

Developing and Deploying Knowledge on a Global Scale

James Borron, David Morales, and Philip Klahr

■ Reuters is a worldwide company focused on supplying financial and news information to its more than 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 North America, Europe, and Asia. The system supports 38 Reuter products worldwide. This article presents a case study of Reuter experience in putting a global knowledge organization in place, building knowledge bases at multiple distributed sites, deploying these knowledge bases in multiple sites around the world, and maintaining and enhancing knowledge bases within a global organizational framework. This project is the first to address issues in multicountry knowledge development and maintenance and multicountry knowledge deployment. These issues are critical for global companies to understand, address, and resolve to effectively gain the benefits of global knowledge systems.

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 numeric and textual historical databases, news, graphics, still photos, and news video. Reuters has 40,000 subscribers and 309,000 user accesses and operates in 154 countries. Reuter 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.

Although Customer-Service Operations have been in place at 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: (1) the data and news content Reuters provides, (2) the technology with which these data are delivered; and (3) 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 that 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 reengineering 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, and development of a strategy to preserve knowledge assets.

The strategy involved building knowledge

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.

systems using case-based reasoning (CBR) (Kolodner 1993). Other customer-support organizations had already reported successes in using CBR technology (as subsequently reported at IAAI conferences [Hislop and Pracht 1994; Nguyen, Czerwinski, and Lee 1993; Acorn and Walden 1992]), 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, authoring a style guide, agreeing on product domains to cover, committing time for case authors (people building the case bases), developing an incentive program, and setting and achieving targets.

In December 1993, RAM deployed its first case base. It had approximately 1200 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 ongoing basis and an expert adviser 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.

However, there were 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 1200 cases would be insufficient. Also, 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 to generating complete and accurate case bases.

While RAM pursued CBR technology aggressively in 1993, other areas at 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 provide guidance on using and integrating CBR technology for problem resolution. Some initial case bases were built at UKI in the 1992 to 1993 time frame that also

helped RAM in its evaluation of CBR in early 1993. A third effort in CBR at Reuters was also simultaneously occurring at the Reuters Middle East and Africa (MEA) organization. In particular, the South Africa help desk quickly and independently developed case bases to support its 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. Its initial objectives were to (1) consolidate the case base projects in the United States, United Kingdom, and Africa; (2) increase domain coverage; (3) more effectively use the best resources for authoring worldwide; (4) establish a global funding mechanism; and (5) 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 given 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 downtown Chicago or 5000 miles! Knowledge can easily be captured wherever it resides, and this conclusion 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 6 months long:

Phase 1 involved initiating and organizing the project, formulating the budget, establishing global procedures, developing a global style guide, generating the initial knowledge base (combining and globalizing knowledge from three existing efforts), and creating prototypes of the initial supporting software utilities.

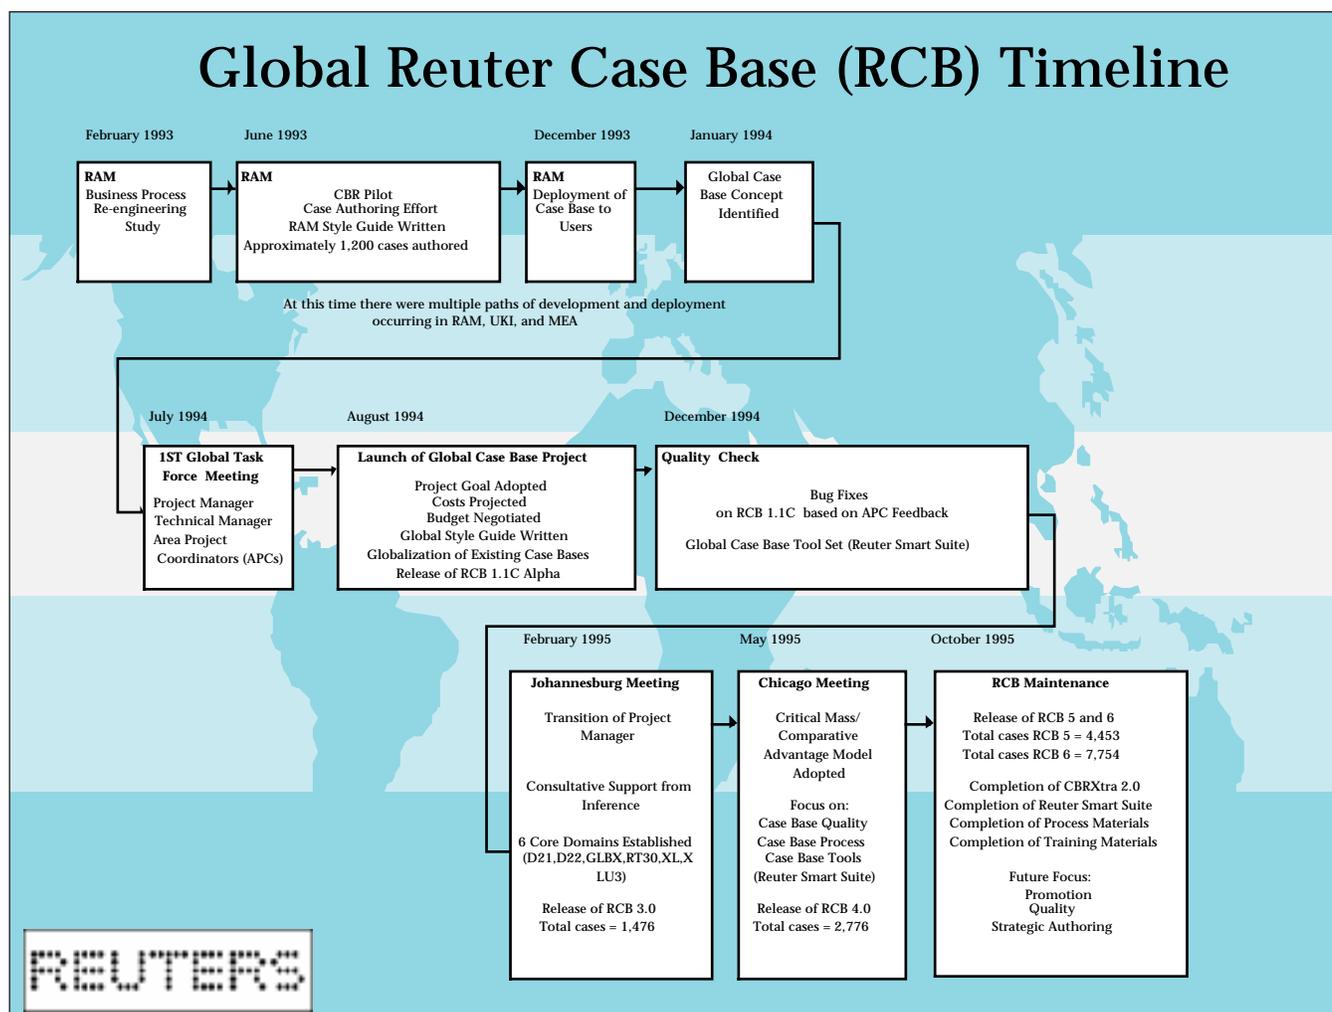


Figure 1. Global Reuter Case Base Project Time Line.

Phase 2 involved completing the first version of software utilities to support global procedures, solidifying quality assurance procedures, expanding the knowledge-building process (here focused on increasing the number of products-domains covered), providing initial user training, deploying and beta testing initial versions of the knowledge base, and computing return-on-investment metrics.

Phase 3 involved further expanding the knowledge building to a critical mass (now focused on increasing the depth of knowledge for products-domains covered), completing the second version of the software utilities, quality testing the complete system, completing the documentation (utilities, global procedures, training materials), and completing global rollout with a more robust knowledge base. Figure 1 is an overview of the project time line.

The 18-month effort is now reaching its

completion. It has achieved all its objectives. As of November 1995, the Reuter case base contains 7754 cases, covers 38 Reuter products and services (figure 2), is installed at 29 sites around the world (figure 2), and 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: (1) RAM, (2) UKI, (3) Reuters Asia (RA) and Reuters Continental Europe (RCE), (4) MEA (RCE and MEA have since been combined), and (5) Reuters corporate. The Global Steering Committee meets quarterly to review project milestones, issues, and directions. Bi-weekly conference calls provide interim sta-

The Reuter Case Base

The Reuter Case Base is a global knowledge base, created by Reuter staff, that grows with user interaction. It acts as a tool for Help Desk personnel in solving customer queries quickly, putting expertise at their fingertips.



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

tus updates and address immediate issues and needs.

The principal global organization structure includes five logically defined job-skill positions: (1) global project manager, (2) global technical manager, (3) area project coordinators, (4) domain authors, and (5) case authors.

The *global project manager* led the project and coordinated the development and deployment efforts in the five regions. This challenge was certainly one of the more difficult of the project. Not only are there significant time-zone differences among the five regions but also differences in culture, software development practices, structure and management organization, and even project objectives. In addition, from a knowledge perspective, there were also differences in operations and business practices that resulted in issues of *localization*, that is, differences in knowledge among the regions, that had to be

factored into the knowledge development processes. Thankfully, foreign-language issues were minimized. Globally, the standard language for financial information 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 (for example, a region contributing 20 percent of Reuter yearly revenues funded 20 percent of the project). Thus, contributions were required from all parts of the globe for both financial funding and case authoring. This approach ensured appropriate management attention and buy in worldwide and minimized the common "not-invented-here" syndrome. Every region was participating.

The *global technical manager* was responsi-

ble for supervising all technical aspects of the project, including (1) developing a single, standard global style guide for the knowledge base; (2) leading all software development efforts, including customizing the CBR EXPRESS authoring environment to support the Reuter style guide and building new utilities to support Reuter global procedures; (3) supervising technical deliverables of external consultants; (4) approving case bases for global distribution; (5) maintaining a central library (repository in Chicago) of global cases; (6) distributing global case bases or updates to the five regions worldwide; and (7) technically supporting the five regions.

The *area project coordinators* were responsible for the overall operation of their local region, including (1) managing the hardware and software infrastructure locally, (2) authoring knowledge bases (building case bases and supervising the domain authors and case authors), (3) training authors and end users, (4) testing and approving locally built knowledge bases, (5) transmitting knowledge bases to the central site in Chicago, and (6) receiving global case bases and updates from Chicago and implementing them locally.

The *domain author* was responsible for a particular case base. Each case base contains cases relative to a particular product. This segmentation of knowledge based on products seemed the most natural at Reuters. Customer calls focused on the particular Reuter product that 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 assumed 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 work load and domain expertise.

A domain author could use multiple *case authors* to help author the knowledge in a particular case base. Again, this method was dependent on work-load 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, although simple to present in terms of responsibilities, was 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 or 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 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, different work ethics, different approaches to software development, different business practices, and so on. 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 important lessons to be learned. An initial aggressive plan stalled when authors were not given sufficient time to build global cases; optimistic deadlines to produce working systems resulted in brittle and patchy knowledge bases; and giving early incomplete case bases to end users created negative impressions that were hard to subsequently turn around. Several times over the 18-month period, the project was almost killed because of 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.

Case-Based Reasoning Technology

The use of CBR technology in the customer-support arena has now been established firmly and reported on (Hislop and Pracht 1994; Nguyen, Czerwinski, and Lee 1993; Acorn and Walden 1992). The current Reuter 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 case of Reuters, these CBR knowledge bases have been built in a distributed global environment and deployed in multiple countries.

The current Reuter 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.

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; that is, 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.

CBR has been an active area of AI research for over a decade (Allen 1994; Kolodner 1993; Riesbeck and Schank 1989). In the United States, DARPA-sponsored research in the mid-to late-1980s established CBR as an active research discipline (Bareiss 1991; Hammond 1989; Kolodner 1988). In Europe, a series of workshops helped formulate CBR research directions and highlighted opportunities for applications (Watson 1995; Haton, Keane, and Manago 1995; Wess, Althoff, and Richter 1994). Worldwide, this work led to CBR's first international conference in October 1995 (Veloso and Aamodt 1995).

Applying CBR technology in the customer-support help-desk environment is an appropriate use of the technology. A Reuter customer calls 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 by a local area network to a server that 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-answer 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 the solution to the customer.³

Reuters used off-the-shelf CBR products from Inference Corporation—CBR EXPRESS for building case bases and CASE POINT for deployment. These tools offer a combination of nat-

ural 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 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, it gradually becomes clear when a high-scoring matched case doesn't surface 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 in a single database at the end of the project. When building question-answer sets, it is important to keep in mind all possible uses of the questions to prevent building too many questions that might only differ in shades of meaning.

The difficulty with the 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 system 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 Reuter style guide provided a single uniform template for cases, questions, and actions. The graphic user interface front end to CBR EXPRESS was modified to create an authoring environment that enforced the Reuter 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 Reuter 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 because steps were skipped or performed in the wrong order. On a few occasions, authors' work was lost when files became corrupted. Subsequently, development began on the Reuter 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 Reuter SMART SUITE includes five utilities to support global authoring and distribution:

First, SMART DIVERGE determines, for a recently modified case base (for example, new cases added to a case base or existing cases modified), those cases, questions, and actions that are new (and should be appended to the current global case base) and those that are modified or updated (and should replace those in the current global case base).

Second, 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 a separate file).

Third, SMART SORT sorts cases, questions, and actions (for example, by product domain) and cross-references them.

Fourth, SMART COLLATE aggregates *unresolved cases* (that is, 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 these unresolves (that is, 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.

Fifth, *architect* creates a tab-separated file showing the entire structure of a case base that 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

Although 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. 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: (1) authoring, (2) distribution, and (3) localization. Although 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. It is important to recognize what the issues are and what the pros and cons are of the various alternatives. We list here some of the more generic issues and how Reuters addressed them.

Knowledge Authoring

The knowledge-authoring issues are (1) a centralized versus a distributed knowledge process, (2) the segmentation of knowledge bases, (3) case bases for different users, (4) global versus local knowledge, (5) single global style versus multiple styles, (6) expert authoring versus knowledge engineer authoring, and (7) ongoing maintenance.

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

The issues Reuters faced, and we believe need to be addressed in any global knowledge effort, are authoring, distribution, and localization.

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 that has responsibility for generating, testing, 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 can be on the help desks of customers or even customers directly (through case bases distributed to customers on CD-ROM or through the Internet). Case base design needs to consider the end users and accommodate their (sometime conflicting) needs.

Global knowledge versus local knowledge: Ideally, global case bases should contain information that is pertinent worldwide. In some cases, however, it can be important to include local information within global cases (see also the discussion under Localization). For example, in a 24-hour global support strategy (for example, 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 an Asian 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 might be on a different continent).

Single global style versus multiple styles: Although Reuters created one global style for all cases worldwide, other efforts might require regional styles because of different business practices and requirements. The regional approach creates issues in translating case bases based on style (for example, one region might enter detailed descriptions of problems and then focus on a solution in a couple of confirmation questions; other regions might require leading the representative through a well-defined, ordered question-and-answer dialog).

Expert authoring versus knowledge engineer authoring: Reuters has used both approaches because of work load constraints in various regions. In RAM, cases are authored by domain experts. The number of authors has varied from 6 to 12 who work between 4 and 8 hours a week authoring cases (the rest of the time they are on the phone solving customer problems). In UKI, two 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.

Ongoing maintenance: Case bases are maintained (authoring new cases and updating-modifying existing cases) using the same global organizational structure; that is, 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 earlier to correctly identify the appropriate region to send new unresolved cases to author.

Knowledge Distribution

The knowledge-distribution issues are (1) central library versus distributed library, (2) updates versus whole cases, (3) the frequency of updates, (4) foreign-language translation before or after distribution, (5) the delivery mechanism, and (6) the extent to which the process is automated.

Central library versus distributed library: Reuters chose a central library to be a single repository for case bases. This alternative seemed the best to more readily manage the distribution of updates. The alternative is to have each region responsible for distributing its 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).

Distribution of updates versus whole case bases: Reuter global procedures allow the 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 these local changes). In addition, communication bandwidth can significantly be reduced. However, 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, transmitting, and incorporating updates at the local level. Although 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 versus after distribution: Foreign-language-translation requirements complicate the distribution process. Issues of whole case bases versus updates and the frequency of distribution are all affected. Translation could occur before distribution (centrally) or after distribution (locally). This order of events depends on what translation process is 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: Because Reuters has not yet established a worldwide standard for all databases, it chose to distribute cases as text files that 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 by diskettes and mail, which, of course, are cumbersome and slow processes. Reuters is moving toward electronic communication. Other options are being considered because case bases are being made available to other Reuter organizations and 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. Although Reuters has accomplished some automation through the SMART SUITE (for example, 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

Localization issues are (1) the incorporation of country-specific knowledge, (2) designs, (3) foreign-language translation, and (4) different infrastructures or integration requirements.

Incorporation of country-specific knowledge: Reuters allows individual regions to customize cases for their own use. (As mentioned earlier, there can be good reasons for actually incorporating local knowledge within global cases.) Each region or country can have differing business practices (for example, determining when to send a

field engineer), safety regulations (for example, allowing customers to replace parts or components on their own), legal issues, cultures, and so on. 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 because additional updates sent from the global master site can affect 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, in the case of text files versus compiled database records (that is, 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. Working at the case base level, rather than the database level, creates inefficiencies in distributing database records and index files directly. Issues of standardizing on databases also factor in here. An example effect of localization relative to case base design would be a design where authors create global actions but allow local changes only through file attachments (and not through direct changes to the global cases and actions themselves). Thus, a global action can 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 in the design process. Seeing how local regions want to localize can be disastrous. The more localization that is required, the more complicated the design is. At the extreme, if everything is localized, there is little reason to create global knowledge bases.

Foreign-language translation: This issue is important 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 (automated versus manual translation or the possible authoring of cases in multiple languages concurrently) all need to be addressed. Decisions need to be made on who maintains already translated

Not only are these case bases being built by the true product experts, but this knowledge is being made available to the customer-support organizations in time for them to use the knowledge base on the first product call!

cases, that is, 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, and so on, all need to be factored in for all the sites to be serviced. This infrastructure affects 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 up front 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 adequately been established for accumulating statistics worldwide. Nevertheless, feedback both from existing numeric metrics and qualitative evaluations shows improvements in numerous areas. A return-on-investment model has been established that focuses on four key benefit areas:

First is the support of first-call clearance (that is, more calls are being resolved on the first call): The knowledge in the case base is providing an expert assistant to the phone agent, which has enabled the agent to resolve more calls, saving escalation costs, repeat calls, and field-visit costs.

Second is lowered reliance on second-level technical support. Experts save time not having to deal with redundant problems and call back customers to provide the solutions.

Third is a reduction in field dispatches. Solving problems on the phone and not having to send field engineers to visit customer sites saves on average of \$400 to \$500 for each site visit.

Fourth is new-hire training reductions of 33 percent. New hires become more productive more quickly now. Turnaround in helpdesk employees can be significant, so this benefit alone justifies the project's cost.

Although these benefits are measurable, many other important benefits are more difficult to measure. However, these qualitative benefits often significantly outweigh the tangible benefits listed previously: (1) capturing

and aggregating knowledge that is distributed worldwide into a single knowledge library that can then be distributed to any site around the world; (2) sharing this knowledge to other organizations besides customer support, for example, to sales, marketing, and field-service divisions; (3) providing consistency and high-quality solutions worldwide; (4) 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); (5) enhancing customer satisfaction; and (6) retaining customers and achieving 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 who is directly experiencing a problem is less natural when explaining to someone else or trying to record this 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 important benefit has resulted from the development of case bases, this time within the Reuter product development organization. Case bases are now being built and delivered simultaneously with new product launches. Not only are these case bases being built by the true product experts (those who 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 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 the release of product documentation.

Summary

The Reuter Global Case Base Project is one of

the 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 ongoing 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 simultaneously 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 (United States). 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 (United States), Steven Sim (Asia), Andre Vanderschiff (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.

Notes

1. We use the term *author* to describe the knowledge-engineering-knowledge acquisition process. Authoring is the customary term used in the case-based reasoning customer-support community. It provides a more concise, active description of the entire process (acquiring the knowledge and encoding it). Later, we discuss the alternatives of having an expert author cases directly or having a knowl-

edge engineer interview experts and then encode the knowledge.

2. Language issues are still a concern at Reuters; in some areas (for example, Japan), help-desk operations are provided principally in the local language. Other unpublished non-Reuters case-based reasoning efforts have addressed the language issues (for example, where support is provided in the local language only) and have set up translation processes to create multilanguage knowledge bases. Issues of maintenance and update are even more significant in these efforts.

3. There are now many examples of the use of CBR directly by customers for self help, where customers solve their own problems using 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.

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*, eds. C. Scott and P. Klahr, 3-18. Menlo Park, Calif.: 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 Third DARPA Case-Based Reasoning Workshop*. San Francisco, Calif.: Morgan Kaufmann.
- Hammond, K. J., ed. 1989. *Proceedings of the Second DARPA Case-Based Reasoning Workshop*. San Francisco, Calif.: Morgan Kaufmann.
- 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*, eds. E. Byrnes and J. Aikins, 63-74. Menlo Park, Calif.: AAAI Press.
- Kolodner, J. 1993. *Case-Based Reasoning*. San Mateo, Calif.: Morgan Kaufmann.
- Kolodner, J. K., ed. 1988. *Proceedings of the First DARPA Case-Based Reasoning Workshop*. San Mateo, Calif.: Morgan Kaufmann.
- Nguyen, T.; Czerwinski, M.; and Lee, D. 1993. Compaq QUICK SOURCE: Providing the Consumer with the Power of Artificial Intelligence. In *Proceedings of the Fifth Innovative Applications of Artificial Intelligence*, eds. P. Klahr and E. Byrnes, 142-151. Menlo Park, Calif.: AAAI Press.
- Riesbeck, C. K., and Schank, R. C. 1989. *Inside Case-Based Reasoning*. Hillsdale, N.J.: Lawrence Erlbaum.
- 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.

IDA 97

Intelligent Data Analysis

Program Chairs:

Paul Cohen
 Department of Computer Science
 Lederle Graduate Research Center
 University of Massachusetts, Amherst
 Amherst, MA 01003-4610, USA

and

Xiaohui Liu
 Department of Computer Science
 Birkbeck College, Malet Street
 London WC1E 7HX, UK

IDA Steering Committee:

M. Berthold, Univ. of Karlsruhe, GER
 P. Cohen, Univ. of MA, Amherst, USA
 F. Famili, National Research Council, CA
 D. Fisher, Vanderbilt University, USA
 A. Gammernan, Royal Holloway London
 D. Hand, Open University, UK
 W. Hsu, AT&T Consumer Lab, USA
 X. Liu, Birkbeck College, UK
 D. Pregibon, AT&T Research, USA
 E. Simoudis, IBM Almaden Research, USA

For enquiries:

IDA-97 Administrator
 Department of Computer Science
 Birkbeck College, Malet Street
 London WC1E 7HX, UK
 E-mail: ida97-enquiry@dcs.bbk.ac.uk
 Tel: (+44) 171 631 6722
 Fax: (+44) 171 631 6727

Call for Papers 2nd International Symposium on Intelligent Data Analysis

Birkbeck College, University of London
 4th-6th August 1997

in cooperation with

AAAI, ACM SIGART, BCS SGES, IEEE SMC, and SSAISB

<http://web.dcs.bbk.ac.uk/ida97.html>

The focus of IDA-97 will be "Reasoning About Data". We are interested in intelligent systems that reason about how to analyze data, perhaps as human analysts do. IDA-97 will be a single-track meeting featuring a careful selection of high-quality papers. All submissions will be reviewed by at least two referees of the international program committee on the basis of relevance, originality, significance, soundness, and clarity. Of particular interest to IDA-97 are papers concerning algorithms, applications, general principles, theory, and tools. A detailed list of **topics of interest** as well as guidelines for **paper submission** can be found on the IDA-97 web page.

Submission Deadline: February, 1st, 97.

Publication: Papers which are accepted and presented at the conference will appear in the IDA-97 proceedings, to be published by Springer in its *Lecture Notes in Computer Science* series. Authors of the best papers will be invited to extend their papers for a second round of review for a special issue of "Intelligent Data Analysis: An International Journal".

For **additional information**, please refer to our web page.

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.

James Borron, MLIS, is the director of Customer-Service Systems at Reuters America. He has 10 years' experience designing, developing, implementing, and supporting online information systems for corporate users and customers at Reuters, *The Financial Post*, and libraries in Canada.



David Morales graduated from Concordia University with a bachelor of arts in communications in May 1992. Directly out of college, he pursued a career with Reuters America, where he worked his way up from a telephone-support position, then moved on to the

Global Case Base Project, and now manages a staff of 11 for Reuters Customer-Support Help Desk in the Americas.



Philip Klahr is vice president of customer quality at Inference Corporation, where he has spent the last 10 years. He is responsible for ensuring customer success in building knowledge-based systems, including advising on technology, management, processes, and methodology. Prior to Inference, he was director of Rand Corporation's Information-Processing Systems Research Program. He has published more than 50 articles on AI and knowledge-based systems, particularly in the use of the technology in applications. He has received six Innovative Applications of AI awards and, in 1994, became a fellow of the American Association for Artificial Intelligence for lifetime contributions in applying AI to real business problems.