Issue #23 • February 2013

Tree•times

A Publication of Treehouse Software, Inc.

FINAL PRINTED ISSUE

Snippets From the Treehouse President's Blog

"I remember squashing all the Procedure Division code onto as few cards as possible so that the deck would fit in my shirt pocket with all the colored markers."

Read more of George's Blog at: www.treehouse.com/blog.

Visit us on Facebook
www.facebook.com/TreehouseSoftware

LinkedIn
www.linkedin.com/company/39849

Treehouse Software

The Last of the Real World Series

by Wayne Lashley, Chief Business Development Officer

Those of you who have followed Treetimes (and its predecessor, Treetips) over the years will know that we have had a recurring feature under the banner of "Real World Series", describing case studies involving Treehouse customers achieving business goals through our data replication solutions. It has been my privilege over the years to have authored several of these—and to have participated in some of the actual project implementations described therein.

It has always been important to us to focus on the "real world" aspect. Our customers come to us with real challenges, and need solutions that work on platforms beyond just PowerPoint. We have seen other vendors come and go with spectacular-sounding ideas, but these ideas don't always scale, or can't always cope with the real world that our customers inhabit.

In our customers' real world, mainframe cycles are still an expensive resource to be managed carefully. Networks and servers sometimes break down, and sometimes they are slower than desired. Mission-critical databases are supporting mission-critical operational applications, and can't be burdened with complex, performance-killing ad-hoc queries and "data exploration". Platforms and applications require periodic maintenance, necessitating outages.

What our customers need is a highly fault-tolerant way for heterogeneous applications and databases to act in a coordinated fashion to accomplish business goals.

These days, "data virtualization" is a hot topic. The objective of data virtualization is to abstract data sources across the enterprise to present the client (data consumer) with access to logical, composite "business entities" whose technical implementation—location, structure, access method, etc.—are deliberately obscured. This is a noble and lofty goal, but it is really just an extension of the concept of "data federation"—which has been with us for at least two decades.

Federated databases are those where middleware enables clients to access heterogeneous data sources transparently, e.g., to issue a SQL query which JOINs mainframe IMS/DB segments with Oracle tables on Linux. This often

(continued on page 3)

Editor's Notes

by Joseph Brady, Marketing and Documentation Manager

Become a Beta Test Site for New Treehouse Product Releases

We are looking for current North American customers to become Beta test sites for new releases of Treehouse products. If you are interested in becoming a Beta test site, please contact us at www.treehouse.com/contact-us-0, and a Treehouse representative will contact you.

Find out the Latest on Treehouse Products

To find out about current versions of Treehouse products, compatibility (operating systems, languages, etc.), and support information for all of our products, view the TSI Product Status Matrix on-line at https://support.treehouse.com/SUPPORT/prodstatus.shtml.

Online Product Demos

Would you like to see Treehouse products in action before you request a trial? To set up a live, online demonstration of any product, simply fill out the short form on the Treehouse Software website at www.treehouse.com/request-a-demo.

Treetimes

Editing, Writing, and Design
Joseph Brady

Contributors
Phyllis Doran, Wayne Lashley, George Szakach

Distribution
Terri Hammerschmitt

Back issues are available on the Treehouse website. Documentation for all products is available on CD.

Hard Copy Circulation: 8,000

Next Issue: 0

Snippets From the Treehouse President's Blog

"I used to think, 'why not 6 Presidents and one VP?'"

Read more of George's Blog at: www.treehouse.com/blog.

When Treehouse Software founder **George Szakach** returned from the Kingdom of Saudi Arabia after an ADABAS performance consulting assignment in 1982, he had a vision for a company he would name Treehouse Software. Having established the company as the premier provider of ADABAS migration solutions, today Treehouse offers solutions for all mainframe data migration and integration requirements, business intelligence and analytics, and application modernization.

WHITE PAPER

Legacy Data Migration: DIY Might Leave You DOA

Best Practices for Implementing a Data Modernization Initiative White Paper

In any application migration/renewal project, data migration is usually a relatively minor component in terms of overall project effort. Yet failure of the data migration process can and does cause failure of the entire project.

This white paper touches on those most critical tactics and features of a comprehensive legacy data renewal strategy and highlights a relatively low-cost solution that mitigates the risk of having it run aground on data migration -- especially as compared to the effort, complexity and risk entailed in a "Do-It-Yourself" solution.

Jumpstart your legacy renewal project by reading this white paper first.

To download the White Paper, go to: www.treehouse.com/legacy-data-migration-white-paper.

WHITE PAPER

Data Modernization as the Gateway to Legacy Modernization

Best Practices for Implementing a Data Modernization Initiative White Paper

Strategies for approaching a data modernization effort vary, but regardless of approach, there are myriad benefits for the organization. Most methodologies have the end goal of deploying replication technology to disseminate information "locked up" in disparate legacy applications to low-cost, readily-accessible local databases, with low cost and minimal impact on IT and legacy applications.

This white paper describes best practices when undertaking a comprehensive legacy modernization project and provides insight to an approach that is a means to demonstrate success quickly and deliver short-term ROI.

Download your copy now and get started on the right path to implementing a strategy for legacy modernization. To download the White Paper, go to: www.treehouse.com/Data-modernization-white-paper.

requires a metadata/management layer to manage and collate field/column names, proprietary datatypes, etc.

Improvements in processor speed, cost/performance ratios and software sophistication have made data virtualization more capable than data federation used to be. Nevertheless, knitting together disparate data sources and attempting to harmonize them into seamlessness is subject to the age-old principle that the system is only as robust as its weakest component. Thus, any suboptimal aspect of the interface to a given data source tends to render the whole system brittle. When it works, it can work very well; when it breaks, it's a nightmare.

Let's take the example of a popular nonrelational legacy database, Software AG's ADABAS. Thousands of mainframe sites worldwide run high-performance, bet-your-business applications using ADABAS as the back end. ADABAS, which has been around for over 40 years, manifests an interesting feature: the result set returned to a query may contain data that has been updated but not yet committed—indeed, updates that may never be committed if they are subsequently backed out. Therefore, a client issuing a query that does not take this possibility into account may receive "invalid" results—results that, when issued only a short time later (after the backout), would be different.

Modern relational databases offer configurable "transaction isolation" to help assure accurate-to-the-client and consistent query results, but ADABAS has no concept of transaction isolation. This is the way the database has always worked.

A robust data replication system, such as is offered by Treehouse Software's **tcVISION**, avoids such issues by synchronizing each data source with a readily-available, consistent target database, first by doing a bulk load from source(s) to target and then by replicating only changes—only committed changes—from source(s) to target. So there can never be ambiguity as to whether a query against the target database involves uncommitted data.

By definition, data replication necessitates redundancy and increased storage costs. A couple of decades ago, storage was limited and expensive; today—with everyone talking about "Big Data"—storage is abundant and cheap.

tcVISION: Mainframe Data Replication and Integration

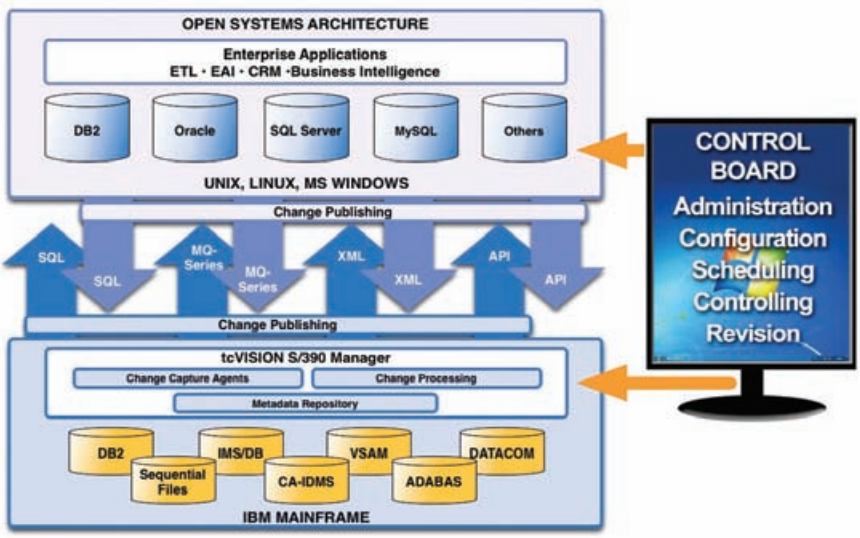

What's in a Word (or Two)

At Treehouse Software, we cite "data replication" as one of our key competencies. But the term—like so many others in the IT arena, including some that are mentioned in this article—suffers from lack of clear definition. In the early days of our flagship ADABAS-to-RDBMS solution, tRelational/DPS, we coined the terms "materialization" and "propagation" to refer to what are often dubbed "ETL" (Extract-Transform-Load) and "CDC" (Change[d] Data Capture). We have also characterized our multiple offerings in this space as "data migration" solutions, though we prefer "data transfer", since many of our customer implementations involve ongoing coexistence of source and target systems rather than one-time migrations. At times we have used "data synchronization", and even made that the basis for naming our DPSync product.

Searching the Internet for "data replication" not only yields results for Treehouse and our competitors, it may also return information on disk-mirroring hardware and the like.

Today we offer customers robust and comprehensive solutions that move data between a multitude of heterogeneous sources and targets, unidirectionally or bidirectionally, in bulk or in real time, within and across mainframes, open systems and Windows platforms throughout the enterprise. Until somebody comes along with a better idea, we'll continue to refer to these as "enterprise data replication" solutions.

Data replication solutions over the years have been criticized for being inconsistently implemented and difficult to manage. This was mainly a function of the tools used—and of the tools not used. We have long held that "our real-world customers are our biggest competitors", in that we often find that customers have done in-house development of their own data replication systems. Certainly, there is nothing magical

(continued on page 4)

about such systems; our customers employ skilled developers who are fully capable of creating them. And often the initial implementation works well, and becomes entrenched such that other solutions are resisted even when the in-house system must be extended beyond its original capabilities and starts to break down.

These days, choosing to implement a commercial data replication software system is a "no-brainer", as long as it is one that comprehensively covers—efficiently, natively and scalably—the full gamut of data sources and targets that the organization might reasonably contemplate using. There is enormous benefit to leveraging a single solution to consistently design, implement and manage replication scenarios and processes. Data replication becomes simply another service offered by the professional IT center.

Savvy IT organizations also recognize that data replication and data virtualization do not represent an either/or choice. Hybrid implementations can be effective: virtualize directly to the system of record where that can be done readily and manageably (and with acceptable performance impact on the source system), and virtualize to a replica where direct access is problematic or the workload is better handled on the replica.

Note that well-architected replication solutions not only incur minimal overhead on the source system, they also enable the "heavy lifting" of data transformation and application to target to be placed on the platform where capacity is most available and manageability is maximized. For example, **tcVISION** can be configured to replicate from a variety of mainframe data sources with only the change-capture process being executed on the mainframe and all other processing done in a "lower-TCO" environment.

We were recently working with a customer in the real world that is phasing out legacy mainframe systems. At the same time, however, the customer wants to move to a Service-Oriented Architecture (SOA). The two initiatives were in conflict, because exposing mainframe-based processes as services would have increased mainframe capacity needs, necessitating an unacceptably-costly upgrade. However, by replicating the mainframe data to a Microsoft SQL Server environment, services could be exposed and the processing contained to the Windows environment—where, ultimately, they would need to reside anyway, after the decommissioning of the mainframe.

(This customer scenario actually added another layer of complexity, since some of the services being contemplated were read/write in nature, meaning that updates would also have to flow from SQL Server back to the mainframe in real time. Fortunately, such a bidirectional synchronization scenario is fully and easily supported by **tcVISION**.)

The foregoing discussion has, to a degree, implicitly focused on "heterogeneous replication", synchronizing data from one data source to an unlike target, e.g., from CA-IDMS to DB2 LUW. However, many large enterprises have been doing "homogeneous replication" for years in synchronizing data to a remote "hot" or "cold" disaster-recovery/high-availability site, using products like E-Net Corp.'s Remote Recovery Data Facility (RRDF) and Enterprise Data Replicator (EDR) for mainframe data sources.

Today's sophisticated data replication products provide both heterogeneous and homogeneous replication capabilities, expanding the use cases and ultimately increasing the value of the customer's investment in them. In our real world, I recently visited with a customer overseas that faced a conundrum: critical business requirements could only be satisfied by implementing a new feature of the database (in this case, ADABAS), yet

tcVISION "Control Board" provides a single point of administration.

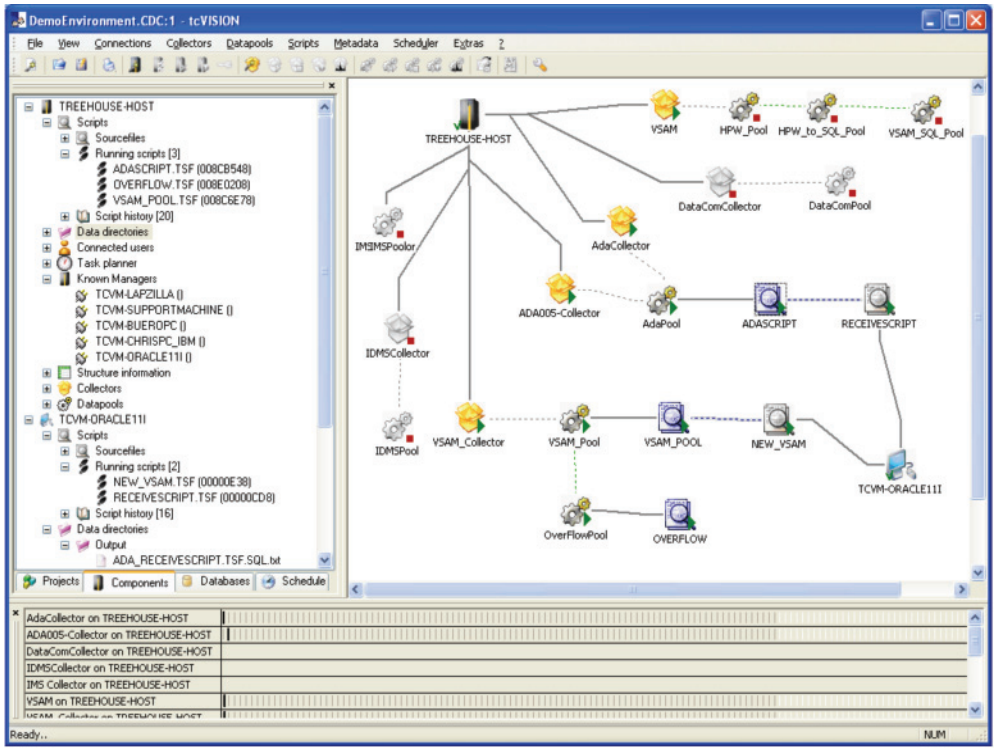

(continued on page 5)

this implementation could not be undertaken because it would necessitate a complete reorganization of the entire database—estimated to take weeks in duration, during which the database could not be used! Obviously, such an outage would not be tolerable. Yet, using a solution like **tcVISION**, database files can be migrated to a replica—a new ADABAS database, having the new feature implemented—in phases, and synchronized using replication techniques, with the process continuing for as long as necessary until all data are migrated and only a brief interruption to switch systems (a "zero downtime" cutover) can be executed.

Many data centers are facing the same demands for high availability. Using homogeneous replication involving a separate machine partition, or even another machine, enables extended maintenance windows to be scheduled and fault tolerance to be provided. In fact, our customer sees this scenario as a further benefit to the contemplated "reorganization replication": after the database reorganization is completed and the switchover to the new database is complete, they can replace the original database with a replica of the new one, providing a parallel failover environment in support of establishing 7x24 operations.

This approach can actually lend robustness to virtualization systems, in that a virtualization that involves "direct interface" to the system of record—as long as it can be configured to switch to the "mirror" environment dynamically and quickly—alleviates the possibility that the service using the virtualized data may find it "out of service".

While this issue represents the end of the line for the Treetimes newsletter and the "Real World Series", Treehouse Software will continue to serve our customers with real-world solutions. Watch our blog "The Branches" at www.treehouse.com/blog for continuing news, opinion and real-world case studies. I'm looking forward to adding a few of my own sprouts there. •

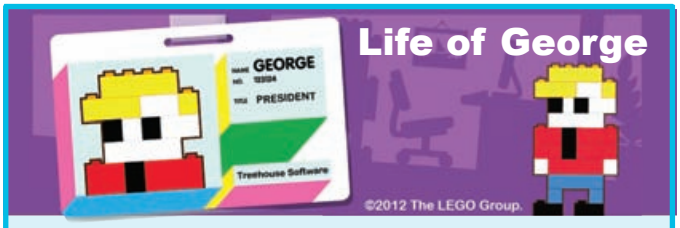

Snippets From the Treehouse President's Blog

"Yes, dogs were allowed in the office – and his master insisted this dog was smarter than the average COBOL programmer…"

Read more of George's Blog at:
www.treehouse.com/blog.

tcVISION Quick Facts

tcVISION is a product that enables enterprise real-time data replication through change data capture, while providing easy and fast data migration for mainframe application modernization projects.

- Provides modeling and mapping via a GUI modeler.
- Enables bi-directional replication between mainframe databases and DB2, Oracle, and SQL Server running on Unix/Linux/Windows.
- Change capture facilities identify changes in mainframe and relational data sources without programming effort.
- Transfers the changed data to target mainframe and relational databases and applications in real time, in time intervals or in response to events.
- Guarantees transparent and auditable data transfer between sources and targets.
- Provides powerful processors to perform fast, efficient and reliable transfers of mass data.
- Change Data Capture reduces costs: less processing time, less hardware and less human resources are needed for exchange of information.
- Reduces the amount of data that must be transferred between the different systems to an absolute minimum. Only data that has been changed is transferred--only data that is really needed will be moved.

tcVISION considerably simplifies mainframe data exchange processes. The structure of the existing mainframe data is analyzed by **tcVISION** processors, then automatically mapped to the target. The data mapping information is presented in a user-friendly and transparent format – even for users with no mainframe knowledge.

All information is maintained in a meta data repository hosted in a relational database, and can easily be made available to other applications. The Windows-based Control Board of tcVISION provides an easy-to-use facility to administer the data flow. **tcVISION** provides a variety of interfaces to allow seamless integration with ETL or EAI solutions.

It's Not the End of Treetimes, But a New Branch of Communication

by Phyllis Doran, Director of Marketing

As mentioned in the last *"Real World Series"* feature article, this issue represents the end of the line for the Treehouse Software Treetimes newsletter. However, in keeping with the times and advancing with new channels of communication, including social media trends, we are excited to announce that we will continue to serve our customers ongoing thought leadership, product and industry news, real-world case studies, and unabashed opinions via our new blog, *"The Branches"*.

According to a recent Forbes blog by **Ken Makovsky**, there are LOTS of good reasons for one or more members of company leadership to blog. These include:

- Humanizing the company
- Enhancing visibility
- Building credibility and trust
- Establishing industry expertise
- Promoting products and policies
- Addressing important issues
- Defending the company against its critics
- Generating leads … and business

Makovsky's blog also references *The State of Inbound Marketing in 2012*, a report published by marketing automation vendor Hubspot, which was based on a 2012 survey of 972 professionals familiar with their business' marketing strategy. The key takeaway was that social media use across the spectrum has grown significantly. In fact, 62% of companies surveyed said that social media had become more important in the past six months than any other type of direct marketing.

So while we are attempting to keep current with emerging trends, we also know how important it is to communicate with our customers. The blog allows us to do this more frequently, and it enables our customers and followers to participate in the conversation: the newsletter and other direct mail messages are one-to-many, but the blog is interactive, and allows for comments, questions, and dialogue among all its subscribers.

We invite our customers and *Branches* readers to visit the site often, but it is best to subscribe to *The Branches*, so that you are notified whenever a new article is posted. As we continue to build our subscriber base, we anticipate the posts will also become more relevant and colorful; after all, this is what social media is all about! In today's highly mobile, "linked", and "friendly" world, everyone has a voice and the methods to comment, tweet or share are continuously expanding.

Perhaps the silver lining to this new social media is in the color itself. As consumers and businesses alike have become more eco-conscious, we're greener than ever -- using less paper and relying more on digital methods to interact. While some of us will always prefer a pen and ink, it's time for Treehouse to branch out with our next phase of marketing, communications and customer dialogue. If you really need to hold something in your hands to read our newsletter, you can always print back issues of Treetimes on our site: **www.treehouse.com/treetimes-newsletter**.

In the meantime, we'll see you online! **www.treehouse.com/blog**. •

Business Intelligence—Anywhere

With **Cubeware** from Treehouse Software, you can easily build solutions that meet the unique information needs of different types of users throughout your organization. The powerful underlying technology allows you to integrate information from any data sources effortlessly and without any need for programming.

Since the software is intuitive to use, you can quickly build management cockpits and dashboards in different languages and customize your reports to reflect your corporate design. In no time at all, you can analyze data ad hoc, create your first reports, transform budgeting into a collaborative process and automate the distribution of reports throughout your business.

Thanks to **Cubeware**'s usability and scalable design, you can quickly extend the solution to meet your changing needs. This makes **Cubeware** a secure, worthwhile investment.

More information can be found on our website at: **www.treehouse.com/business-intelligence-analytics**.

From the Treehouse Blog: Are You a Data Hoarder? *by Phyllis Doran*

I moved recently. I've heard it said that moving is one of life's most stressful experiences...right up there with death. I would have to agree.

At some point during the process of evaluating every item I owned to determine whether to keep, toss or donate, it occurred to me that I could be the next star of The Hoarders reality show. It hit me especially hard when we unearthed the boxes that were tucked away under the staircase storage space that contained documents (or other essentials like my son's artwork from kindergarten) that I hadn't unpacked -- let alone touched -- since my last move over seven years ago. Over the years, I had accumulated a houseful of stuff -- all of which at some point in time had been essential.

It took weeks to accomplish packing it all. Each item that I was going to move to the new place had to be carefully wrapped, boxed, and labeled, minimizing any breakage or losses and the amount of time the movers would have to spend on the other side. Not only would this potentially save us hundreds of dollars, but would expedite the unpacking and ensure my stuff in the new digs would be well organized and clutter-free. Looking back, this was time well spent.

So what does this have to do with your data?

Just as stuff inside a home can be hoarded, so can data! One of the long-running complaints about corporate IT is how data gets “siloed”, which constrains organizations in acting across internal boundaries. Enterprises generate huge amounts of structured and unstructured data stored in various data repositories across the organization and in production systems. There are many reasons behind data hoarding, but if you have a data archiving strategy, you can ensure that inactive data -- especially that which may be inside legacy systems applications -- is stored, managed, secured or destroyed and that the data you keep can be accessed for any reason at any time.

For the purposes of this article, let's talk about inactive data that is rarely or lightly used (or may not be used at all except on occasion when someone needs to look something up). While there may be myriad methods to access this data, including writing customized SQL queries, these may vary greatly in quality and sharability and you may need training to support them. Maintaining a mixed bag of individually knocked out SQL queries may prove to be a headache. What is needed is an agile, simple solution.

Introducing **DataNovata**.

DataNovata is a Rapid Application Development tool that instantly generates a secure, read-only, web-based application to query and navigate any database in any conceivable way. Simple focused and cost-effective tool, **DataNovata** is a perfect way to get a standard architecture in place as applications are being retired or decommissioned while the data still needs to be available for various purposes—regulatory/statutory, audit, historical analysis, etc.

The feature-rich, web-enabled applications generated by **DataNovata** are suitable for end users and give them powerful data interrogation facilities to facilitate finding the information they need quickly—with minimal training and technical expertise required. Organizations are realizing enormous cost savings by leveraging the power of **DataNovata**.

DataNovata is used for a variety of purposes such as:

- Application Retirement
- Archived Data Management
- Enhance Access to Mainframe Data
- Satisfy Legal Requirements
- Platform and Application Rationalization
- Applications Portfolio Management
- Information Lifecycle Management
- Forensic Analysis
- Analytical Databases
- Application Renovation
- Testing Facility

Because of its universal use, **DataNovata** does not sit in any specific vertical market niche. However, its ability to provide users access to legacy data makes it well-suited for use by the financial, pensions, banking and insurance industries.

Now, if there were only something as easy as **DataNovata** that would help me unpack those last miscellaneous boxes, I could have my housewarming party.

Read more fun and informative Treehouse Software Blogs at: www.treehouse.com/blog. •

TREEHOUSE SOFTWARE, INC.
2605 Nicholson Road, Suite 230
Sewickley, PA 15143 USA

POSTAGE
INFORMATION
HERE

Treetimes

Treehouse Software Products & Services

Legacy Data Modernization:

DPS - ADABAS-to-RDBMS data materialization (ETL), replication, and propagation (CDC) software

DPSync - Real-time ADABAS-to-RDBMS data propagation (CDC) software product set

tRelational - ADABAS modeling, mapping, and data analysis tool; DPS parameter generator

tcACCESS - powerful integration platform for users of IBM mainframes, allowing a transparent integration of mainframe data sources and mainframe programs into open system applications

tcVISION - Data replication product that focuses on changed data when transferring information between mainframe and workstations or open systems

NatQuery - GUI-based tool that intelligently generates NATURAL code to handle all of the complexities of data extraction from ADABAS

NatCDCSP - Add-on to NatQuery designed to create immediately-usable data out of the ADABAS PLOG

Application Portfolio Management and Modernization:

DataNovata - A powerful software tool designed to rapidly build a software application, giving access to the data within any relational database in order to satisfy any user or statutory requirements.

Business Intelligence:

Cubeware - Everything needed for easy-to-use, self-service business intelligence for decision makers, managers, and departmental users

Services:

Consulting and Remote DBA Services - Our expert consultants can provide ADABAS/NATURAL performance analysis, tuning, and optimization; development, training and implementation; help with special routines or user exits; product installation, upgrades, and training; data transfer/integration; and remote DBA services for ADABAS/NATURAL on the mainframe and open systems

Software AG Related:

ADAMAGIC - Tool for converting mainframe ADABAS files into ADABAS for UNIX/Linux/Windows, flat file, or comma-delimited formats

ADAREORG - File reorganization tool for ADABAS

ADASTRIP - Data extraction utility for ADABAS

AUDITRE - Generalized ADABAS auditing facility

CHART for NATURAL - NATURAL application analysis and documentation tool

N2O - NATURAL application change management system

N2O/3GL - 3GL support within N2O for PANVALET, LIBRARIAN, ENDEVOR, and PDSs

PROFILER for NATURAL - NATURAL quality assurance and testing tool

SECURITRE - ADABAS and NATURAL security interface to RACF, ACF2, and TOP SECRET

TRIM - ADABAS and NATURAL performance monitor

Mainframe Emulation:

SEDIT - XEDIT and ISPF/PDF compatible editor for UNIX and Windows

S/REXX - REXX-compatible language for UNIX and Windows

S/REXX Debugger - Optional graphical debugger for S/REXX programs

Phone: 724.759.7070 **Fax:** 724.759.7067 **E-mail:** tsi@treehouse.com **Web:** http://www.treehouse.com