Proxy serves as a repository for all alert documents and a proxy to other web servers on the network. As a general convention, all documents that are retrieved from the local repository can be safely replicated to the client's local cache and therefore be viewed by the user when he or she is not connected to the network. However, content that is retrieved from other web servers on the network cannot be guaranteed to have the proper elements for offline browsing and therefore are never replicated to users' local cache. One example of this is a secured document that requires special authentication procedures stored in the user database on the ContactPoint server. Having these documents routed through the Content Proxy relieves the user from seeing these authentication procedures.

User Interest Profiling Agent:

During a user study of an early version of ContactPoint, it became apparent that alerts sent to users must contain information that is both interesting and useful to them. To accomplish this goal user profiles were developed for ContactPoint users. Many different approaches have been found for capturing user interests for different types of applications [1] [3]. Based on the results of our study, we found that ContactPoint users' interests could be defined using three distinct, independent types of constraints. The three types of constraints on interests are represented as the three axes in Figure 3. The y-axis represents geography. Many of the users that we interviewed were only interested in information that pertains to the geographies in which their customer(s) reside while other users were interested in all geographies. The x-axis represents industries in which the users were interested. Some users wanted information relating to just one industry, while others were interested in multiple industries. The z-axis represents the types of articles that users are interested in receiving. Some of the types of articles identified include: new products, events to which customers can be invited, industry research reports, strategies, executive messages, and success stories. Users' interests in the various types of articles depended both on their jobs and their personal preferences. Generally, an article had to pertain to at least one interest on each axis to be of interest to the user. Using the three different categories of interests, we represent each user's interests as a disjunction of interest points. Interest points are the points that are created by the intersection of three planes in which one of the three planes is parallel to the xy plane, one is parallel to the $y-z$ plane, and one is parallel to the xz plane. The set of users who share a given interest point form an interest group, so that there is one interest group at each possible intersection point. These interest groups are represented by the ovals in Figure 3.

The ContactPoint profiles are developed using a combination of methods. First, users are given the opportunity to specify interests. Users interests are stored by adding users to interest groups (native LDAP groups) based on these inputs. For example if a user specifies that he or she is interested in market research reports, new products, banking and finance industries, and only in articles relating to the United States, that person would be added to four interest groups: 1) market research on banking in the US, 2) market research on finance in the US, 3) new products that can be used for banking in the US, and 4) new products that can be used in finance in the US. It is important to note that one article may belong to several of these categories. If the user does not submit his or her interests, he or she will initially be added to all groups and receive all alerts, until the second mechanism, an interest learning

module, has a chance to learn something about the individuals interests. Based on the articles that users actually choose to view, they are asked if they would like to update their interests. For instance, if a user is a member of a group, but never opens any of the articles that match that specific interest group, the user will be asked if he or she is really interested in the subject.

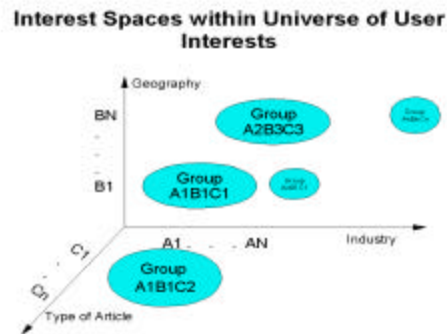


Figure 3 – Space of User Interests

Collaboration and Messaging Services:

One of the most important functions of ContactPoint is the ability it gives users to identify and collaborate with subject matter specialists when they are interested in a given article. ContactPoint provides three methods for supporting this collaboration: instant messaging, e-mail, and phone information. Instant messaging services are particularly convenient if the specialist is available because it allows the user to have quick, informal discussions with the specialist [2]. Each ContactPoint user is given a Sametime user id when they subscribe, and is logged on whenever ContactPoint is running on the system. If a particular specialist is either not available to take instant messages or does not wish to receive instant messages, ContactPoint provides users with an e-mail form which uses javamail and an SMTP server to send email to the specialists automatically. The specialist's e-mail address is automatically entered into the mail form so that the user does not need to spend any time searching for e-mail addresses. The third option is a direct access to a database that provides the user with the specialist's phone number.

USING CONTACTPOINT TO SUPPORT A LARGE MARKETING ORGANIZATION

The organization that we worked with is a large, distributed financial marketing organization whose members are spread over five continents. The individual members are marketeers whose customers are banks, insurance companies, investment firms, and other financial institutions. Within the organization is a group whose role it is to keep the members informed of developments within the industry, customer-related information, and corporate news. Given the fast pace of the financial industry and the geographic distribution of

the organization, this is a challenge. Prior to ContactPoint, the group published a newsletter that was distributed by email and maintained several websites for the purpose of internal communication. ContactPoint has been used to augment and complement these sources of information. The first version of ContactPoint was launched in September, 2000. There are currently about 700 users of ContactPoint. In this section, we discuss how ContactPoint has met many needs for this group.

Many users are overwhelmed by the amount of e-mail they receive. As one ContactPoint user explained, "I received 75 e-mail messages yesterday and I had to reply to two thirds of those. Do you know how long it takes to go through all of those." Time critical and other important information within a newsletter can be overlooked with so much email. ContactPoint brings the information to the user's attention by pushing alerts to the user's desktop. In this marketing organization, people were particularly interested in knowing about things such as customer events and industry happenings.

When members of the organization read the article or watch the video linked to an alert, they may have specific questions about the subject of the alert. If the topic is an unfamiliar one, it can be "very difficult", according to one user, to find someone with the appropriate expertise. This is likely due to a combination of factors: the size of the organization, its geographic distribution, and the fact that members spend little time in the office, but are instead visiting customers or traveling. The expert feature of ContactPoint provides the name of a subject matter expert for collaboration and, thus helps members get the answers to their questions in a timely manner. By providing asynchronous and synchronous channels for communication, ContactPoint supports members in the same time zone or a different time zone as the expert.

The first version of ContactPoint did not have a user interest profile (we had anticipated the need for this feature but had not implemented it in the first version). Consequently, all users got all alerts. Users expected the alerts to be relevant to their interests. For example, a member in Hong Kong might not be interested in a customer event in Paris. Too much irrelevant information diluted the value of the tool. Thus, interest profiles were designed and implemented for the second version of ContactPoint. Users were always on the look out for material that was relevant to their customers, such as success stories or solutions that might be applicable. However, if the information did not arrive at the time it was needed, members of the organization wanted to be able to save it for possible future reference. Otherwise, it was likely to be forgotten. As one user said, "I am interested in solutions for customers, but the

timing is not relevant to me. I would like to be able to browser through the success stories that are sent. " The past alerts feature provides access to past alerts, and in addition, subject matter experts for past alerts.

CONCLUSIONS

In this paper, we have presented ContactPoint, a collaborative, web based application that allows members of a large distributed organization to stay informed about what is going on within the organization and to collaborate with subject matter specialists regarding this information. ContactPoint provides these functions using a thin java client which has three main functions: retrieve new alerts when notified of their existence, notify the user of this new information, and provide clients for the collaboration and messaging services. The server provides all the rest of the functionality. The server distributes alerts after they have been filtered for user interests, updates user interests based on both user input and activity, provides collaboration functionality, and shields the user from extra security steps he or she might otherwise have to take. ContactPoint is being used by a large, distributed marketing organization to help members to both stay up-to-date and collaborate about subjects that are of interest to them. This application has not been used to replace more traditional methods such as newsletters or websites but to augment them by allowing users to be alerted to updates and then provided with automatic links to both useful websites and subject experts, thereby helping members of the organization to be effective in their jobs.

REFERENCES

1. Kramer, J., Noronha, S., and Vergo, J. A User-Centered Design Approach to Personalization. In *Communications of the ACM* V43 N8 August, 2000 pages 45-48.
2. Nardi, B. A., Whittaker, S., and Bradner E., Interaction and Outeraction: Instant Messaging in Action in *Proceedings of CSCW2000*. (Philadelphia, PA), ACM Press, 79-88.
3. Tu, H. and Hsiang, J. An Architecture and Category Knowledge for Intelligent Information Retrieval Agents. In *Decision Support Systems* V28 N3 May, 2000 P255-268.
4. Underwood, G., Maglio, P, and Barrett, R. User-centered Push for Timely Information Delivery. *Computer Networks and ISDN Systems* 30 (1998) 33-41.
5. URL:<http://www.backweb.com>
6. URL:<http://www.infogate.com>
7. Whittaker, S. and Sidner, C. Email Overload: Exploring Personal Information Management of Email. In *1996 Conference Proceedings on Human Factors in Computing*. (Vancouver, Canada), 276-283.