TREEHOUSE SOFTWARE

TREETIPS

A Publication of Treehouse Software

Issue #19 - February 1995

tRelational Today

Improving access to corporate data has become an important focus for organizations seeking to gain a competitive advantage. Some sites want to move certain functional areas or applications to other platforms and relational database management systems (RDBMS). Some need to consolidate information from various sources into a Data Warehouse. Some may want to re-engineer legacy systems. Others may want to propagate mainframe data to other platforms and keep them somewhat synchronized. The objective of these efforts is to provide easy access to data for analysis and decision support. However, data transformation can be an overwhelming task. That is why TSI created **tRelational**.

tRelational is a tool-assisted ADABAS-to-RDBMS migration service. Our consultants use sophisticated software, known as the **tRelational Tools**, to ensure that migration is performed correctly, efficiently, and within a reasonable time frame. The **tRelational Tools** help the consultants to automate many time-consuming and difficult migration tasks, and make it easier for consultants to identify which data should be implemented in the RDBMS, how that data should be implemented (normalized or not), and in what format the data should be implemented.

Analysis of ADABAS File Structures

tRelational examines ADABAS FDTs, NATURAL DDMs, COBOL Copybooks, and other generated PREDICT objects to identify:

- Discrepancies between DDEs, FDTs, and userviews
- The need for field name changes
- Format, length, and null suppression conversion requirements
- Repeating fields eligible for normalization
- Files with multiple record types through analysis of userviews
- Unnecessary ADABAS descriptors
- Possible indices for RDBMS tables

The consultants work with the client to resolve these issues, taking into account the current and future uses of the data.

(continued on page 8)

Inside This Issue

TSI Helps Client Propagate Data

A recent **tRelational** client expressed an immediate need to propagate ADABAS data to multiple RDBMSs. The client has a large application built on ADABAS. This application creates and maintains data for the client's many customers. These customers would like their subsets of the ADABAS mainframe data to be distributed to relational databases on their local UNIX servers, where they can easily access it using SQL-based tools. Our client's customers want the data closely synchronized with the mainframe (i.e., UNIX data should be no more than 30 minutes out of synchronization with the mainframe). The client had considered another vendor's "transaction propagation" tool, but that tool did not meet their needs, which included propagating data to a non-ADABAS RDBMS. They asked us for help.

The client and TSI discussed how a combination of **tRelational** and **AUDITRE** could form the basis of a solution that would meet their needs:

- Assist with data modeling and mapping,
- Transform the current ADABAS data into relational format,
- Identify and filter relevant updates to the mainframe ADABAS data,
- Transform all updates into relational format,

(continued on page 3)

Editor's Sproutings

by Michael Salsbury

Preparing for the Year 2000

Customers have asked about the ability of our software to operate in the year 2000 and beyond.

Some TSI products do not use the date for comparisons or calculations, so there will be no problem. Some TSI products process all dates with 4-digit year values, or process the STCK time value (which is still valid after 1999). Other TSI products require changes which our developers have already identified.

Comparisons for today's date which assume that it will be higher than yesterday's date could fail at year 00. Parameters specifying date ranges from year 99 to year 00 will appear to be invalid, or produce invalid results (e.g., 1999 to 1900 instead of 1999 to 2000). Software which "postdates" transactions to occur a day, week, month, etc., later could have problems when the date nears 2000. We will address these and other situations in our software.

Users are advised to check with all of their software vendors regarding this century change. "Don't just ask the vendor if the software will work," says George Szakach, President of Treehouse

(continued on page 9)

TREETIPS

Editing, Writing, Design, and Layout
Michael Salsbury

Production and Distribution
Terri Hammerschmitt
Heather Mengel

Back issues available upon request. Documentation sets for all products also available.

Circulation: 12,400

Continuing to Travel

Since the last issue of *TREETIPS*, our people have visited Virginia, New York, Alabama, Texas, North Carolina, South Carolina, California, Oklahoma, DC, Germany, Canada, Massachusetts, and other sites in Pennsylvania. We hosted well-attended seminars in Boston, Charlotte, and Philadelphia. As this article is being written, TSI is already on the road, busily working at a customer site. (See the article "TSI Helps Client Propagate Data" on page 1 for more details.)

Sewickley Gets a Famous Houseguest

Actor/comedian **Sinbad** filmed the movie "Houseguest" in TSI's hometown, Sewickley, Pennsylvania. Scenes were shot in the streets within 100 yards of the TSI offices. We enjoyed seeing Sinbad walking around Sewickley, and visiting our shops and restaurants. When we last heard, Houseguest was the number 3 film at the box office.

TSI Affiliate Week

During the week of October 31, 1994, representatives from our affiliates in Spain, Germany, France, Belgium, Brazil, and Australia joined us for training and to exchange information. **Jim Lumpp** demonstrated and discussed the service and software aspects of **tRelational**, Treehouse Software's ADABAS-to-RDBMS migration service. The TSI developers discussed product internals. We also discussed technical documentation, support issues, sales/marketing, marketing materials, and other relevant issues.

First Annual International TSI Shuffleboard Tournament

During the Affiliate Week in October, we invited our guests to an American Dinner Party. During the party, the First Annual International TSI Shuffleboard Tournament was held. In the final round, the team of **Emilie Szakach** (of TSI) and **Wagner Martins** (of FYT - Brazil) defeated the team of **Joyce Maguire** (of TSI) and **Carlos Usin Lacarcel** (of Quark Software & Services - Spain) to claim the title of International Shuffleboard Champions!

Joyce Takes Charge

We are very pleased with the job that Joyce Maguire has been doing for TSI. Joyce had earlier directed the development of **PROFILER**, **SECURITRE**, and **DynaDoc**. On September 1, Joyce became TSI's Director of Technical Development. In this new position, Joyce takes on the leadership responsibility for all of our current products, and for any new products we may develop or acquire. You will soon begin to see the positive effects of Joyce's leadership in all our products.

Donna Helps a Site Improve Performance

A site in Texas wanted a written review of their ADABAS/NATURAL environment, focusing on several specific performance issues, such as physical dataset placement, avoiding format translations, and setting programming standards to improve performance. **Donna Heimbaugh** of TSI visited them, examined their environment, and prepared a report which outlined several things the site could do to improve performance. The site was very happy with Donna's report, and plans to implement her suggestions soon.

(continued on next page)

Editor's Sproutings

(continued from previous page)

TSI Helps Client Propagate Data

(continued from page 1)

- Transmit the original data and subsequent updates over a network to the customers' UNIX servers, and
- Apply the updates via software on the customers' UNIX servers and send confirmation back to the mainframe.

TSI currently has several people working at the client's site to create a detailed specification for the development of the desired data propagation solution. Our technicians are considering important issues such as the potential impact of running certain ADABAS utilities (e.g., mass updates, reloading of data), network traffic, modifications to ADABAS file formats or RDBMS table layouts, verification that updates have been sent by the mainframe and applied by the customers' servers, and many more issues that can impact the propagation process. Once the full data propagation solution has been developed, we know that many ADABAS sites will be interested in it, and we will make it available to all of them.

Due to the changing nature of many sites' computing environments, we anticipate that more and more sites will be looking for ways to make data accessible to end users. Data migration solutions (like tRelational) and data propagation solutions (like the one discussed in this article) are sure to become increasingly important. Treehouse Software is continuing to develop and enhance products to meet the changing needs of its customers. Watch *TREETIPS* for more news.

TSI Wins Award!

In the last issue of *TREETIPS*, we reported that Treehouse Software was a finalist in the 1994 MCI/Inc. Positive Performer Awards competition, which recognizes a company's commitment to customer satisfaction. **George Szakach**, **Emilie Szakach**, and **Michael Salsbury** traveled to New York City in late August to attend the awards ceremony. Forty companies from around the country received an award. Only three of them were high-technology companies. **TSI was the only software company to receive an award.**

USAir Tragedy in September 1994

Several of you called us in September after the crash of the USAir flight between Chicago and Pittsburgh. We did have people returning from Chicago at the time, but none of our people (or their families) were aboard that flight. Thank you for your concern. Our hearts go out to the families of those who died in the tragedy.

SAG Client/Server Conference

Jim Rock and **Jim Lumpp** represented TSI at the Software AG Client/Server Case Studies Conference held November 6-9, 1994, in Washington, DC. An estimated 530 people attended the conference, including customers, Software AG personnel, Software AG business partners, and others.

Software AG's new Data Warehouse program was a major theme of the conference. Building a Data Warehouse involves extracting data from various database management systems and placing it into a single repository. Users may access this Data Warehouse without impacting production systems. The first step in Data Warehousing is to select, map, and transform (migrate) data into the Data Warehouse (i.e., exactly what **tRelational** is designed to do). The next step is to make this data accessible to the end users for querying and "analytical discovery".

Fact Sheets Available

TSI recently produced a Fact Sheet for each of our products. Fact Sheets provide a one-page summary of the basic information about each product. To receive a Fact Sheet for any or all of our products, please contact us at your convenience.

Please Return Your Old TSI Release Tapes!

You may find a number of old Treehouse Software product release tapes or cartridges sitting around your data center, gathering dust. If so, please return them to us. We can erase and recycle these tapes for future product releases. Recycling helps us to keep your maintenance costs low, and helps to protect the environment. Thank you for your help! ▲

PROFILER Version 3.1.0

PROFILER for NATURAL, Treehouse Software's Quality Assurance and Testing Tool for NATURAL, offers a substantial increase in functionality and ease of use in its latest release, Version 3.1.0.

Limitations Removed

The previous **PROFILER** limitations on the collection of statistics for very large programs have been removed. Statistics may now be collected and reported for up to 2,018 executable statements per NATURAL program. Up to 255 **PROFILER** sessions may be defined. Up to 44 database accessing statements per program may be profiled. These maximums should handle the largest NATURAL applications.

Background Monitoring Capability

In **PROFILER** 3.1.0, the password-controlled Background Monitoring facility allows authorized users to cause a profiling session to be automatically activated for themselves or for a third party, based on library/program masks, NATURAL object types, and date/time. Background Monitoring parameters are specified on the screen shown below:

```
99-12-31                *** PROFILER for NATURAL ***                    USER24
10:00:00             Define Background Monitoring Record                PROLIB

     Define Background Monitoring Record for USER-ID USER24

     Session:                  PAYTEST
     Library Mask:             PAYROLL
     Program Mask:             PAY*

     Mask Options:             * Matches any character in remainder of string
                               ? Matches a single character
                               . Matches any numeric character
                               @ Matches any non-numeric character

     Object Types:             X Program      X Subprogram   X Subroutine
                               _ Map          _ Helproutine

     Starting Profiling Date:  12 31 1999 (mmddyyyy)    Time: 11 00 (hhmm)
     Stop Profiling Date:      12 31 1999 (mmddyyyy)    Time: 13 30 (hhmm)
                                                                   BCK DEF
Enter-PF1---PF2---PF3---PF4---PF5---PF6---PF7---PF8---PF9---PF10--PF11--PF12-
      HELP        END                                                  EXIT
```

Enhanced Reporting Features

With the **PROFILER** Enhanced Reporting Facility, statistics for multiple sessions, User-IDs, libraries, programs, object types, and program catalog timestamps, can be merged on **PROFILER** Enhanced Reports.

The **PROFILER** Enhanced Reporting Facility saves time and effort. Consider a site where a team of users tests the same application. Each user tests a specific function or component of the application and uses a separate **PROFILER** session. Using the Enhanced Reporting Facility, the team members can merge their statistics on one report to show that the entire application has been fully tested.

The layout of the **PROFILER** Enhanced Reporting File is provided to facilitate the development of site-specific reports which display **PROFILER** statistics in any desired format.

(continued on next page)

PROFILER Behind the Scenes

A number of "behind the scenes" changes have been made to PROFILER, increasing its power and ease of use, The changes include the following:

- **Improved Ease of Use**: Session names, rather than session numbers, now identify **PROFILER** sessions. Improved navigation and interactive capabilities have been added. The last session and User-ID specified are used as the default on pop-up windows.
- **Enhanced Statistics Collection**: Profiling and reporting may be performed for specific users and the specific NATURAL object types of interest to the users (Programs, Subroutines, Subprograms, Maps, Helproutines).
- **Improved Reporting**: Left/right scrolling is available in appropriate on-line reports. The beginning statement of each block of NATURAL Optimized Code (NOC) is identified on the Program Listing and Enhanced Program Listing Reports. Input parameters filter report contents by library, starting program, ending program, object type, and User-ID.

TSI Internet Policy

Treehouse Software is connected to the Internet, enabling us to send and receive E-Mail and subscribe to discussion lists, such as the SAG Discussion List hosted by the University of Arkansas (SAG-L@UAFSYSB.UARK.EDU).

Some organizations have flooded the Internet with the electronic equivalent of junk mail. Their actions have upset many Internet users. TSI does not want to follow their bad examples, so we have developed a policy regarding the Internet and how we will use it.

Treehouse Software respects the non-commercial nature of the Internet. We will not use the Internet to deliver "electronic junk mail". If a specific user site asks to receive information from us via the Internet, we will respond directly to that specific person, and not to a general discussion list. If a site appears to be unaware that a solution exists in a TSI product, we may call or write that individual at that site - outside the Internet.

As a SAG customer, we may contribute to SAG-related discussions on the Internet when there is a technical question we can answer, or a problem that we can help a user to solve. If the solution requires us to mention our products, we will say as little as necessary about them.

Internet messages are visible to many current and potential customer sites, as well as to software vendors. As a result, Treehouse may contribute to an Internet discussion if our company or our products are discussed in a message, and we feel that the message contains inaccurate, incomplete, or misleading information. Our response in such cases will attempt only to correct any misinformation in the original message, and will not be used as an excuse to deliver a "sales pitch".

We believe that ours is a responsible policy, which protects the non-commercial nature of the Internet and allows us to participate in a meaningful manner. We hope that other vendors adopt similar policies for their Internet usage.

PROFILER Version 3.1.0

(continued from previous page)

Improved Handling of NATURAL Copycode

In **PROFILER** Version 3.1.0, statistics for NATURAL copycode are collected separately from those collected for the program containing the INCLUDE statement. The NATURAL Program Listing report displays the INCLUDE statement and expands the copycode when the user places the cursor on the INCLUDE statement and presses a PF-Key. When the user is finished viewing the statistics for the copycode, **PROFILER** will return to the Program Listing Report for the program which included the copycode.

The **PROFILER** Enhanced Program Listing Report, part of the new Enhanced Reporting Facility, automatically expands copycode, identifying it with the letter "C" following the statement number, as shown below (note the lines following the INCLUDE at "Stmt Nbr" 0060):

```
99-12-31                  *** PROFILER for NATURAL ***                  USER24
10:00:00               Enhanced Program Listing Source Code             PROLIB
Session: TESTPROF                                                       Page 2
User-ID: USER24
Library: PAYTEST  Program: PAYEMPLP  Version: 99-12-31 09:37:26
>top_
Statement   Total        Average          Stmt
Count       CPU (msec)   CPU (msec)  N/E  Nbr   NATURAL Statement
----------  -----------  ----------- ---  ----  --------------------------------
                                          0010  DEFINE DATA LOCAL USING PAYEMPLL
                                          0020  END-DEFINE
       442       14.551         .033      0030  READ EMPLOYEES-VIEW
<<DATABASE     20054.928       45.373>>
       441        2.816         .006      0040  DECIDE FOR FIRST CONDITION
                                          0050   WHEN ZIP = MASK ('9')
        10         .022         .002      0060    INCLUDE PAYEMPLC
        10         .068         .007      0010C SET KEY PF3
        10         .180         .018      0020C INPUT USING MAP 'PAYEMPLM'
        10         .084         .008      0030C  IF *PF-KEY = 'PF3'
                                       *  0040C   FETCH 'PAYENDP'
                                          0050C  END-IF
       441        2.672         .006      0070  WHEN ZIP = MASK ('0')
        22        8.617         .392      0080   WRITE 'UPDATE ZIP FOR EMPLOYEE'
--------------------------------------------------------------------------------
Enter-PF1---PF2---PF3---PF4---PF5---PF6---PF7---PF8---PF9---PF10--PF11--PF12---
      HELP  COPY  END                      -     +           <     >    EXIT
```

The field next to the ">" character above the Statement Count column is a command line. By entering a specific statement number here, the user instructs **PROFILER** to position the report to that statement. Navigation commands (such as top, bottom, +, ++, -, or --) may also be entered on the command line, allowing the user to easily page through the report.

PROFILER Verify Program

PROFILER 3.1.0 includes a special program (PROVRFY) which verifies that **PROFILER** is properly installed. PROVRFY ensures that NATURAL is properly configured for **PROFILER**, that the **PROFILER** modules are in place, that the proper statistics file is installed and accessible, etc. It then mimics a profiling session, verifying that statistics can be collected and stored in the statistics file. It also provides diagnostic messages which identify possible installation problems and suggests how to resolve them.

PROFILER continues to grow and improve in response to customer requests. If you develop NATURAL applications, you need **PROFILER**. Contact TSI or your TSI affiliate today!▲

N_2O Tracks Projects!

Over the past several months, **N_2O** has taken a quantum leap forward. **N_2O** Version 3.3 will show the results of the **N_2O** Team's extensive efforts in a variety of new features, including a **Project Tracking Subsystem**.

The N_2O Project Tracking Subsystem

In TREETIPS 17, we discussed the concepts behind a Project Management feature under consideration for a future **N_2O** release. The developers received suggestions and information from a number of users, and incorporated these ideas into the **N_2O Project Tracking Subsystem**.

N_2O sites can use this subsystem to track a project from the moment it is proposed through to its approval and use in production. The screen below shows how a Project is defined to **N_2O**:

```
95-02-15                    N2O PROJECT DEFINITION                    TSIUSR01
09:38:00                                                              TSI019

  Project    :  PAYROLL
  Updated    :  95-02-15  09:38:00  TSI0171
  Short Desc:  THIS IS A SAMPLE PROJECT

                                                    Priority     Impact
 Task Groups    Stages                              Values       Values
 ------------   -------------------------------     ----------   ---------
  1. BENEFITS    1. PRE-DESIGN   11.                 High A1      High H
  2. TAXES       2. INITIATED    12.                      A2           M
  3. INSURNCE    3. DEVELOPED    13.                      A3           L
  4. MISC        4. APPROVAL-1   14.
  5.             5. APPROVAL-2   15.
  6.             6. COMPLETED    16.
  7.             7.              17.
  8.             8.              18.
  9.             9.              19.
 10.            10.              20.                 Low          Low

Enter-PF1---PF2---PF3---PF4---PF5---PF6---PF7---PF8---PF9---PF10--PF11--PF12---
      ----  ----  END   ----  ----  ----  ----  ----  ----  ----  ----  STOP
```

Each Project has certain tasks associated with it. Task Groups are used to logically separate related tasks within a particular project. Stages relate to the logical progression of a task from beginning to end (e.g., milestones). Each task can have a Priority Value assigned to it, to indicate its relative urgency. A task may also have an Impact Value associated with it, to indicate the impact of a particular task on the Project.

The values for the Task Groups, Stages, Priority Values, and Impact Values are all user-defined values. The **N_2O Project Tracking Subsystem** supports up to 10 Task Groups, up to 20 Stages, up to 10 Priority Values, and up to 10 Impact Values per Project. This allows for complete flexibility. For example, one site may wish to have Task Groups which relate to a specific functional area, such as "Documentation", "Programming", and "Testing". Another site may wish to have Task Groups such as "Enhancements", "Maintenance", and "Training". One site may wish to have Priority Values which are numeric (e.g., 1=low priority, 10=high priority). Another site may wish to have mnemonic priority values, such as "H" for "High", "M" for "Medium", and "L" for "Low".

(continued on next page)

Reporting Enhancements and Other Improvements

A number of enhancements have been made to the Reporting Subsystem to make it easier to get information from the **N_2O** audit trail. Many reports have been enhanced, and new reports have been added. One of the new reports is the "Events Processed by Date" report, which identifies all Events processed during a specific date range (e.g., 95-01-17 to 95-01-18). All **N_2O** reports are now offered both on-line and in batch, and allow the user to view additional "Event Details" and "Object Details".

A number of other improvements have also been made. It is now possible to specify the order in which individuals must approve an Event. **N_2O** now uses LFILEs set aside by Software AG for the exclusive use of **N_2O**. This will ensure that our use of LFILEs will not cause conflicts with Software AG products which also use LFILEs.

N_2O is also getting ready for the year 2000 and beyond. Starting with Version 3.3, **N_2O** will store all year values as 4-digit numbers to ensure that all processing and reporting will be unaffected by the change from 1999 to 2000. Conversion routines will convert old-style audit trail information into the new format.

N_2O Tracks Projects

(continued from previous page)

The **Project Tracking Subsystem** includes a number of reports that make it possible to view project tracking and status information. For example, consider the "History of a Task" report shown below:

```
95-02-15                N-2-O PROJECT TRACKING REPORTS                TSIUSR01
09:38:00                HISTORY OF A TASK  BENEFITS 000001            TSI019
                        Project:  PAYROLL                  Page: 1

Stage       Date      Time      User-ID   Comments
----------  --------  --------  --------  --------------------------------
COMPLETED   95-01-25  18:35:25  TSI0171   TASK COMPLETED TO SPECIFICATIONS
APPROVAL-2  95-01-23  09:34:38  TSI0197   APPROVED BY SYSTEM MANAGER
APPROVAL-1  95-01-22  10:10:31  TSI0191   APPROVED BY DEPT. LEADER
DEVELOPED   95-01-20  14:05:54  TSI0174   ADJUSTED TO STANDARDS
REJECT      95-01-18  11:05:54  TSI0191   DID NOT MEET CODING STANDARDS
DEVELOPED   95-01-08  17:32:35  TSI0171   CODING COMPLETED FOR NEW BENEFITS
INITIATED   94-12-27  08:32:05  TSI0171   BEGAN CHANGES FOR NEW BENEFITS
DEFINED     94-11-21  15:07:20  TSI0171   ADD NEW BENEFITS TO PAYROLL APPL.



Enter-PF1---PF2---PF3---PF4---PF5---PF6---PF7---PF8---PF9---PF10--PF11--PF12---
      ----  ----  END   ----  ----  ----  ----  ----  ----  ----  ----  STOP
```

This report shows the progression of Task number 000001 for the BENEFITS Task Group in the PAYROLL Project. From the information in the Comments column, we see that this Task related to changes needed in the PAYROLL Project to reflect the addition of new benefits. In this report, we see the history of this task from the moment it was defined on November 21, 1994, up to the point that it was officially completed on January 25, 1995.

No Price Increase!

The **Project Tracking Subsystem** is a new component of **N_2O**, and something unavailable in other NATURAL change management products. In spite of the substantial leap in functionality that **N_2O** Version 3.3 represents, there will be no immediate price increase, and no additional charge to existing **N_2O** customers for the **Project Tracking Subsystem**. (However, TSI reserves the right to charge future **N_2O** customers separately for the **Project Tracking Subsystem**.)▲

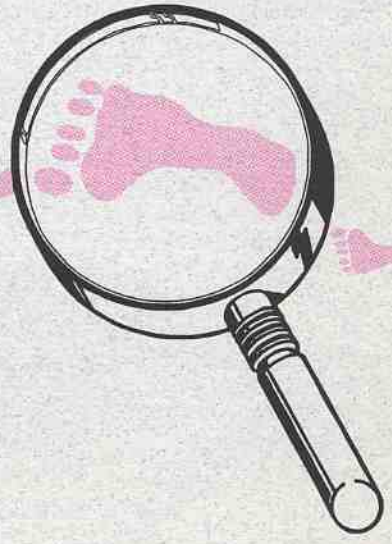

tRelational

(continued from page 1)

TSI and HCL: A Perfect Combination

We are working on a data propagation project (described on page 1) with HCL America.

HCL America was founded in 1988 and currently has over 250 full-time consultants. It is a member of the HCL worldwide information technologies group which is comprised of 11 companies operating in 5 countries, and has over 4,500 employees. The group's 1994 revenue was over $275 million. HCL's customer list includes AT&T, Novell, Oracle, Silicon Graphics, Hewlett Packard, Sun Microsystems, Sybase, and Autodesk. In addition to on-site consulting services, HCL also offers offshore development services at a state-of-the-art software factory in India. Their people have worked on a number of data migration, conversion, and propagation projects.

"We were very pleased when Treehouse approached us to work with them on this project, "said Shiv Kumar of HCL America, "We have developed similar data propagation capabilities in the past, and have considerable experience with UNIX and RDBMS products. Treehouse has the in-depth knowledge of ADABAS and NATURAL along with the critical software tools needed to complete this project. We look forward to working with them on many future tRelational and data propagation projects."

Richard Jacobson, Vice President of Treehouse Software, told *TREETIPS*, "HCL is definitely the right company to work with. Their people seem to be very in-tune with what we are trying to accomplish for the customer, and they are very eager to make sure that the customer is happy. Their large pool of talent is just what we need to bring tRelational services to the sites who need them."

The **tRelational Tools** are used to resolve discrepancies between the way a particular field is defined in the ADABAS FDT, the PREDICT DDE, and the various PREDICT userviews, as shown in the screen below:

```
99-12-31            RESOLVE FIELD DISCREPENCIES                    TREE03
11:38:00            FILE: EMPLOYEES                                TSI703

Implement?  Y     MIDDLE-NAME                          AC

                  ATTRIBUTE          PREF        FDT        DDE
                  -----------        ----        ---        ---
                  Format                A          A          A
                  Length             20.0       20.0       16.0
                  Suppression
                  Descriptor            D          D          D
                  Uniqueness
                  Field-type
                  PE Member
                  Super Usage

Userview                             Userview Field Name                Length
EMPLVW1                              MIDDLE-NM                          16
EMPLVW2                              MIDDLE-INITIAL                     1


Enter-PF1---PF2---PF3---PF4---PF5---PF6---PF7---PF8---PF9---PF10--PF11--PF12---
      ----  ----  ----  VIEW  ----  ----  ----  ----  ----  ----  ----  ----
```

In the above example, **tRelational** provides relevant information about the attributes of an ADABAS field (MIDDLE-NAME, field AC in the ADABAS FDT), and attempts to identify a preferred (see the "PREF" column) way to implement this field in the RDBMS. (Although it cannot be shown here, all discrepancies are highlighted on the actual **tRelational** screen.) The "MIDDLE-NAME" field is defined to ADABAS as an alphanumeric field of length 20. The PREDICT DDE for the file (and the userview EMPLVW1) identifies the length of the field as 16. The userview EMPLVW2 refers to the field as "MIDDLE-INITIAL" and identifies its length as 1.

Based on information from the client, the **tRelational** consultant determines whether or not to implement the field in the RDBMS, and enters the preferred name, format, length, etc., in the appropriate places on the screen.

Using the **tRelational Tools**, the consultant may learn that although the field is defined to ADABAS as a 20-character field, only one character is actually being used. In this case, it might be appropriate to implement the field as a single-character column in the RDBMS. **tRelational** would illustrate this situation in the report shown below, which displays the results of an analysis of 10,000 records on the EMPLOYEES file:

```
99-12-31            VARIABLE CHARACTER FIELD STATISTICS             TREE03
11:38:00            File: EMPLOYEES                                 TSI703
                    Records Analyzed: 10000

                                                  *****  Lengths  *****
 S  DB Field Name                          Ty PE D  Defined Highest Average
 -  -- -------------------------------    -- -- -  ------- ------- -------
    AB FIRST-NAME                                      20.0      18     5.8
    AC MIDDLE-NAME                                     20.0       1     1.0
    AD LAST-NAME                                    D  20.0      17     6.8
    AI ADDRESS-LINE                            MU      20.0      20    15.8
    AJ CITY                                         D  20.0      20     8.6
    AP JOB-TITLE                                    D  25.0      25    12.4
    AT CHILD-NAME                                      30.0      11    11.0
```

tRelational also looks at the actual data to be migrated, to determine how repeating fields are being used (e.g., highest and average number of occurrences), the range

(continued on next page)

tRelational

(continued from previous page)

of values present for a descriptor, the uniqueness of the descriptor, etc. In addition, data collected by **TRIM** (or other performance monitoring tools) is analyzed to determine which descriptors are not being used, so that they can be eliminated. The analysis even helps determine which descriptors are used most often, making it easier to establish the order of the physical data (i.e., "clustering").

The screen below shows statistics gathered about multiple-valued fields:

```
99-12-31                REPEATING FIELD STATISTICS                    TREE03
11:38:00                FILE: EMPLOYEES                               TSI703
                        Records Analyzed: 1000

                                                    #Records #Records * Occurrences *
S  DB Field Name                          Ty PE D    w/Mult   w/Any   Highest Average
-  -- ------------------------------ -- -- --  -------- -------- ------- -------
   AI ADDRESS-LINE                   MU          660      1000       4     2.1
   AQ CHILD-INFO                     PE          680       725      11     2.5
   AW LEAVE-BOOKED                   PE          400       410       3     1.7
   AZ LANG                           MU    D     200       750       3     1.0
```

It would appear that the "LANG" field could be implemented as 3 columns in the RDBMS. Only a small percentage of records contain multiple values in this field, and the highest number of occurrences for the field is 3.

The site may wish to normalize the "CHILD-INFO" periodic group, since many records contain multiple occurrences for this group.

Suppose that the consultant decides to examine the CHILD-INFO periodic group and determines that it consists of CHILD-BIRTH, CHILD-SSN, and CHILD-NAME fields. The CHILD-BIRTH field contains the child's date of birth. The CHILD-SSN field contains the child's Social Security Number or other identification code number. The CHILD-NAME field, however, contains the child's first, middle, and last names combined into one field. The PREDICT userview used to access this field automatically redefines it into fields containing the first name, middle name, and last name. After conferring with the client, the consultant may learn that the field should be implemented as three RDBMS columns in order to meet future data access needs.

Development and Physical Implementation of a Logical Model

Once the file structure and usage information has been gathered and analyzed, the **tRelational Tools** help the consultants design a logical model. The development steps of a logical model include identifying Data Entities, defining the relationships between the Data Entities, establishing standards and naming conventions (including abbreviations, default data types, default lengths, and null conversion), identifying a Primary Key (based on a unique descriptor or a newly-defined index), establishing referential integrity through foreign keys, and assigning the attributes for each entity (based on the fields of a file, expansion of redefined fields, or assigning new attributes). Data Entities can be based on an ADABAS file, a repeating group, a userview, or other parameters. Data Entities can also be newly-created entities with no relationship to the ADABAS data.

The **tRelational Tools** are then used to implement the physical model, identify alternate keys for each Table, calculate storage parameters, enforce referential integrity (through rules, triggers, and stored procedures), and generate Data Definition Language (DDL) for objects such as Tables and Indices.

(continued on next page)

Preparing for the Year 2000

(continued from page 2)

Software, "Ask about specific situations, like date comparisons, postdating, or displaying a list of data in chronological order (where 00 would appear before 99). Ask how the vendor has solved these potential problems. Most importantly, plan to bring your systems down on December 31, 1999, in the evening, and back up on January 1, 2000. I predict havoc for those companies running software through this date change. You can probably imagine many of the problems you might see, like dual log swapping problems based on which log is the 'oldest', etc."

George adds, "In 1961, while working on date subroutines, I became concerned about this date problem that will occur near the end of the century. I quickly calculated that I would be 57, retired, and I would be long gone from data processing in 1999. A few years ago, I predicted that this date confusion would get huge and ugly, and would spawn new businesses and software dedicated to solving the problem."

George was right. Several products are on the market already, and a number of companies claim to have some kind of solution for automatically changing all software to operate correctly into the next century. However, he was probably wrong about retiring by 1999.

TRIM tRELOG Feature

The **TRIM** Development Team has been busy adding new functions to **TRIM** which support **tRelational**, our tool-assisted ADABAS to RDBMS migration service.) This ne. **TRIM** functionality provides **tRelational** with a picture of how the ADABAS data is currently being used, to help a site predict how the data might later be used in the RDBMS. It is also useful for sites with no RDBMS migration plans. We thought you might like to know more about this new functionality, called TRELOG.

TRELOG is an optional **TRIM** feature which provides **tRelational** with information about ADABAS data usage. TRELOG is activated for specific files on specific databases to monitor the contents of the ADABAS Control Block and certain buffers for each command issued by a user. TRELOG accumulates and summarizes this information in **TRIM** User-Exit-4, in a manner very similar to the current PRESUM feature. Periodically, the TRELOG information is written to the ADABAS Command Log for later analysis and reporting by **TRIM** batch facilities.

The **TRIM** batch facilities recognize the TRELOG records and produce summarized or detail datasets containing statistics which can be used as input to **tRelational** or homegrown programs, and identify:

- Which descriptors were used on each file during the monitoring period
- How often each descriptor was used
- The volume of update activity versus retrieval activity for each file

These statistics can help a DBA make a number of intelligent decisions during ADABAS to RDBMS migrations, and can also help sites improve ADABAS performance.▲

tRelational Partnerships

With respect to **tRelational**, Treehouse Software has established relationships with several major vendors, including:

- Oracle: Business Alliance Program
- IBM: DB2 Developer Assistance Program, Business Partner
- Informix: InSync Software Partnerships Unlimited

We are currently speaking with other vendors regarding partnering issues. We want to ensure that **tRelational** is a viable, tested option for migrating data from ADABAS to all major RDBMS and Data Warehouse products.

tRelational

(continued from previous page)

Migration of the Data

With the logical model developed and implemented in the RDBMS, **tRelational** can migrate the ADABAS data into the RDBMS objects. **tRelational** extracts the data from the ADABAS file(s) and writes it to work files representing the base and dependent tables. **tRelational** then sorts these work files, based on information gathered during the prior analysis of summarized Command Log data. Finally, **tRelational** generates the appropriate SQL statements to load the sorted data into the target RDBMS.

The migration process includes a filtering or "cleansing" process, which removes unwanted data, corrects invalid field values, translates field formats (e.g., EBCDIC to ASCII), adjusts field lengths, etc. This ensures that the data is ready for use when it is loaded into the RDBMS.

tRelational supports Software AG's new relational DBMS, ADABAS D, and all SQL-based RDBMS products, including Oracle, Sybase, Informix, Ingres, and DB2.▲

tRelational and Data Warehousing

Data Warehousing, the concept of merging all of an organization's data into a single database (or "warehouse"), is gaining a lot of attention. Several vendors have announced Data Warehouse products, including IBM and Software AG. Customers have asked if **tRelational** can help them to get their ADABAS data into a Data Warehouse. The answer is yes.

SECURITRE Multi-Platform

SECURITRE customers requested a way to restrict access to mainframe ADABAS data from calls originating on other platforms, such as UNIX, Windows, and OS/2. We felt that it was important to address this issue as quickly as possible.

On the mainframe, a **SECURITRE** user-exit to the ADABAS Link Routine places information about a user (such as the user's User-ID, NATURAL Library, etc.) into the User Information Area for each ADABAS call. (The User Information Area is part of the User Buffer which is attached to the call by the ADABAS Link Routine as the seventh parameter.) The **SECURITRE** User-Exit-1 to ADABAS extracts this data from the User Information Area and uses it for security checking. No User Information Area is currently available for calls originating on non-mainframe platforms (i.e., there is no seventh parameter). We understand that by mid-1995 Software AG plans to provide a User Information Area as part of non-mainframe ADABAS calls. Until then, TSI has developed an interim solution.

The interim solution consists of source code (written in C) for a **SECURITRE** user-exit to the ADABAS Link Routine on the non-mainframe platform. The source code contains a Software AG supplied function call which determines the user's User-ID. Some sites may choose to replace this function call with a home-grown function which substitutes the individual's mainframe User-ID, a LAN User-ID, etc. The desired User-ID is placed into the ADDITIONS-3 or ADDITIONS-4 field of the ADABAS Control Block before the call is sent to the mainframe. **SECURITRE** on the mainframe will extract the User-ID and perform the security checks.

When this solution is implemented at a site, **SECURITRE** will limit access to *mainframe* ADABAS data from calls originating on both mainframe and non-mainframe platforms.▲

TREEHOUSE SOFTWARE

Affiliate Directory

AL-FALAK
(Persian Gulf)
973-723-001

Blenheim Software Ltd.
(United Kingdom & Eire)
0753-571770

Computer Consultants Australia
(Australia & New Zealand)
61-3-416-3377

Fairware
(France)
33-1-39-65-06-88

FYT
(Brazil)
55-11-253-2288

Information Technology Services
(South Africa)
27-11-886-7690

MaK DATA SYSTEM Kiel GmbH
(Germany)
49-431-3995-135

MegaSoft
(Venezuela & Caribbean)
58-2-693-3211

Metastore
(Benelux)
03-238.87.60

Quark Software Services
(Spain & Portugal)
34-350-18-55

Softeco Sismat
(Italy)
39-2-7386501
(Switzerland)
41-21-311-3737

System Networks (Asia) Ltd.
(Southeast Asia)
852-530-5480

MetaStore Announces SQL Command Analysis

SQL Command Analysis operates on SQL access modules and on source programs written in NATURAL, COBOL, CSP, etc., with SQL calls. It is an indispensable tool for SQL programmers to understand how a command will be executed and to evaluate it from a quality assurance and performance viewpoint. SQL Command Analysis performs automated EXPLAIN of the program's SQL commands and the Text Analysis component scans the program for suboptimal SQL command clauses. SQL Command Analysis monitors access path changes in stored SQL access modules.

For more information about this product, contact Leo Van Dongen of METASTORE at +32-3-238.87.60.

Publication of this announcement in TREETIPS does not constitute endorsement by Treehouse Software or its other affiliates

TREETIPS

Treehouse Software products and services include:

tRelational - Tool-assisted ADABAS to RDBMS migration service
PROFILER for NATURAL - NATURAL quality assurance and testing tool
DynaDoc - NATURAL application documentation tool
N_2O - NATURAL application change management system
N_2O/3GL - 3GL support within N_2O
SECURITRE - Centralized security administration for ADABAS/NATURAL
TRIM® - ADABAS/NATURAL performance monitor
AUDITRE - ADABAS auditing tool
AUTOLOADER - ADABAS file automatic unload/reload/dump utility
LumeNAT - NATURAL monitoring tool, coming soon from Treehouse

Consulting and Customized on-site classes designed to meet your unique needs.

Phone:
(412) 741-1677

Fax:
(412) 741-7245

IBM Business Partner

ORACLE BUSINESS ALLIANCE PROGRAM

SOFTWARE AG INTERFACE PARTNER for PROFILER for NATURAL

INFORMIX InSync *Software Partnerships Unlimited*

TREEHOUSE SOFTWARE
409 Broad Street, Suite 140
Sewickley, PA 15143-1511

FIRST-CLASS MAIL
U.S. POSTAGE
PAID
Sewickley, PA 15143
Permit No. 89