

# Network Map: Applying Knowledge to the Strategic Selling Process

Jeffrey D. Kenyon and Peter L. Briggs

e.coe (E-Business Center of Excellence)  
U S WEST Information Technologies, Inc.  
1801 California Street, Suite 1560  
Denver, Colorado 80202-1984  
[jkenyon|pbriggs]@uswest.com

## Abstract

The Network Map application leverages an existing corporate data repository to provide a synoptic view of the telecommunications services provided by U S WEST to a large business client. The phased architecture uses a server to generate data portfolios (consisting of HTML and hyperbolic tree data files) on a weekly basis; at the end user's convenience, the portfolios are downloaded on to their local computer, where they may be viewed through a standard WWW browser. By providing these data as part of the standard sales support given to all account representatives, we reduce the sales cycle by eliminating time formerly required to manually assemble a synoptic view, and allow the account representative to focus upon building a consultative relationship with the customer.

## Introduction

This paper describes the Network Map application, a support tool for account representatives engaged in selling telecommunications products and services to U S WEST's largest business customers. The application emerged from a desire to re-engineer the related business process; the account representatives rarely had the data they needed on a client's current voice and data services, and the sales cycle was elongated by the time required to pull various custom reports, assemble them into a synoptic view, and analyze that view to identify opportunities.

The Network Map application takes a phased approach to the above problem. The first phase targeted knowledge management; using an existing corporate repository as the data source, it assembled all data on the client's voice and data services, and created graphical and textual displays that are easily viewed on the account representative's laptop computer.

The second phase (currently in development) uses the data assembled in the first phase and a rule-based system to perform an analysis of the client's services and identify the high-value sales opportunities. The third phase will identify and integrate additional data sources, to improve both the utility of the display, and the quality of the

opportunities generated.

Anecdotal evidence suggests that the portions of the application already delivered, by eliminating data retrieval and assembly time, reduce the sales cycle by weeks, significantly improves the quality of our customer presentations, and allows sales professionals to enact coherent sales strategies.

## Problem Description

U S WEST is a Regional Bell Operating Company (RBOC) formed at the breakup of the Bell system in 1984. Currently U S WEST operates in a 14-state region in the western, southwestern and central United States, serving approximately 15 million customers. Its annual revenue is approximately \$13 billion. One of its market units is Business and Government Solutions (BGS), whose account representatives, account managers and sales managers focus on serving large business, local and state government, and federal clients that typically have many offices throughout the U S WEST territory. BGS account personnel have long-standing, high-level familiarity with their customer's businesses and accounts. However, one large client may have scores of locations and thousands of phone numbers, each with a plurality of products and features. Forming the desired consultative relationship with clients to help them modernize, optimize, and intelligently manage their telecommunications services requires a detailed yet synoptic view of the client. Synthesizing this view is a daunting task in the current environment.

BGS Account Representatives generally spend their time in the field, selling products and services, rather than working out of an office. To support the computing needs of this highly mobile population, they were given laptops; unfortunately, the data about their clients and about the U S WEST network facilities supporting those clients reside on the U S WEST network, and accessing the data requires a network connection. In addition to access, retrieving the relevant data also requires some technical skill in accessing the data (which resides in so-called "legacy" systems, a subset of which is available through a corporate data repository); such skill is in short supply among the mostly non-technical salespeople. Once data is

retrieved, the only way to obtain a graphical representation of the data is to build it yourself, in a tool like Microsoft PowerPoint, using the retrieved textual or tabular data as a starting point. We may summarize this as business problem #1: ***There are presently no tools available to the Account Representatives that provide a concise, usable summary of the client's voice and data services.***

A second major obstacle lies in the application of the data, once it is retrieved. While experienced sales people have developed tacit business rules that they apply to strategically selling new products and services to clients (e.g., if the client currently is paying for a large number of additional lines at residential locations, to support telecommuting, present the benefits of and upgrade path for xDSL service), these rules are not known to newer salespeople, or they may not be able to recognize the relevant patterns in the retrieved data. This is business problem #2: ***There are presently no tools available to the Account Representatives to perform preliminary analysis of the client's current products and services, and identify strategic sales opportunities.***

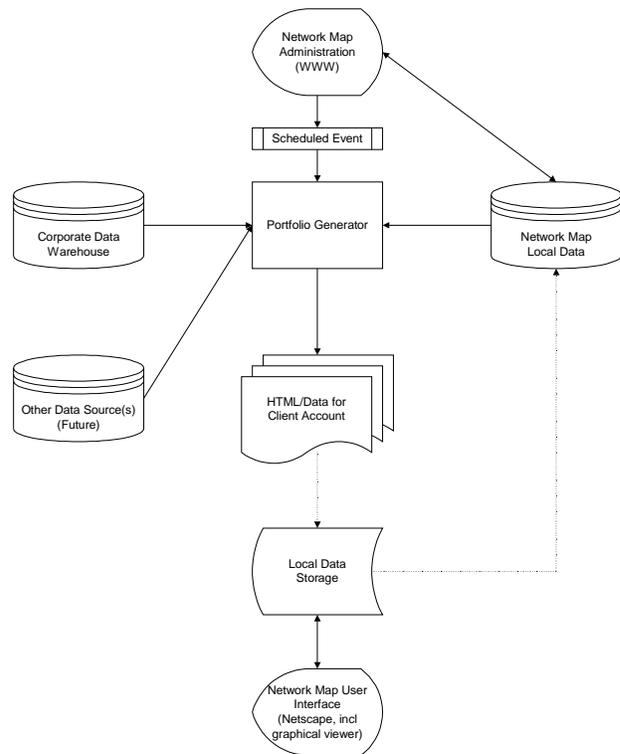
A third major obstacle is the missing connection between the products and services that USW sells, and the available USW network facilities at a specific location that would be required to sell a specific product or service. The result is that many sales proposals must first be examined by the System Design Center (and sometimes Tactical Planning), to determine whether the facilities and spare capacity exist to sell the service. This is business problem #3: ***There are presently no tools available to the Account Representatives that provide even a preliminary answer to the question "Is there inventory available so that I can sell this product or service to the client?"***

### Application Description

The application has been structured in a manner that allows a phased deployment, in which each delivered phase provides business value. The basic architecture is a simple client-server model (Figure 1).

On the server side, an account representative support person sets up accounts, corresponding to one or more customer IDs in the corporate data repository. One or more accounts is then gathered into an account module, and an account representative designated for the module.

At weekly intervals, the portfolio generator runs. This program, written in object-oriented Perl, iterates through the account modules and, for each account, does a series of queries on a data subset extracted from the corporate data



**Figure 1. System Architecture**

repository (it should be noted that the corporate data repository contain several terabytes of data; our application uses a subset of approximately 4Gb of data relevant to the targeted customer base). These queries are used to assemble a set of HTML and hierarchical (TRE) data files (Figure 2), collectively referred to as a *portfolio*. The completed portfolio is compressed and placed in a directory whence it may be downloaded onto the Account Representative's desktop or laptop computer at any time. Downloading and decompression of the archive is accomplished by double-clicking an icon on the Windows 95/NT desktop, and entering a user name and password at the appropriate login screen.

Once the Account Representative downloads the portfolio, the data may be accessed locally at any time via a WWW browser (at U S WEST, this is exclusively Netscape), without further need for a network connection. They can work with these data to analyze the client's array of products and services in the office or the field.

```
A Very Large Company Indeed, Inc.;AVLCI;;;
Arizona ($6841.07);AZ;;;0
PHOENIX ($994.25);PHOENIX;;;0
  a1063_det_AZ/PHOENIX/2820228TH.html;28202 28TH ($174.76) ~SWITCHNET 56 INTRASTATE...;5;;;0
  a1063_det_AZ/PHOENIX/1000CAMELBACK.html;1000 CAMELBACK ($819.49) ~PBX TRUNKS FLAT ACCESS...;5;;;0
    19.IBSA.120541..MS (S);19.IBSA.120541..MS (S);2;;0
TUCSON ($28832.93);TUCSON;;;0
[etc.]
```

**Figure 2. TRE Data File**

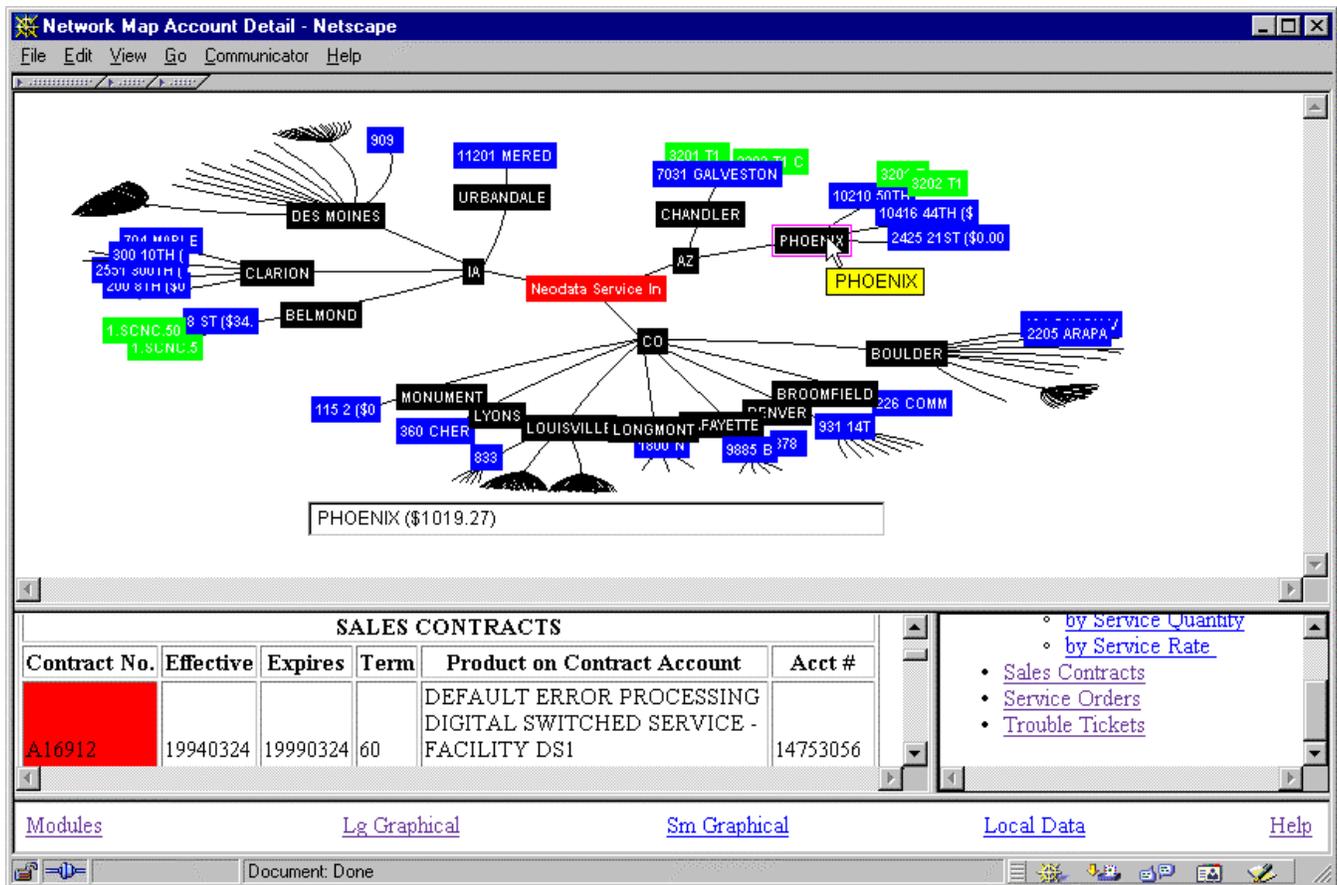


Figure 3. User Interface

**User Interface.** The primary user interface is a WWW browser. Network Map utilizes a frame-based display to present an integrated graphical and textual view (Figure 3). The graphical view is implemented as an active hyperbolic tree, displaying a geographical hierarchy. The layout is alphabetical within each layer, moving clockwise around the radial display. The hierarchy has the following layers, from coarse-grained to fine-grained:

- Central customer-name node
- States (maximum of 15 nodes; 14 states, and “Not Applicable”). Regulations that govern telecommunications are established by state Public Utility Commissions. These regulations, such as the products that may be sold, prices, etc., generally differ in each state. For this reason, this layer of the hierarchy makes a natural grouping both from a regulatory and geographic perspective.
- Cities
- Service Addresses
- Circuit IDs.

When a client’s tree is rotated, the distances between nodes change so that the current area of focus appears relatively uncluttered. This is in contrast to a conventional latitude/longitude map where many rural areas might have

little data while metropolitan areas might be so densely populated with data as to be unintelligible. Also, while it is possible to zoom in on details of a conventional map, the surrounding context is often lost. In the hyperbolic tree display, the context of surrounding parent and child nodes around a node of interest are almost always visible, and hence, the chance of cognitively “getting lost” in the data is reduced. Additionally, the use of color to denote the levels of the hierarchy also helps users keep track of where they are.

The tree itself has two purposes. From a navigational perspective, the tree is used to locate the node for a service address of interest. When that node is double-clicked, the detailed textual summary appears in another pane of the display. From an information display perspective, each state and city node, when highlighted with the mouse, indicates the cumulative monthly recurring charges rolled up to that level. Information “hidden” on each node also facilitates the high-level analysis of the client’s products and services, through the use of the search/spotlight facility.

This facility helps users zero in on specific nodes, particularly useful in large trees of 10,000 or so nodes. An obvious example is for locating a specific service address (e.g., searching on “arap” quickly brings spotlights all

nodes containing that string, such as Arapahoe Street in Boulder, Colorado). A more interesting example would be to search on a specific class of service (e.g., DS1) to get an immediate sense of how embedded a service offering is for a client, and identify opportunities for “upselling.” The location of interesting peripheral nodes can be highlighted even when nodes in that ply are too numerous to resolve; knowing the location, one can then navigate through the display quickly to that specific node.

The applications of the tool to research and to client presentation are not yet fully established. However, early feedback indicates that the tool does not impose an analytical style upon a user; rather it supports any number of approaches to analyzing a customer’s account.

### **Application Development and Deployment**

The first phase of the application (which we would characterize as supporting business process management through knowledge management) went into production at the end of September 1999, and was rolled out to the account representative community in October and November. In parallel with the rollout, development of an *opportunity generator* (a rule-based expert system to explicitly identify sales opportunities in the phase 1 information display) began. The opportunity generator is expected to be integrated as a step within the portfolio generation process, and the output would be an HTML file added to the resulting archive. A prototype focusing on the identification of opportunities to sell the 10-15 products that the corporation has deemed of the highest strategic value is scheduled for delivery in April, 2000.

Original development of the prototype was done in Perl (Wall, Christiansen, and Schwartz, 1996). To keep development costs low, we decided to optimize and deploy the Perl version into production, rather than re-coding in another language. This option will be reconsidered, should performance become an issue.

Development of the rule-based expert system is underway using OPS/R2, the current generation of the OPS programming language (Brownston, Kant, Farrell, and Martin, 1985). Prior to beginning development on this component, we conducted a literature review on the expert systems with the longest histories and published articles, to extract principles and best practices that could guide our development. Specifically, this meant reviewing MYCIN (Buchanan and Shortliffe, 1984), and R1/XCON (Kraft, 1984; McDermott, 1982; McDermott, 1993; Soloway, Bachant, and Jensen, 1987).

### **Discussion**

**Replication.** We believe that the development path used in Network Map may easily be replicated to move data from within a repository to one or more points within the corporation where it may be used by knowledge workers to create value. This belief is based upon efforts already made within our development group on additional prototypes (although none of these prototypes have yet

developed into funded projects). The development steps utilized are:

1. Identify the business problem. As part of this, we identify the nature of the audience with the problem, level of technical expertise, computing environment, etc., since this may limit the accessible data elements, or potential solutions. We feel that it is also important to investigate whether other constituencies within the company have the same, or a similar, problem.
2. Identify (existing) data elements that must be utilized in solving the problem, sources (if any), and exact business semantics for these data.
3. Design the user interface. In Network Map, as well as our other prototypes, our preliminary efforts focus upon a data visualization environment, to give the prospective user appropriate illumination upon the problem.
4. Test with selected users.
5. Release.
6. Repeat. Thus far, we have received considerable feedback and suggestions for additional data that would be useful. All of the feedback requires returning to previous steps; some has even caused us to return to Step 1, re-thinking the original problem, and considering a broadening of the scope of the application.

We should point however that the identification of significant knowledge management opportunities is a non-trivial issue, especially since corporate IT is largely focused upon achieving incremental gains in process improvement, under the direction of the client (as opposed to independently investigating disruptive technologies that may produce much more significant returns). We are attempting to approach problems from a knowledge management perspective and deliberately seek out and promote knowledge management opportunities.

**Measuring Value.** Network Map may be thought of as an information presentation tool, which, on the surface, produces little value. The task it automates is not formally part of the business process; when it is currently performed—either at the account representative’s request, or at the customers—it requires days or weeks of a support person’s time, and we may attribute this cost reduction to the application. These conventional, or mechanical, costs are easily quantified. However, the real value of the system is not reducing costs, but improving revenue generation and customer service. In the hands of an account representative (especially one with significant domain experience), this tool supports business and account analysis for the purpose of consultative selling and evolution of telecommunications services in concert with the future directions of a customer’s business.

The inability, on the surface, to accurately calculate a direct return on investment (ROI) is common among knowledge management applications. To address this issue, we have relied upon educated estimates of the application’s potential to reduce costs and to contribute

towards the generation of additional revenue. In addition, we have been attempting to formulate a business case for the application, based on its ability to advance the business process along a “knowledge scale.” Bohn (1994) has described an eight-level scale for measuring process, from complete ignorance of variables relevant to a process, through complete knowledge (a level never reached in practice). Our conclusion was that the current *information display creation* process was at a low level on this scale (approximately level 2), and through the application, we advanced to approximately level 5, by:

- Providing those data that are most relevant to the sales function;
- Making the time to produce a display predictable;
- Making the displays available to all account representatives, so that population as a whole is more uniformly informed;
- Removing variations in the display production methods; and,
- Enforcing uniform semantics for all data and information elements.

In our work to date, we have made the information display creation process one with an output whose mean (availability and method of production) and standard deviation (data semantics) are well-controlled at a high level of quality. In the next phase, we plan to move the *data analysis* process to a knowledge level where the mean and standard deviation are similarly controlled. We have been successful thus far in arguing that there is value in ascending the scale (where success is defined as achieving project funding and continued management support), and we hope to buttress our argument in 2000 with additional quantitative and anecdotal evidence of ROI.

**Display Technology.** We attribute a major part of the project success to the early selection and incorporation of the Inxight hyperbolic tree technology. The hyperbolic tree proved to be extremely useful in generating and manipulating the synoptic view of thousands of information elements that was necessary if our application was to be useful for performing strategic analysis. At the same time, the ease of use and the engaging nature of the display were valuable in generating early interest and enthusiasm.

The massive amount of data to be handled by the display technology is an ongoing issue only partially resolved in the current release. The current approach is object-oriented, in which thousands (or some cases, hundreds of thousands) of objects are created as the result of the database queries, then processed in order to create the synoptic view. While this view may be created in as little as three seconds in some cases, very large customers may require up to 20 hours.

The decision of what to incorporate, both visible and hidden, into the synoptic view has also been problematic. Early versions of the application incorporated telephone numbers, individual Universal Service Order Codes (USOCs), and other information; however, in some cases, this led to hyperbolic trees containing up to 100,000 nodes

(well beyond the 10,000 node limit recommended by Inxight). Our current release takes the approach that the hyperbolic tree is used for (1) viewing the entire customer, by geography, and (2) answering questions with regard to geographic location of specific services or circuits; accordingly, we have removed everything not compatible with those goals, resulting in displays that more closely conform to the recommended node count.

**Use of Java.** The hyperbolic tree applet uses the Java runtime environment native to Netscape (1.1.5); our other use of Java met with less success. Our original architecture called for a local data applet that would allow members of an account team to record information not available in the corporate repository<sup>1</sup>, exchange account-related messages, and capture issues and feedback for later, automatic transmission. Because of security issues and corporate architectural policies, the decision was made write the applet to the Java 1.2 runtime environment plug-in. Development proceeded well; however, installation and configuration of the plug-in and incompatibilities with other applications in the computing environment where the application was to be deployed led to a decision not to deploy the local data applet in the first phase. Deployment of the component has been indefinitely postponed until either (a) the 1.2/2.0 version of Java is supported by Netscape Navigator, or (b) the Java 1.2 plug-in becomes part of the standard software configuration in the deployment environment.

## Conclusions

We have identified a number of preliminary conclusions at this point in our development efforts.

1. The business problems we have identified must be solved if the organization’s strategic goals (to modernize, optimize and intelligently manage the client’s portfolio of telecommunications products and services) are to be met. Our early success in delivering a limited synoptic view, on time and on budget, has been key in making this point to management.
2. The enthusiasm among the application’s users and their clients, contrasted with the accounting perspective that the application provides little or no value, indicates that the organization is focused upon “return on investment,” rather than upon “return on knowledge.” We (as an organization) have begun to work on ways of quantitatively measuring the latter. In the absence of these metrics, we are relying on our sponsor’s belief that the application’s facilitation of a single sale to a

---

<sup>1</sup> By this, we mean notations on opportunities, prospective sales, account contact information, and the like. Representatives are not permitted to modify the data describing products and services; these changes must be made and validated via established processes, so that the data are corrected in the source of record.

strategic client would cost-justify the development.

3. The use of an innovative display technology (i.e., the hyperbolic tree) has been key, both from a technical perspective (for both information display and navigation), as well as from an application marketing perspective.

In addition to the above conclusions, we have also come to believe that the greatest barrier to projects of this type (i.e., those at the nexus between knowledge management, artificial intelligence, and business process management) are organizational and cultural, rather than technical. All the necessary technology exists; however, proceeding with a knowledge management project (and getting to the point where it can be integrated into the business process, and/or augmented with AI) often requires an *a priori* belief by management that adding knowledge to a process has a contributory effect to the bottom line. Network Map, and similar applications already under development, are establishing that the availability of such knowledge is likely to lead to improvements to or transformation of business processes.

### Acknowledgments

The authors wish to thank the project sponsors in the Business and Government Solutions (BGS) group, specifically Linda Kunzweiler and Sandi Miyaki.

### References

- Bohn, R. 1994. Measuring and Managing Technological Knowledge. *Sloan Management Review* 36(1):61-73.
- Brownston, L.; Farrell, R.; Kant, E.; and Martin, N. 1985. *Programming Expert Systems in OPS5*. Reading, MA: Addison-Wesley.
- Buchanan, B. and Shortliffe, E.H. eds. 1984. *Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project*. Reading, MA: Addison-Wesley.
- Kraft, A. (1984). XCON: An Expert Configuration System at Digital Equipment Corporation. In Winston, P.H. & Prendergast, K.A. (Eds.), *The AI Business*. Cambridge, MA: The MIT Press.
- Lamping, J.; Rao, R.; and Pirolli, P. 1995. A Focus+Context Technique Based on Hyperbolic Geometry for Visualizing Large Hierarchies. In Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems (ACM, May 1995).
- McDermott, J. (1982). R1: a rule-based configurer of computer systems. *Artificial Intelligence*, 19:39-88.
- McDermott, J. 1993. R1 ("XCON") at age 12: lessons from an elementary school achiever. *Artificial Intelligence* 59:241-247.
- Soloway, E., Bachant, J., and Jensen, K. (1987). Assessing the Maintainability of XCON-in-RIME: Coping with the Problems of a VERY Large Rule-Base. In Proceedings, Sixth National Conference on Artificial Intelligence (AAAI-87).

Wall, L., Christiansen, T., and Schwartz, R.L. (1996). *Programming Perl (Second Edition)*. Sebastopol, CA: O'Reilly & Associates.