Developing and Deploying Knowledge on a Global Scale

James Borron
Reuters America
360 Motor Parkway
Hauppauge, NY 11788
james.borron@reuters.com

David Morales
Reuters America
311 S. Wacker Dr., Suite 1000
Chicago, IL 60606
david.morales@reuters.com

Philip Klahr
Inference Corporation
100 Rowland Way
Novato, CA 94945
klahr@inference.com

Abstract

Reuters is a worldwide company focused on supplying financial and news information to its over 40,000 subscribers around the world. To enhance the quality and consistency of its customer support organization, Reuters embarked on a global knowledge development and reuse project. The resulting system is in operational use in the North America, Europe, and Asia. The system supports 38 Reuter products worldwide. This paper is a case study of the Reuters experience in putting a global knowledge organization in place, in building knowledge bases at multiple distributed sites, in deploying these knowledge bases in multiple sites around the world, and in maintaining and enhancing knowledge bases within a global organizational framework. This project is the first to address issues in multi-country knowledge development and maintenance, and multi-country knowledge deployment. These issues are critical for global companies to understand, address and resolve, in order to effectively gain the benefits of global knowledge systems.

Introduction

Reuters Holdings PLC supplies the global financial and news media communities with a wide range of products and services including:

- real-time financial data
- transaction systems for financial trading
- access to numerical and textual historical databases
- news
- graphics
- still photos and news video.

Reuters has 40,000 subscribers, 309,000 user accesses, and operates in 154 countries. Reuters information is accessed through a series of Reuter products. These products, and the real time data feeds, are supported through customer support help desks around the world.

While Customer Service Operations have been in place in Reuters for many years, it is only within the last few that help desks have risen to prominence in the company as a key competitive differentiator. There are three key areas in which Reuters competes:

- in the data and news content Reuters provides
- in the technology with which that data is delivered
- in customer service.

Providing outstanding customer support encourages customer loyalty, supports repeat business, and promotes a reputation for customer orientation. Within the last few years, Reuters has aggressively invested to provide significant improvements in its customer support organization and operations.

As the organic growth of help desks and technical support functions developed from local to continental initiatives, Reuters realized that a great deal of customer and technical support expertise was developing within the company which might effectively be reused. A global steering group was formed, and help desks in the United Kingdom and the Americas began exchanging staff and support materials. The idea of encoding and reusing knowledge is an extension of the initial effort to move people and their knowledge around the world.

Project History

In February, 1993, Reuters America (RAM) engaged Inference Corporation in a re-engineering study based on RAM's Chicago Customer Response Center to explore opportunities for significant improvements in customer support. Key findings of this study included:

- opportunities for improving inefficient internal systems
- an environment oriented around specialists, and not conducive to generalists
- issues in staff training and key competencies

- development of a strategy to preserve knowledge assets.

The strategy involved building knowledge systems using *Case-Base Reasoning* (CBR) [Kolodner, 1993]. Other customer support organizations had already reported successes in using CBR technology (as subsequently reported in IAAI conferences [Acorn & Walden, 1992; Nguyen et al., 1993; Hislop & Pracht, 1994]) and Reuters decided that CBR would be appropriate for developing knowledge systems for use in their help desks as well.

In June 1993, RAM launched its first case-base project. Initial tasks included:

- forming a project team
- initial authoring of a style guide
- agreeing on product domains to cover
- committing time for case authors (people building the case bases)
- developing an incentive program
- setting and achieving targets.

In December 1993, RAM deployed its first case base. It had approximately 1,200 robust cases. It had especially good response and feedback from new hires who were able to quickly use the tool to provide both self-training on an on-going basis, and an expert advisor to increase their competency and ability. The case base was also a strong source for infrequent and complex situations. As users learned the more routine situations, their use of the case base evolved to the more complex and unusual situations. This ability for users to come up to speed on the routine situations was significantly enhanced with the case-base tool.

There were, though, some difficult issues. Developing the case base required resources to *author* the knowledge.[1] Given the dozens of Reuter products available, the generation of knowledge bases would require a more critical mass of authors. The few authors that generated the initial 1,200 cases would be insufficient. Also important was that even though much expert knowledge existed in Chicago, there was significant expertise elsewhere at Reuters -- in London, Continental Europe, Asia, the Middle East and Africa. To exploit this knowledge was vital in generating complete and accurate case bases.

While RAM pursued CBR technology aggressively in 1993, other areas in Reuters were also pursuing CBR initiatives. Reuters United Kingdom and Ireland (UKI) was involved in CBR as early as mid-1992. UKI's principal focus in 1992 and 1993 had been on developing a call tracking and problem management system called CALLS. UKI used Inference International consultants to help build CALLS as well as to provide guidance on using and integrating CBR technology for problem resolution. Some initial case bases were built in UKI in the 1992-1993 timeframe which also helped RAM in its evaluation of CBR in early 1993. A third effort in CBR at Reuters was also simultaneously occurring the Reuters Middle East and Africa (MEA) organization. In particular, the South Africa help desk quickly and independently developed case bases to support their needs. Consultants from Syscon in South Africa supported this effort.

Given the three multiple CBR efforts at Reuters, and the distribution of expertise around the world, Reuters decided in February 1994 to form a Global Customer Support Steering Committee to organize a global case-base effort. Their initial objectives were to:

- consolidate the case-base projects in America, United Kingdom, and Africa
- increase domain coverage
- more effectively use the best resources for authoring worldwide
- establish a global funding mechanism
- organize a global management organization.

Interestingly, the idea of a distributed model for case authoring arose from the Chicago-based product experts' need to work from home. These experts needed isolation from the running of the day to day operation to be most productive. While at their desks, these senior-level experts are continually interrupted, as experts often are, and unable to devote the necessary time to writing good cases. Typically one writes a family of cases in one sitting, and continuous concentration is conducive to a well-designed case base.

As a result, expert authors were outfitted with home PCs and access to their supported products. In addition, they were given detailed instructions for writing cases in such a way that there would be no corruption to the database when their work was merged with the master case base file. From this experience, it dawned on the project team that cases could be written anywhere, whether 5 miles from

[1] We use the term *author* to describe the knowledge engineering/knowledge acquisition process. Authoring is the customary term used in the CBR customer support community. It provides a more concise active description of the entire process (acquiring the knowledge and encoding it). We discuss later the alternatives of having an expert author cases directly or having a knowledge engineer interview experts and then encode the knowledge.

downtown Chicago or 5,000 miles away!

Knowledge can be captured easily wherever it resides, and it was this conclusion that precipitated the drive toward a global case base.

By July 1994, a plan was established for an 18-month global project. This plan was approved by the Global Steering Committee, and the global project began in August 1994. The 18-month period can roughly be divided into three principal phases, each of 6-month duration:

Phase 1: Project initiation and organization, formulate budget, establish global procedures, develop global style guide, generate initial knowledge base (combining and globalizing knowledge from three existing efforts), and prototype initial supporting software utilities.

Phase 2: Complete first version of software utilities to support global procedures, solidify quality assurance procedures, expand knowledge building process (here focused on increasing the number of products/domains covered), provide initial user training, deploy and beta test initial versions of knowledge base, and compute return-on-investment metrics.

Phase 3: Further expand knowledge building to a critical mass (now focused on increasing depth of knowledge for products/domains covered), complete second version of software utilities, quality test complete system, complete documentation (utilities, global procedures, training materials), and complete global rollout with more robust knowledge base.

Figure 1 provides an overview of the project timeline.

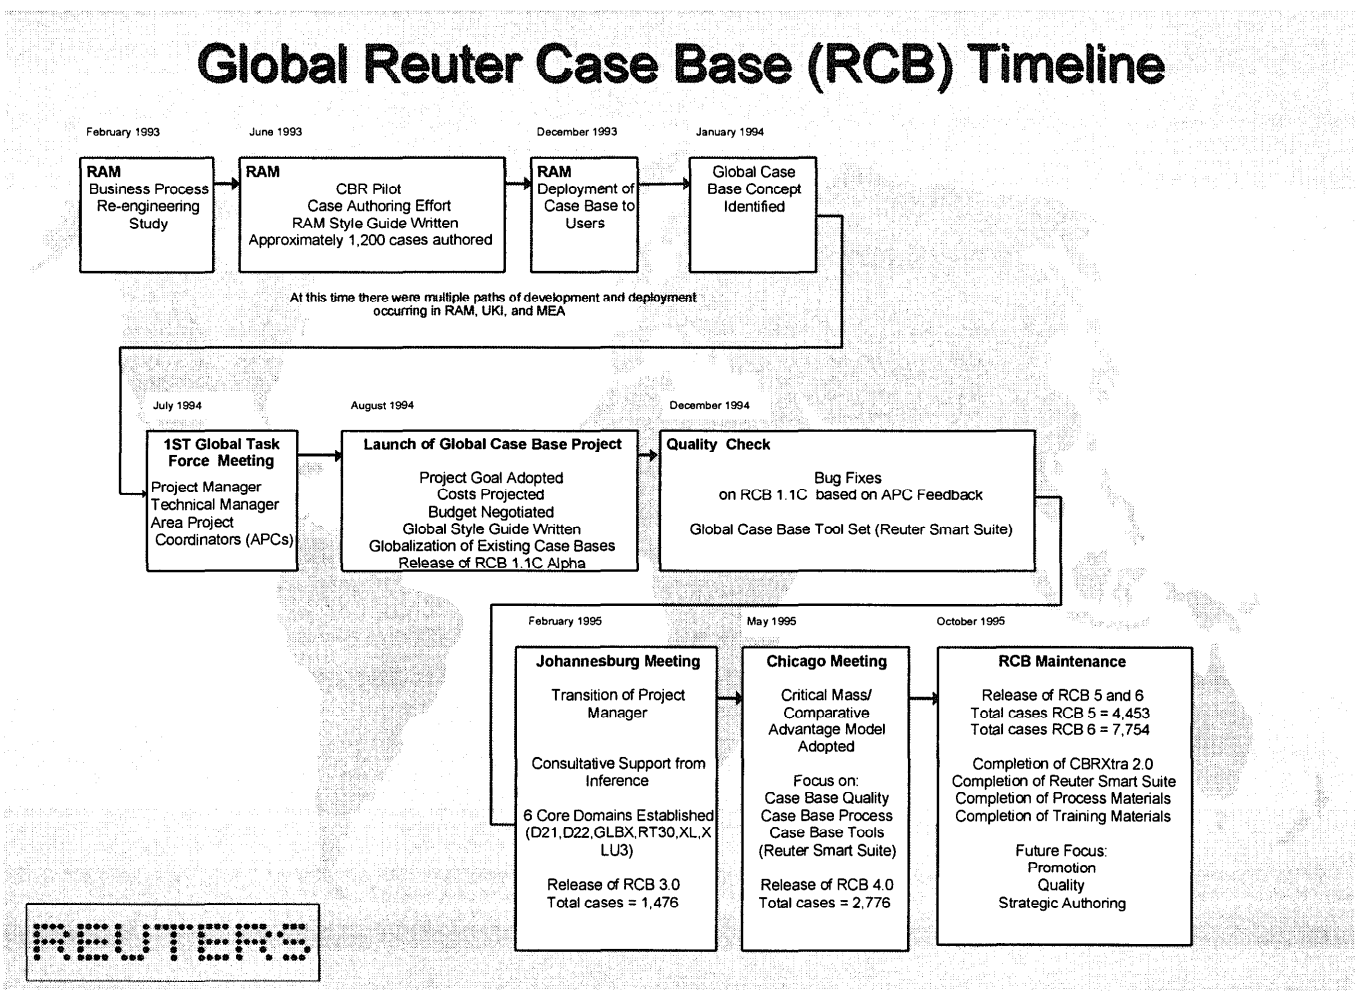

Figure 1. Global Reuter Case-Base Project Timeline.

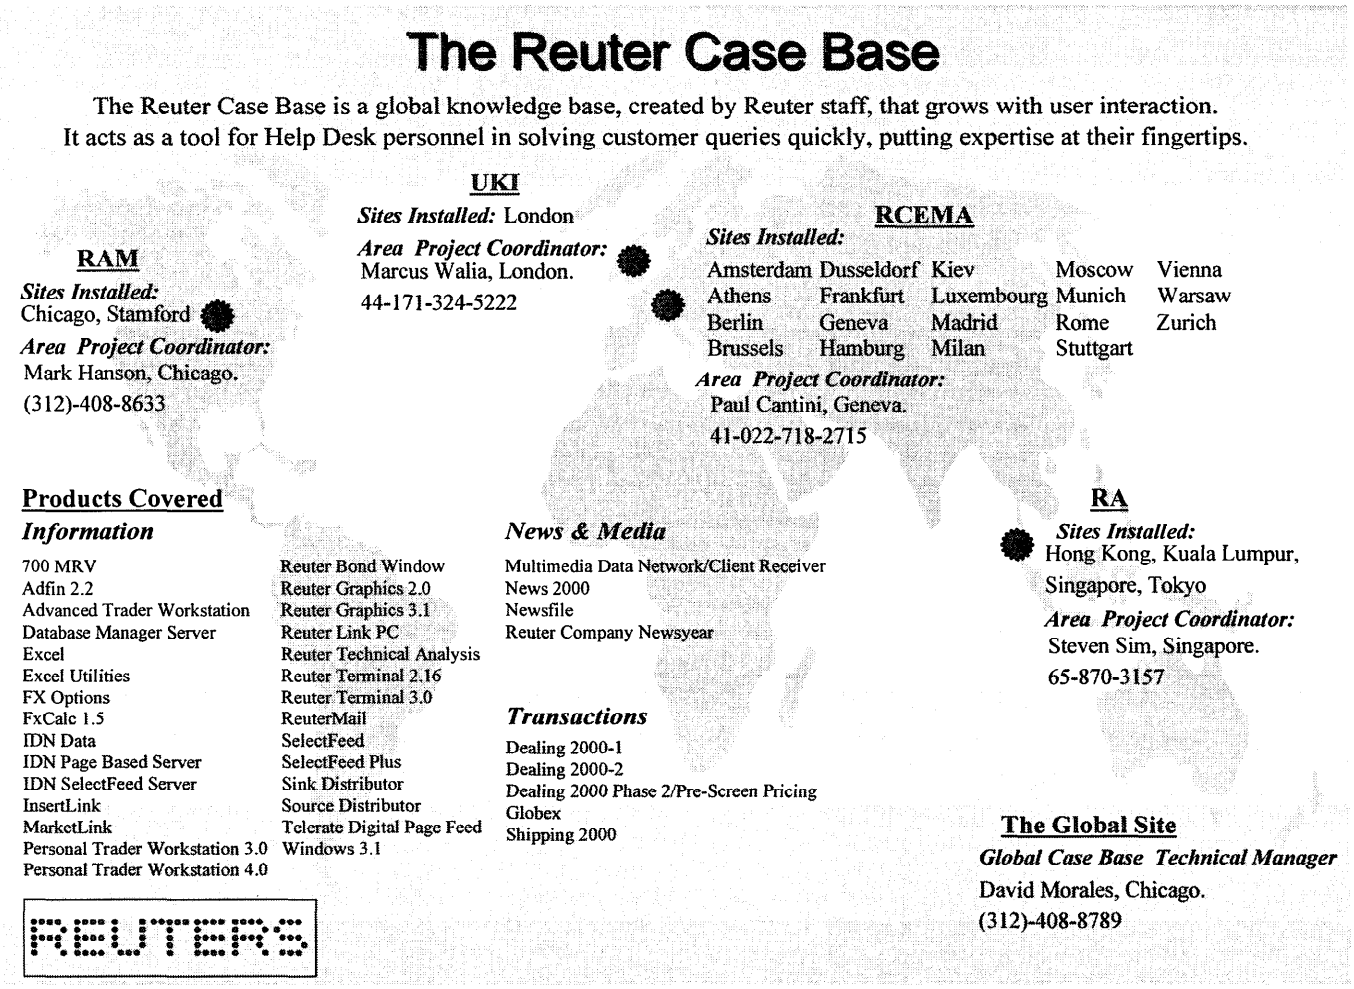

Figure 2. Reuter Case-Base Deployment Sites and Products Covered.

The 18-month effort is now reaching its completion. It has achieved all of its objectives. As of November 1995, the Reuter Case Base

- contains 7,754 cases
- covers 38 Reuter products and services (Figure 2)
- is installed in 29 sites around the world (Figure 2)
- is used by 190 users who are directly servicing Reuter customers.

Global Project Organization

The *Global Steering Committee* consists of the heads (vice president or director level) of customer support from each of the then five (now four) global areas: Reuters America (RAM), United Kingdom and Ireland (UKI), Reuters Asia (RA) Reuters Continental Europe (RCE), and Middle East and Africa (MEA) (RCE and MEA have since been combined), plus a senior customer support representative from Reuters corporate. The Global Steering Committee meets quarterly to review project milestones, issues and directions. Biweekly phone conference calls provide interim status updates and address immediate issues and needs.

The principal global organization structure includes five logically-defined job/skill positions:

- Global Project Manager
- Global Technical Manager
- Area Project Coordinators
- Domain Authors
- Case Authors

The *Global Project Manager* led the project and coordinated the development and deployment efforts in the five regions. This was certainly one of the more difficult

challenges in the project. Not only are there significant time-zone differences among the five regions, but also differences in culture, in software development practices, in structure and management organization, and even in project objectives. In addition, from a knowledge perspective, there were also differences in operations and business practices which resulted in issues of "localization", i.e., differences in knowledge among the regions, which had to be factored in the knowledge development processes. Thankfully, foreign language issues were minimized. The standard language for financial information globally is English, and an English-only knowledge base was sufficient for current global deployments.[2]

The Global Project Manager was also responsible for managing the budget and all external consultants. Funding of the project was global, proportional to each region's revenue contributions to the company (e.g., a region contributing 20% of Reuter yearly revenues funded 20% of the project.) Thus, there were contributions required from all parts of the globe, both for financial funding and for case authoring. This approach insured appropriate management attention and buy-in worldwide, and minimized the common "not-invented-here" syndrome. Every region was participating.

The *Global Technical Manager* was responsible for supervising all technical aspects of the project, including:

- developing a single, standard global style guide for the knowledge base
- leading all software development efforts, including customizing the CBR Express authoring environment to support the Reuters style guide and building new utilities to support Reuters global procedures
- supervising technical deliverables of external consultants
- approving case bases for global distribution
- maintaining a central library (repository in Chicago) of global cases
- distributing global case bases and/or updates to the five regions worldwide
- technically supporting the five regions.

The *Area Project Coordinators* were responsible for the overall operation of their local region, including:

- managing the hardware and software infrastructure locally
- authoring knowledge bases (building case bases and supervising the domain authors and case authors)
- training authors and end users
- testing and approving locally-built knowledge bases
- transmitting knowledge bases to the central site in Chicago
- receiving global case bases and updates from Chicago and implementing them locally.

The *Domain Author* is the person responsible for a particular case base. Each case base contains cases relative to a particular product. This segmentation of knowledge based on product seemed the most natural at Reuters. Customer calls focused around the particular Reuter product the customer was using, and problems or issues associated with the product. A particular product case base was assigned for development to the particular region that had the most expertise in the product. The Domain Author is the individual assuming ownership (content, delivery, maintenance) of a product case base.

In some cases the Area Project Coordinator (APC) was also a Domain Author for one or more product domains. These added responsibilities for the APC varied and were based on the APC's workload and domain expertise.

A Domain Author could use multiple *Case Authors* to help author the knowledge in a particular case base. Again this was dependent on workload responsibilities and product expertise. Thus, many configurations were possible: an APC could fully author a small product case base, or could supervise a Domain Author who had several Case Authors to contribute the knowledge. The Domain Author is ultimately responsible for the knowledge content and organization within the assigned case base. That individual would accumulate cases from the Case Authors, look for redundancies, and ensure consistency and style.

This overall organization, while simple to present in terms of responsibilities, was very difficult to establish and manage. A principal issue that must be addressed in any such effort concerns reporting structure. APCs continued to report to their region's management structure, and not directly to the Global Project Manager nor the Global Technical Manager. Domain and Case Authors had phone responsibilities and allocation of their time for the global case-base project was done regionally, and not assigned

[2] Language issues are still a concern at Reuters -- in some areas (e.g., Japan), help desk operations are provided principally in the local language. Other unpublished non-Reuters CBR efforts have addressed the language issues (e.g., where support is provided in the local language only), and have set up translation processes to create multi-language knowledge bases. Issues of maintenance and updates are even more significant in these efforts.

globally. Thus, timetables and milestone dates were continually modified and various regions had peaks and valleys in terms of their productivity and commitments.

Added to this issue were some of the issues mentioned earlier: working with different cultures, with different work ethics, with different approaches to software development, with different business practices, etc. In view of these issues, the need to be realistic in deliverable time-scales and the need to remain flexible in achieving results were two very important lessons to be learned. An initial aggressive plan stalled when authors could not given sufficient time to build global cases; optimistic deadlines to produce working systems resulted in brittle and patchy knowledge bases; giving early incomplete case bases to end users created negative impressions that were hard to subsequently turnaround. Several times over the 18-month period, the project was almost killed due to missed deadlines and poor acceptance. Nevertheless, at each decision point, these issues were evaluated and, each time, some progress occurred suggesting that the end goal could still be achieved. It was always important to continually focus on the original goals and objectives -- and each time they remained valid, and over time they seemed more and more within reach. Perseverance paid off.

CBR Technology

The use of case-based reasoning technology in the customer support arena has now been firmly established and reported on [Acorn & Walden, 1992; Nguyen et al., 1993; Hislop & Pracht, 1994]. The current Reuters project adds to these efforts in confirming the use of CBR as a viable technology to develop a *global* knowledge repository that can readily be built, maintained and reused. In the Reuters case, these CBR knowledge bases have been built in a distributed global environment and deployed in multiple countries. This environment, and its associated requirements, presented a whole new list of challenges and issues. But first we overview the technology used.

Briefly, for those unfamiliar with CBR technology, a CBR knowledge base consists of a set of past *cases* (situations, problems, inquiries) each of which contains a description and various features that define the situation and its uniqueness. Associated with each case is its applicable action or solution, i.e., given the defined situation, it is advisable to suggest the given solution.

Cases are aggregated into a *case base*, which is then used to search against in response to a new situation or problem. When a similar case is retrieved, it then forms the basis of a solution or response. Case bases evolve as new knowledge is entered or as modifications and updates occur to existing cases.

Case-based reasoning has been an active area of artificial intelligence research for over a decade [Riesbeck & Schank, 1989; Kolodner, 1993; Allen, 1994]. In the United States, ARPA sponsored-research in the mid to late 1980s served to establish CBR as an active research discipline [Kolodner, 1988; Hammond, 1989; Bareiss, 1991]. In Europe, a series of workshops helped formulate CBR research directions and highlighted opportunities for applications [Wess et al., 1994; Haton et al., 1995; Watson, 1995]. Worldwide, this has led to CBR's first international conference held in October 1995 [Veloso & Aamodt, 1995].

Applying CBR technology in the customer support help desk environment is a very appropriate use of the technology. A Reuter customer calls in with a problem or issue. The Reuter customer support representative tries to solve the problem with the help of the case base. The representative has a 486 PC (connected via a LAN to a server which stores the cases) and interacts with the case base to solve customer problems and answer various inquiries. The representative enters a description of the problem, and various features of the problem (entered through a question-answering dialog). Through a process of entering information, searching for relevant cases, and answering questions to help narrow the search, a solution is found. (If no solution is found the situation is then a candidate for a new case to be authored.) The representative then provides that solution to the customer.[3]

Reuters used off-the-shelf CBR products from Inference Corporation – CBR Express for building case bases, and CasePoint for deployment. These tools offer a combination of natural language entry and controlled searching. Every case usually has a textual description, much like an abstract of précis of the problem, describing the problem and its symptoms. When a user initiates a search, this system description alone often presents enough suggestions in the resultant set of possible solutions that the user can find the appropriate solution to use.

When the initial natural language search doesn't provide a

[3] There are now many examples of the use of CBR directly by customers for self help, where customers solve their own problems via a case base, for example, locally available on a CD-ROM, or remotely accessible through the Internet. These deployment strategies can have significant impact on reducing costs and increasing customer satisfaction.

useful result, the user can then begin answering the questions accompanying the first set of retrieved solutions. Each time a question is answered, the set of retrieved solutions is refined with more reasonable solutions. If no solution exists for a problem, this gradually becomes clear through the lack of a high-scoring matched case surfacing as a solution.

The rules for writing effective questions is defined by a *style guide*. At Reuters, the question and answer style definitions were devised by a team of people. The team approach was necessitated because of the breadth of the services offered by Reuters and the need to have all case studies end up in a single database at the conclusion of the project. When building questions/answer sets it is important to keep in mind all possible uses of the questions, so as to not build too many questions that might differ only in shades of meaning.

The difficulty in team approach is that many semantic arguments were endured, which could have been alleviated had only one individual defined the initial question/answer set. There was a great and ongoing debate about whether systems messages were error messages or informational messages, or whether all messages should be considered error messages. It was eventually determined that all messages are error messages (much to the dismay of former application programmers and system administrators on the initial authoring team who had a refined sense of computer messages). The point is that if the breadth of the case base can be understood by a single individual, it is probably more effective for that person to design the question/answer index and be given absolute authority to settle meaning disputes as they arise.

The Reuters style guide provided a single uniform template for cases, questions and actions. The GUI front-end to CBR Express was modified to create an authoring environment that enforced the Reuters style guide constraints. For example, all Reuter cases were required to contain the same two first questions "What product are you using?" and "What is the nature of your call?". The customized system, called CBRXtra, enforced this constraint for every new case generated by any author around the world. Another example of a style guide constraint is the requirement to put in a standardized product name as part of a case title (the product associated with the case). This customized interface created a consistent Reuters case template for use in authoring globally.

Given that case authoring was distributed around the world, various manual procedures were initially developed for correctly merging master case bases. However these manual procedures didn't always work correctly, as steps were skipped or not performed in the right order. On a few occasions authors' work was lost when files became corrupted. Subsequently, development began on the *Reuters Smart Suite* of utilities to automate the knowledge merging procedures making it impossible for authors to make case corrupting errors.

During the specification process of these utilities, some concern arose that what we were trying to do might in fact be impossible to achieve on a reliable basis. After a few days of horror that our assumptions were flawed, and through the persevering efforts of an analyst and the development team, all issues were ultimately resolved and it is now possible to write a case anywhere in the world at any time and have it merged into the Reuter Case Base.

The Reuters Smart Suite includes five utilities to support global authoring and distribution:

- *Smart Diverge* determines, for a recently modified case base (e.g., new cases added to a case base, or existing cases modified), those cases, question and actions that are new (and should be appended to the current global case base) and those that are modified/updated (and should replace those in the current global case base).

- *Smart Shred* eliminates unused questions and actions within a case base (locates defined questions and actions that never occur in any case, and discards them or puts them in separate file).

- *Smart Sort* sorts cases, questions and actions (e.g., by product domain), and cross references them.

- *Smart Collate* aggregates unresolved cases (i.e., situations for which a case was not found during search, and thus good candidates for new cases), into appropriate groups to send to particular regions in the world authoring case bases for those unresolves (i.e., determines to which Area Project Coordinator to send the individual unresolves). This utility is critical for the maintenance process to provide the information to the correct resource for knowledge authoring.

- *Architect* creates a tab-separated file showing the entire structure of a case base which can be printed or imported into Excel.

These customizations allowed Reuters to enhance and tune their authoring and distribution processes to support their requirements.

Global Knowledge Management

While the use of CBR on the help desk is no longer a technical or business innovation, developing a CBR knowledge base in a global framework did present some unique challenges that did require innovative solutions. Some of the management and organizational issues have already been discussed above. Here we address some of the issues and challenges of knowledge management and knowledge capture on a global level.

The issues Reuters faced, and we believe need to be addressed in any global knowledge effort, are grouped here into three primary areas: authoring, distribution and localization. While these issues are extensive, they are in no way insurmountable. We have interacted with other companies embarking on global knowledge efforts and each company solves these issues in different ways depending on their requirements, operations and business objectives. What is important is to recognize what the issues are, and the pros and cons of the various alternatives. We list here some of the more generic issues and how Reuters addressed them.

Knowledge Authoring

centralized vs. distributed

Certainly the easiest approach is to centralize the knowledge building process, i.e., in one central site with a group of knowledge engineers co-located in one area to share ideas, approaches, issues, etc. Management of the knowledge process is also simplified. At Reuters, however, expertise is distributed around the world and it was important to leverage all of that expertise into the most accurate and complete knowledge base possible.

segmentation of knowledge bases

Knowledge (case bases) is organized around Reuter products. Each product case base is assigned to (and owned by) a particular region who has responsibility for generating the case base, testing the case base, and maintaining the case base.

versions for different audiences

In developing case bases for global use, one needs to consider the users, most typically customer support telephone representatives – their skill level, mode of operation, and business practices. In other cases, users may be on the help desks of customers, or even customers directly (through case bases distributed to customers on CD-ROM, or available on the Internet). Case-base design needs to consider the end users and accommodate their (sometime conflicting) needs.

global knowledge vs. local knowledge

Ideally global case bases should contain information that is pertinent worldwide. In some cases, however, it may be important to include local information within global cases (see also the discussion in the Localization section below). For example in a 24-hour global support strategy (e.g., a London customer is connected to a London support center during office hours, but is connected to a U.S. support center in the evening, or to an Asia support center in the early morning hours before London is open), it might be important to give the correct local information to the customer (who may be on a different continent).

single global style vs. multiple styles

While Reuters created one global style for all cases worldwide, other efforts may require regional styles due to different business practices and requirements. The latter approach creates issues in translating case bases based on style (e.g., one region may enter detailed descriptions of problems and then focus in on a solution in a couple of confirmation questions; other regions may require leading the representative through a well-defined ordered question and answer dialog).

expert authoring vs. knowledge engineer authoring

Reuters has used both approaches due to workload constraints in various regions. In RAM, cases are authored by domain experts. The number of authors has varied between 6 and 12 authors who work between 4 to 8 hours per week authoring cases (the rest of the time they are on the phone solving customer problems). In UKI, 2 knowledge engineers (one being the UKI Area Project Coordinator) interviewed domain experts and created all the cases. Both approaches have been successful. The decision on which approach to adopt should be based on skill levels and time commitments of the people involved.

on-going maintenance

Case bases are maintained (authoring new cases and updating/modifying existing cases) using the same global organizational structure, i.e., the regions that built the case bases maintain them and optimally use the same domain and case authors to add and update the knowledge base. This process is facilitated by the Smart Collate utility described above to correctly identify the appropriate region to send new unresolved cases to author.

Knowledge Distribution

central library vs. distributed library

Reuters chose a central library to be a single repository for case bases. This seemed the best alternative to more readily manage the distribution of updates. The alternative is to have each region responsible for distributing their own case bases. In this alternative, regions would be getting updates on the varying (currently 38) case bases from multiple regions, necessitating more overhead in incorporating updates locally. With centralized distribution, updates are received from one site (Chicago).

distributing updates vs. whole case bases

Reuters global procedures allow distribution of updates (new or modified cases), and do not require sending whole case bases each time. Distributing updates is particularly recommended in those situations where local regions make local changes to cases (in which case you want to minimize their rework in making those local changes). In addition, communication bandwidth can significantly be reduced. On the other hand, in some cases, simply replacing a whole case base with another can be easiest if there are no local changes and communication bandwidth is not an issue.

frequency of updates

Ideally as soon as there is new knowledge, it should be distributed immediately. However, distributing updates requires testing the new release, packaging cases, transmission, and incorporating updates at the local level. While automating much of this process is desired and achievable, issues in version control and software management need to be considered in deciding on distribution frequency. Currently Reuters distributes updates twice each month.

foreign language translation (if required) before vs. after distribution

Foreign language translation requirements complicates the distribution process. Issues of sending whole case bases vs. updates, and frequency of distribution are all impacted. Translation could occur before distribution (centrally) or after distribution (locally). This depends on the particularly translation process used and who is designated to manage and control it. Cost and time issues can be significant as well, which would limit the turnaround time for issuing updates.

distribution format

Since Reuters has not yet established a worldwide standard for all databases, it chose to distribute cases as text files which are then read into each region's own environment to create a local database. Standardizing on a single database would allow database records to be distributed and thus further minimize the work required at the local level.

delivery mechanism

Various options exist here: ftp, e-mail, diskettes, CD-ROM, Internet. Initially distribution was done via diskettes and mail. This, of course, is cumbersome and slow. Reuters is moving toward electronic communication. Other options are being considered as case bases are being made available to other Reuter organizations and to outside customers.

extent to which the distribution process is automated

No surprise here – automate, automate, automate. The more that can be done computationally without human involvement the better. While Reuters has accomplished some of this through the Smart Suite (e.g., aggregating unresolved cases, determining what cases are new or modified, creating text files for distribution), much more is needed to make the whole process error-free, secure and shorter.

Localization

incorporating country-specific knowledge

Reuters allows individual regions to customize cases for their own use. (As mentioned above there may be good reasons for actually incorporating local knowledge within global cases.) Each region or country can have differing business practices (e.g., determining when to send out a field engineer), safety regulations (e.g., in allowing customers to replace parts or components on their own), legal issues, different cultures, etc. Cases are distributed to local sites as text files making it straightforward to edit cases, questions and actions. Local regions need to keep track of their changes, as additional updates sent from the global master site may impact cases modified for local use. New local cases can also be added regionally.

designing with localization in mind

If there are localization requirements, alternative case-base designs and distribution mechanisms need to be evaluated. For example, whether to distribute text files vs. compiled database records (i.e., CBR Express creates database tables for storing cases) -- it is easier for authors to work directly at the case-base level rather than at the database level for making local modifications. This is at the expense of efficiency (in distributing database records and index files directly).

Issues of standardizing on databases also factors in here. An example impact of localization relative to case-base design would be a design where authors creates global actions but allow local changes only through file attachments (and not through direct changes to the global cases and actions themselves). So a global action may specify that the customer should call another phone number, where the actual phone number is then stored in a local file. The global action is the same (to call in), but the local information is in a file created and maintained locally – the global action has the "hook" that allows the local action to be attached. The whole issue of localization needs to be understood, and requirements for localization need to be defined early on in the design process. Building global case bases and then seeing how local regions want to localize could be disastrous. The more localization required, the more complicated the design. At the extreme, if everything is localized, there is little reason to create global knowledge bases.

foreign language translation

This is a very important issue for every global effort. Reuter cases are all authored, distributed, and used in English (United Kingdom English was selected for worldwide use over American English). Issues of where translation is done (centrally, locally), when (before distribution, after distribution) and how (to what extent automated vs. manual translation, or even possibly authoring cases in multiple languages concurrently) all need to be addressed. Decisions need to be made on who maintains already translated cases, i.e., whether the local country should maintain a case base once translated, or continue to have updates translated and distributed from other sites.

different infrastructures or integration requirements

Again, understanding local requirements (and plans!) is critical – client computers, servers, networking, bandwidth, databases, etc. all need to be factored in for all the sites to be serviced. This infrastructure impacts system design and automation and distribution alternatives.

To note again, these issues all need to be considered in any global knowledge effort. Like any good software development effort, understanding all the requirements upfront is important. To the more common system requirements, we have added the knowledge requirements.

Benefits Achieved

Benefits from the Reuter Global Case Base Project have been significant. In some cases, however, exact measurements have been difficult to obtain. Processes have not been adequately established for accumulating statistics worldwide. Nevertheless, feedback both from existing numerical metrics as well as qualitative evaluations show improvements in numerous areas. A return-on-investment model has been established that focuses on four key benefit areas:

- *supporting first call clearance* (i.e., more calls are being resolved on the first call) -- the knowledge in the case base is providing an expert assistant to the phone agent, and that has enabled the agent to resolve more calls -- this saves escalation costs, repeat calls, and field visit costs.

- *lowered reliance on second-level technical support* saves experts' time in not having to deal with redundant problems and having to call back customers to provide the solutions.

- *reduction in field dispatches* – solving problems on the phone and avoiding field engineers having to visit customer sites saves $400-500 average for each site visit.

- *new hiring training reductions of 33%*; in addition, new hires become more productive more quickly than previously. Turnaround in help desk employees can be significant and this benefit alone justifies the project's cost.

While these benefits are measurable, many other very important benefits are much more difficult to measure. Yet these qualitative benefits often significantly outweigh the tangible benefits listed above:

- capturing and aggregating knowledge that is distributed worldwide into a single knowledge library that can then be distributed to any site around the world

- sharing this knowledge to other organizations besides customer support, e.g., to sales, marketing, and field service divisions

- providing consistency and high quality solutions worldwide

- enabling 24-hour service to any customer worldwide – with the same intelligent response and solution (processes to implement this strategy are currently in the planning phase; they involve automatic call transfer to other countries when local offices are

closed.)

- enhanced customer satisfaction
- retention of customers, customer loyalty and repeat business.

In addition to deploying a global knowledge asset, the process of building the case base has also increased awareness of the discipline of problem solving among case authors. Some of the experts on the project commented that their methods and means of explanation to other less experienced staff improved because of the rigor of the authoring process. What seems intuitive and natural to a support expert when talking to someone directly experiencing a problem, is less natural when explaining to someone else or when trying to record that knowledge for later reuse.

When building cases, authors need to be precise about their meanings and the order of events in the troubleshooting process. Because of this care and attention in the case building process, authors themselves emerge more experienced in improving the efficiency with which they solve problems.

Finally, yet another very important benefit has resulted from the development of case bases – this time within the Reuters product development organization. Case bases are now being built and delivered simultaneous to new product launches. Not only are these case bases being built by the true product experts (those that designed and built the product), but this knowledge is being made available to the customer support organizations in time for them to use the knowledge base on the very first product call! The Reuter Bond Window product is the first Reuter product launched concurrently with a supporting case base. The release of case bases for products should become as commonplace as releasing product documentation.

Summary

The Reuter global case-base project is one of the very first projects to focus on building a knowledge base from expertise existing in many areas around the world. This knowledge is authored in multiple global regions and stored in a central master library. Knowledge is then distributed to multiple Reuter sites worldwide that need and want this knowledge. To support this enterprise, Reuters has established a global organization, global procedures, and supporting software to make this process effective. The project has now transitioned to an on-going maintenance process. Currently the system is in use in 29 sites by 190 users. The user base is expanding, not only into additional Reuter help desks, but also to other Reuter organizations. Product development is now building case bases to launch simultaneous with new product releases. Other plans under consideration include providing the knowledge base directly to customers' own help desks, providing the case base over the Internet, and leveraging the case base to provide 24-hour worldwide service to Reuter customers around the world.

Acknowledgments

There were many contributions from all over the globe. At various times in the project, the Global Steering Committee, which monitored and funded the project, included Phil Arnett (Reuters Corporate), Steve Arthers (Continental Europe, Middle East and Africa), Graeme Barbour (Asia), Marco Bernasconi (Continental Europe), Iian Burgess (Middle East and Africa), Steve Grigg (United Kingdom and Ireland), and Dan Rooney (America). Jim Borron, Chuck Schwartz, and Jack Bucsko were the principal project leaders at various times during the 18-month period. Phil Klahr provided global management and technical support. David Morales was the global technical manager. Area Project Coordinators included Paul Cantini (Continental Europe), Mark Hanson (America), Steven Sim (Asia), Andre Vanderschyff (Middle East and Africa), and Marcus Walia (United Kingdom and Ireland). External consultants included Ron Bewley, Doug Laney, Larry Mond, Steve Porvin, Savita Raj, Samir Rohatgi, and Mike Smith.

References

Acorn, T. L., and Walden, S. 1992. SMART: Support Management Automated Reasoning Technology for Compaq Customer Service. In *Innovative Applications of Artificial Intelligence 4, Proceedings of IAAI-92* (Scott & Klahr, eds.). Menlo Park, California: AAAI Press.

Allen, B. 1994. Case-Based Reasoning: Business Applications. *Communications of the ACM* 37(3):40-42.

Bareiss, E. R. (ed.) 1991. *Proceedings of the 3rd DARPA Case-Based Reasoning Workshop*. San Mateo, California: Morgan Kaufman Publishers.

Hammond, K. J. (ed.) 1989. *Proceedings of the 2nd DARPA Case-Based Reasoning Workshop*. San Mateo, California: Morgan Kaufman Publishers.

Haton, J. P., Keane, M., and Manago, M. (eds.) 1995. *Advances in Case-Based Reasoning, Proceedings of the*

Second European Workshop (EWCBR-94). Berlin: Springer-Verlag.

Hislop, C., and Pracht, D. 1994. Integrated Problem Resolution for Business Communications. In *Proceedings of the Sixth Innovative Applications of Artificial Intelligence Conference* (Byrnes & Aikins, eds.). Menlo Park, California: AAAI Press.

Kolodner, J. K. (ed.) 1988. *Proceedings of the 1st DARPA Case-Based Reasoning Workshop.* San Mateo, California: Morgan Kaufman Publishers.

Kolodner, J. 1993. *Case-Based Reasoning.* San Mateo, California: Morgan Kaufman Publishers.

Nguyen, T., Czerwinski, M., and Lee, D. 1993. Compaq QuickSource: Providing the Consumer with the Power of Artificial Intelligence. In *Proceedings of the Fifth Innovative Applications of Artificial Intelligence* (Klahr & Byrnes, eds.). Menlo Park, California: AAAI Press.

Riesbeck, C. K., and Schank, R. C. 1989. *Inside Case-Based Reasoning.* Hillsdale, New Jersey: Lawrence Erlbaum Associates.

Veloso, M., and Aamodt, A. (eds.) 1995. *Case-Based Reasoning Research and Development, Proceedings of the First International Conference on CBR (ICCBR-95).* Berlin: Springer-Verlag.

Watson, I. D. (ed.) 1995. *Progress in Case-Based Reasoning, Proceedings of the First United Kingdom Workshop.* Berlin: Springer-Verlag.

Wess, S., Althoff, K., and Richter, M. (eds.) 1994. *Topics in Case-Based Reasoning, Proceedings of the First European Workshop (EWCBR-93).* Berlin: Springer-Verlag.