TREETIPS

A Publication of TREEHOUSE SOFTWARE

Issue #9 - July 1991

TREEHOUSE SOFTWARE ANNOUNCES DynaMacs!

For many years, the data processing industry has been plagued by a number of critical issues:

- Long software development cycles
- A shortage of adequate development tools
- A rapidly changing business environment
- The difficulty of adequately determining user requirements
- The need to reduce the maintenance overhead on systems
- The need to extend the life cycle of application systems
- Tight schedules impacting on the quality of the end product
- Redundant and obsolete documentation

Until now, there was no cost-effective, tailor-made solution for the NATURAL 2 community. Treehouse Software released DynaMacs to fill these needs. DynaMacs is a family of productivity tools which complements NATURAL 2 by providing a number of benefits, including:

- Automated Mainframe Based Prototyping
- Self-Documenting Systems
- Self-Constructing Menus
- Comprehensive Functional Access Control
- Consistent User Interface Across All Application Systems
- HELP Text Management
- Automatic JCL Generation
- Model-Driven Program Generation
- A Programmer's Copybook and Utilities
- A Standard for Functional Specification, Coding, and Documentation
- Practical CASE Tools

These facilities make application developers **more productive** by eliminating repetitive tasks in design, programming, and documentation. The use of templates, skeletons, and generated programs **dramatically reduces development time and maintenance problems.**

DynaMacs is a set of four products. These products are:

- DynaDoc- For creating self-documenting applications
- DynaMenu - For dynamic menu and access control
- DynaGen - For NATURAL 2 code generation
- DynaShell - For automated mainframe prototyping

These products are designed to **maximize the utilization of all work effort** in application development. For example, text entered during program specification using DynaDoc and PREDICT is reused (where appropriate) in user manuals and technical manuals. Maps created during prototyping are reused in the finished application and are automatically integrated into user manuals.

Every step has been taken to eliminate redundant effort in the application development process through the use of DynaMacs.

INSIDE...
More on DynaMacs
on pages 5-7!

Table of Contents

Editor's Sproutings

by Michael Salsbury

The TSI Technical Staff Gets Published

As our customers know, we have considerable ADABAS/NATURAL knowledge and experience. The media is just beginning to take notice of us. During the past year, we have had new product announcements in Datacenter Manager, 370/390 DATA BASE Management, Technical Support, Mainframe Journal, Software Magazine, and other respected publications. However, this will be the first year that **TSI employees will have technical articles published in these magazines.**

Look for "ADABAS/NATURAL Performance Monitoring" in the July issue of Technical Support magazine, "Change Management in the NATURAL Environment" in the August issue of 370/390 DATA BASE Management, and "Securing ADABAS and NATURAL Through CA-TOP SECRET, CA-ACF2, or IBM's RACF" in the October issue. "NATURAL Application Development Using Integrated Productivity Tools" will appear in a future issue.

Regional Meetings - Treehouse On the Road Again

We've been to several user group meetings in the past three months, including: Delaware Valley, New Jersey, Washington D.C., Sacramento, Denver, Salt Lake City, Hawaii, Alaska, the Big D Caucus in Dallas, and the NATURAL 2 Conference in Boston. Also, two of our staff had the hardship of traveling to England, France, and Germany to work with our affiliates there.

At some of these, we spoke about our products specifically. At others, we gave generic presentations. The response has been overwhelmingly positive. The major complaint is about the generic talks. Most of the attendees say they **would have gotten more out of a more specific, technical presentation.**

As a token of our appreciation of the support we have received from the University/College user community, we made a special offer to all CAUCUS members. All CAUCUS members ordering a trial of (or purchasing) one or more of our products before the end of May 1991 were **eligible to win a discount of 50%** on those products. In addition, TSI is donating **$1,000 to the CAUCUS General Fund** and $1,000 to the general scholarship fund of the winner's institution.

TSI has selected its winner, an institution trying several of our products. After we finalize the sale, we will publish the name of the winner in the next TREETIPS.

NATURAL 2 Conference

Boston University hosted this year's NATURAL 2 Conference. TSI had three attendees and gave two presentations. TSI's booth at the Vendor Exhibition was well received.

We've noticed a downturn in attendance at the SAG conference and at the University CAUCUS, but an increase at the NATURAL 2 conference. According to Dennis Hamilton, a member of the NATURAL 2 Conference Planning Committee, "Many more **users prefer to attend the NATURAL 2 Conference** than the CAUCUS or SAG International Conferences because they feel they walk away with a greater amount of information that they can use in their jobs."

In This Issue

This issue highlights **DynaMacs,** a new product offering from Treehouse Software. You probably received one of our recent pink flyers about it. By the way, the pricing for two of the tools was incorrect in the flyer. Switch DynaMenu and DynaGen and it will be correct.

We are pleased to present **three user-written articles,** two on DASD migration. It means a lot to us and our readers when people take the time to send us their articles and interesting information. Remember, you'll get $100 from us when your article is published!

We Need Help

The growth of TSI has caused the need for:

- sales/marketing professionals with technical background
- senior technical people versed in ADABAS, NATURAL, DB2, Assembler, CICS, CASE
- NATURAL applications programmers
- educators experienced in ADABAS, NATURAL, and DB2
- technical writer

Most of these assignments are for our headquarters near Pittsburgh, PA.

(continued on page 18)

Current Breezes
by George

DynaMacs

The popularity of this new family of products has kept three of us on the road giving presentations and demos. We cannot agree on why the popularity is so high. Some say it is the pretty pink color of the flyers. Others say it is the nostalgic feeling one gets looking at the 50s family on the flyers.

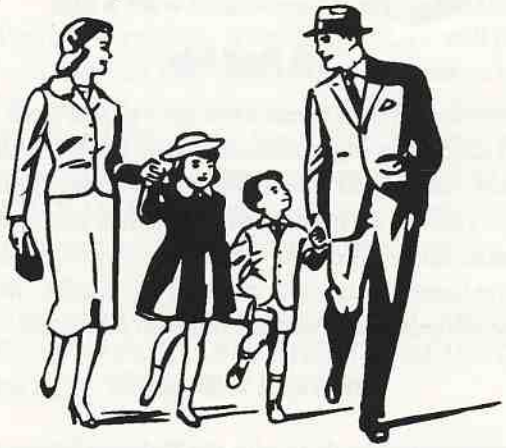

Why the old-fashioned family portrayal? Because the first argument in data processing history occurred in the early 50s with Eckert saying to Mauchly **"Who is going to document this mess?"** The cry is still heard today.

What does that have to do with **DynaMacs**? One of the family, **DynaDoc**, will go a long way toward documenting your NATURAL applications. Have we completely automated the boring task of documenting? No, we cannot create verbiage. But, **DynaDoc**, pulling things together from several sources, **produces high quality technical specifications, user manuals, and other documents.**

Of all the tools, **DynaDoc** is the easiest to install and use. It can document an application very quickly, printing the manual on your mainframe. Optionally, you can send us a tape from **DynaDoc**, and we will print it here on our Macintosh using Ventura. You will be impressed with the quality. Or, send us the source code for one of your NATURAL applications and its PREDICT information. We will run **DynaDoc** against your application to produce a user manual which will convince you to buy **DynaDoc** immediately.

N_2O

Strides continue to be made on N_2O. The **development staff of five** (that is pure development - not support or documentation) continues to keep N_2O far ahead of the competition. The heart of N_2O is in excellent condition, along with its central nervous system and other critical parts. We are adding some interesting and important appendages and features, such as a MOVE/COPY option, migration of PREDICT objects, and more utilities like the popular Program Compare, etc. N_2O will come to the next party with its "top hat" - **migration of 3GL objects**. You can believe that if we are biting off the 3GL piece, we must have the NATURAL side pretty well handled.

Would you like to talk to a user with 375 programmers(!) using N_2O daily, and totally dependent on N_2O? We can arrange it. Or maybe somebody who tried N_2O when it was primitive and now sings its praises? **Call me directly** and we will set you up with all the supporting evidence management needs, and all the necessary tools and training to get you rolling. And, if you do not like the product, **we will give you your money back**. We are that confident.

SECURITRE

Here is another winner. The people who have **SECURITRE** cannot imagine that other mainframe ADABAS users have data lying around unsecured. **SECURITRE** is a rock solid product. It is an easy sell from a technical standpoint. You have ADABAS data relatively unsecured, NATURAL applications uncontrolled, and no way to integrate any of it with ACF2, TOP SECRET, or RACF. **SECURITRE is a perfect match-maker**, getting these forces talking to each other. Speaking technical to the security administrator, **SECURITRE** interfaces with your single-rule base, provides user-id based security, handles the entire SAG line of products, and easily installs, i.e., it is a "sale".

From a marketing perspective, it is tougher. The DBA often resists handing over a part of the database domain (security) to someone else. (Good DBAs find that being relieved of the security function leaves more time for tuning and other important tasks.) The NATURAL Security Administrator fears losing a job.

(continued on next page)

Current Breezes

(continued from previous page)

Sometimes the corporate security administrator has not heard of us, or does not want more work. (Good, professional administrators relish the extra recognition and control.) And management? In 1991, management is computer knowledgeable, and places an increasing level of importance on securing their corporate information. They realize it's not a good idea to wait until 1993 to secure their data and applications. Your management needs to know about **SECURITRE**. Please tell them.

Some of the interesting things we are hearing include users who want us to build an ADABAS-based, NATURAL-accessed Security system, because they have no ACF2, etc. These people say, "**SECURITRE** sounds great, but unfortunately we have no system security facility. Why doesn't Treehouse just make one for us?" One site wanted to know if **SECURITRE** will work across all platforms. Management has dictated that security on the mainframe, VAX, and PCs should be identical! IBM's RACF on DEC's VAX? What is the world coming to?

TRIM and AUDITRE

New releases of **TRIM** and **AUDITRE** are available. Users appreciate that we still have **four people at work improving TRIM**. The new Batch Slowdown feature is well received. User input is giving us excellent ideas for improvements in this feature to make it "more dynamic" (if there's little TP activity, don't slow down batch), and selective, i.e., slow down (or do not slow down) certain batch jobs.

Some of the previously advertised new **TRIM** features have been delayed to place emphasis on some defensive measures in **TRIM**. The Environment Screens and Memory Display are two such measures. The Environment Screens allow the user (or us) to determine what different environments (**TRIM** version, CICS and batch link routines and exits, etc.) are in use, and therefore why certain statistical information may be lacking.

The Memory Display feature lets users check out ZAPs for proper application to the active ADABAS/**TRIM** environment. Users are guided through its use by our support staff. Users have a one-word comment for this feature, "**Slick**".

Manuals

The TSI policy is that any SAG user, SAG office, or SAG consultant may have a single copy of any of our manuals, free. **Simply call and ask for one.** We only ask that you do not copy them (the manuals are copyright protected) unless authorized in writing to do so. We also plan to call users to find out why they want a manual only. They should also want a trial.

TSI Field Trip

Seventeen of TSI's finest went on a recent field trip to the #1 chicken wing establishment in the USA. It's the ***Quaker Steak and Lube*** in Sharon, PA (sorry Buffalo). Before you can try their "Atomic" wings you must sign a release form! We also visited the largest shoe store and the largest chocolate store in the world. Most of our developers wear shoes, and eat chocolate.

The way it was, and is

Remember the big deal with the Belgian fellow trying to get his hands on ADABAS - for sale to the **Russians**? This happened about 12-13 years ago.

There were hidden microphones and other clandestine movements to get the guy behind bars. It made the papers and John McGuire was on T.V. A big advertisement appeared in computer publications, announcing that **"Even the Russians want ADABAS"**, complete with hammer and sickle! Ah, marketing at its finest.

Now there are hundreds of ADABAS sites behind what was once the Iron Curtain. We're not quite ready to hire someone fluent in Russian, but when they start queueing up for TRIM we'll start looking.

Next issue:

SAG company Shark revisited.

TREEHOUSE
SOFTWARE

DynaMacs Q&A

Q. How difficult is it to install DynaMacs?

A. DynaMacs is loaded using the standard NATURAL 2 INPL procedure. A single ADABAS file is loaded (empty) using ADALOD. The installation should take less than 30 minutes.

Q. Once physically installed, what about setup time?

A. There is very little. Your NATURAL will have to be at level 2.1.4 or higher. DynaDoc uses a remote job entry facility (e.g., RJE, NATRJE, CMSBATCH, etc.), but only to submit reports from within an online environment. The single ADABAS file we use contains all the "control information" DynaMacs ever needs.

Q. How does DynaMacs relate to NATURAL CONSTRUCT, ARCHITECT, PREDICT, and PREDICT CASE?

A. DynaMacs is a family of NATURAL 2 productivity tools consisting mainly of DynaGen, DynaDoc, DynaMenu, and DynaShell.

DynaGen is a lower-cost alternative for model-based code generation. CONSTRUCT models are frame-based and require a special control file. DynaGen models are source-based and easily adapt to a site's needs. It does not require an ADABAS control file. DynaGen makes extensive use of copycode, which allows model changes to be implemented by recompiling instead of regenerating programs. Code-reusability is made practical as code segments are generated directly into editor workspace.

DynaDoc makes use of PREDICT, NATURAL source code, help text, maps, and developer/end-user input in creating self-documenting applications. DynaDoc is the only tool with the ability to produce an up-to-date, typeset user manual at the press of a button.

By the way, DynaDoc can document programs generated by CONSTRUCT as well as programs generated by DynaGen.

DynaMenu offers dynamic menu and access control facilities, neither of which are available with CONSTRUCT or PREDICT CASE. (CONSTRUCT can generate hard-coded menu programs without any access controls. CONSTRUCT maintenance models have no option level access control, i.e., add, change, delete, select, inquire, etc.)

DynaShell offers automated mainframe-based prototyping using NATURAL 2 maps as input. This unique facility is not available with either CONSTRUCT or PREDICT CASE.

Q. How do I know when to use DynaShell, DynaMenu, DynaGen, or DynaDoc?

A. DynaMacs offers the tools to assist in the functional specification and development of applications, and the ongoing management of those applications. The chart at the end of this article indicates how the various modules support the systems development life cycle.

Notice that DynaDoc is used throughout the development and maintenance of an application to provide up-to-date documentation. In addition, all of the modules are used at varying times throughout development and all but DynaShell are used in the ongoing management and maintenance of the application.

Q. Is DynaMacs a methodology?

A. No! DynaMacs is not tied to any methodology. The productivity tools will support any methodology. DynaMacs will also allow you to develop your own methodology making it easy to institute your own coding and documentation standards.

Q. Does DynaDoc interface with PREDICT?

A. Yes! DynaDoc fully integrates with PREDICT, and enhances its functionality by providing high quality reports based on PREDICT data and XREF information.

(continued on next page)

DynaMacs Q&A

(continued from previous page)

Q. The chart mentions DynaHelp and DynaJCL. How do these relate to the four main modules?

A. DynaHelp is bundled in with DynaMenu and DynaDoc. So, if you purchase one of these, you get DynaHelp. Both end-users and developers will be able to easily create and maintain help screens for their applications. Help screen text would then be automatically integrated for reference manuals or user documentation.

DynaJCL is a powerful template based JCL generator included with DynaDoc, for purposes of submitting its own batch reporting. Programmers may use DynaJCL to submit batch jobs without having to code JCL.

Q. How do DynaMenu access controls relate to NATURAL SECURITY (NSS) and Treehouse Software's SECURITRE package?

A. SECURITRE is intended for use by DBAs and security administrators to grant access to ADABAS files, fields, and utilities, NATURAL libraries, programs, DDMs, etc. NSS also provides some NATURAL security and control. DynaMenu controls access *within* applications. Senior end-users, rather than DBAs, can efficiently grant and restrict access of their staff to applications. Various user-exits are provided to compliment site-specific requirements without compromising DynaMenu's ease of use.

Q. Do you have any sites currently using DynaMacs?

A. DynaMacs has been installed in a variety of sites, including:

- 3 Insurance Companies
- 2 Retailers
- 1 University
- 2 Software Houses
- 2 Manufacturers
- 1 Bank
- 8 Local Governments
- 6 Federal Agencies

Q. Must I purchase the entire family, or is DynaMacs available unbundled?

A. DynaMacs four main modules are available unbundled. However, purchasing the entire family at one time lowers the purchase price and probably keeps other costs down, such as training, extra installs later, etc.

Q. What training is available for DynaMacs?

A. TSI provides the following DynaMacs courses:

COURSE	LENGTH
DynaMacs Workshop	3 days
DynaMenu Workshop	1 day
DynaMenu for End Users	1/2 day
DynaDoc Workshop	1 day
DynaShell Workshop	1/2 day
DynaGen Workshop	1 day
DynaGen Model Writing	(Call)
DynaMacs Internals*	1 day

* For DBAs and senior developers. Covers user-exits, data models, and integration with other systems.

Additionally, Treehouse offers a full range of consulting services to aid customers in developing models with *their* coding styles, programming logic, and user interfaces.

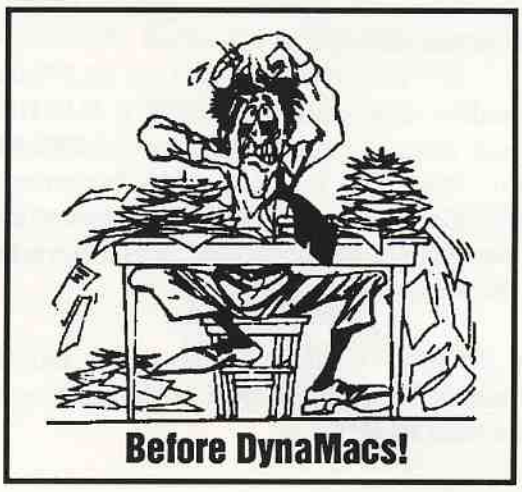

Before DynaMacs!

(continued on next page)

DynaMacs Q&A

(continued from previous page)

DynaMacs Overview Chart

	DynaDoc	DynaHelp	DynaMenu	DynaShell	DynaGen
System Analysis	Requirements Definition				
	Working Documents		Self-Constructing Menus		
System Specification & Design	Detail Spec: Maps & Reports	Initial Draft HELP text			
	Tech. Spec				
	Sign-off Document			Prototypes Generated from Spec & Maps	Develop Models (Only If Required)
	Generate Skeleton Programs			Concrete Model of the System	
System Building	Up to Date Documentation		Suggested Coding / Naming Standards		Generate Programs from Models
	Utility Programs & Subroutines				Library of Reusable Code
	Technical Documentation with XREF		Standard Copycodes		
	DynaJCL for Programs Submitting Batch Tasks		Tailor User Exits		
System Testing	Updated Tech. Doc.				Generate Active HELP Routines
	Functional Acceptance Document		Initial Access Controls		
System Implementation & Maintenance			Create User Access Profiles		
	Up to Date User Manual (Integrates HELP)	End-users Define & Maintain HELP Text	Create User Groups		
	Specifying Changes	Full XREF Capability	Maintain Access Controls		
	Changes Automatically Documented	Automatic Window Sizing & Placing	Temporary Suspensions		
	Ongoing Application Management				Model Changes via Copycode

TREEHOUSE SOFTWARE

N_2O Answers Customers Concerns

With the proliferation of NATURAL applications in many shops, Change Management is a growing concern. **N_2O** is Treehouse Software's Change Management System for NATURAL. **N_2O** features on-line and batch migration of NATURAL source and object programs across local and remote nodes, with multiple levels of authorization, audit trails, reporting, target compilation, XREF integration, automatic archiving, and checkout/checkin features. **N_2O certainly is a mature product** which every NATURAL shop needs.

Here are some questions typical customers have:

Q. We need a Change Management solution which is easy to use with online help screens. Will N_2O meet our needs?

A. Ease of use has been a goal in the design of N_2O and the product has been vastly improved since its creation in 1989. The N_2O interface is a series of online programs designed to make the setup and operation of the product as easy and trouble-free as possible. It consists of several types of menu-driven screens, PF-Key definitions, supports the use of direct commands, has an online HELP facility, and has an error-trapping and reporting mechanism.

Q. Our shop is dissimilar from most. We've got three machines, only two talk to each other, several NATURAL FUSER files, and lots of applications. We have our own unique naming standards. We don't have the NATURAL Security product, and don't want it. We don't want a Change Management system that will change our ways. What about N_2O?

A. Your shop is not so different. N_2O Batch Migration will cover your multiple machines, even the ones not communicating. You'll quickly discover that N_2O works for you. N_2O does not require:

- changes to current FUSER files
- other products such as NATURAL Security
- special repositories for your NATURAL source code
- changes to the way you develop and maintain applications

N_2O provides five user-exits that allow a site to tailor the way migrations are handled, for example allowing the N_2O administrator to reject individual programs from a migration, (for enforcement of **your** naming standards).

Q. What security measures are in N_2O?

A. N_2O contains its own access control system that can effectively control migration activity. NATURAL Security can be used to control access to N_2O itself. However, NATURAL Security is not a requirement.

Via SECURITRE, another product of Treehouse Software, it is possible to secure N_2O through ACF2, TOP SECRET, or RACF.

Q. Does N_2O provide change management across multiple platforms (i.e., mainframe, PC, and Wang)?

A. If a full implementation of NATURAL 2 and ADABAS exists for a particular platform, N_2O should operate on that platform.

Q. We are primarily a NATURAL shop now, but we still have some COBOL code. Will N_2O handle these along with JCL, PROCs, etc?

A. N_2O will provide specific support for 3GL languages in version 2.4. Initially, migrations will take place in batch. However, we intend to provide full functionality comparable to the rest of N_2O very shortly afterward.

Q. Does N_2O handle migration of DB2 plans and DBRMs?

A. N_2O is expected to include the capability to migrate these objects in version 2.4. The DBRM/Bind process will also be supported in N_2O version 2.4.

(continued on next page)

N_2O Answers...

(continued from previous page)

Q. Does N_2O provide flexibility for the online approval process (e.g., number and type of approvers tied to projects, group, type of change, etc.)?

A. Yes. The number and type of approvers may be specified through the DELAY parameter, which has three options: NONE, AUTH, and SERV. Migration Profiles determine the DELAY option in effect, along with other characteristics of migrations between these environments, such as batch or online, autocompile at target or not, etc.

N_2O uses Approval Profiles to specify which migrations a particular group of users is authorized to approve.

Q. Does N_2O provide object selection menus (i.e., allow selection from a list)?

A. Yes, and selected items are stored for future use, and can be modified. Selection lists provide extensive scrolling capabilities.

Q. To meet our auditor's requirements, all elements should have a "locking" facility to prevent concurrent changes. They also want ownership built into the system. Is this possible with N_2O?

A. Checkout/Checkin is a method of "locking" or controlling and monitoring changes in the development life cycle for a new program, or the maintenance life cycle of existing programs. It is designed to protect the integrity of a program throughout the cycle, and to provide an audit trail of the Checkout process.

When a program is selected for migration, N_2O will mark the program as checked out. The program may continue to migrate through other development environments (e.g., system test, training, quality assurance). The cycle ends when the program migrates back to the original "base" environment. N_2O will then mark the program as checked in.

Through the Checkout/Checkin facility, it is possible to control which users have access to the code for a given program or system.

Additional security is provided through the use of N2OEDIT. This is a front-end to the NATURAL EDIT module which follows Checkout/Checkin rules to determine which users may EDIT a given NATURAL program. The site determines which libraries are completely, partially, or not restricted.

Q. Does N_2O provide a facility for making emergency changes, giving notification to appropriate personnel?

A. N_2O provides a capability for making emergency changes or migrations. Notification that these changes have taken place will be supplied on audit reports which can be given to appropriate individuals. If EMAIL based notification is desired, the N_2O user-exits may be used to provide this capability.

Q. Does N_2O provide the ability to determine if other components are affected by (or required to complete) a change (i.e., use of PREDICT XREF information)?

A. Through the use of PREDICT XREF information, N_2O is able to determine if a given program being migrated affects other programs, or requires other programs for a successful recompile. Affected components may optionally be added to the migration request.

Q. We need to allow for changes to be scheduled (i.e., batch scheduling capability). Does N_2O incorporate such a feature?

A. Post dating of a Batch Migration can be done to ensure that program movement does not occur until a desired date and time. A post dated migration time is specified at migration request time. The selection program which is part of the batch process will not migrate programs in an event until after the date/time specified. CA-7 or another scheduler can be used to submit N_2O batch JCL periodically, looking for any eligible event.

(continued on next page)

N_2O Answers...

(continued from previous page)

Q. Can N_2O ensure that changes made to source members are also made to executable members (when applicable)?

A. The most reliable way to ensure that source and object match is to use the Automatic Target Compilation feature of N_2O. Thus, when source code is migrated, N_2O may be used to compile the migrating programs in the target environment.

N_2O version 2.3 will provide for comparing date and time stamps in source and object code to verify that a given object program reflects the changes made to the associated source code.

Q. Will N_2O delete migrated code from the source environment?

A. N_2O currently copies code from the source environment to the target environment. The "Move or Copy" option is available in N_2O version 2.3. When "move" is specified, programs will be copied from the source environment to the target environment, then deleted from the source environment.

In addition, N_2O will support "Deferred Moves" which allow the programs to be copied, then deleted at a later time (for example, a week later).

Q. We would like to reduce source library size by keeping deltas, rather than entire versions. Will N_2O do this?

A. N_2O contains a program compare function which identifies changes for programmers or auditors. This is the first step toward storing only program changes (deltas). The current archiving system allows site specification of the number of versions of archived programs to retain. We anticipate optionally archiving deltas in N_2O version 2.4.

Q. Does N_2O provide backout and restore procedures for a specific version of an application?

A. N_2O provides a complete archival and restore process for both online and batch migrations.

N_2O version 2.3 is expected to include the option to automatically recover objects from the archive when a migration attempt has failed.

Q. Notification of "approval needed" and "move completed" is important to us. How does N_2O provide such notification?

A. Through its reporting mechanism, N_2O allows approvers to list the events awaiting their approval. This list can be displayed online, or printed for later reference. Auditors and others may produce reports showing migrations completed. N_2O provides extensive reporting selection capabilities, allowing the site to generate reports based on criteria such as the following:

- Range of Events,
- From a NATURAL Definition,
- From a Library,
- To a NATURAL Definition,
- To a Library, or
- From an Event Date.

In addition, it is possible to see all programs migrated by a particular event, or the change history for a particular program. Many other types of reports are available. The format of the various N_2O information files is provided to allow the site to code any non-standard reports which might be desired.

Q. Can anyone view these reports?

A. The N_2O administrator determines which users are authorized to access the reporting facilities where audit trail information is available. All users may be granted access to this information if desired, or access may be limited to specific users or groups of users.

TREEHOUSE
SOFTWARE

SECURITRE News

Version 2.0 Released

SECURITRE 2.0 has been released and several shops are in the process of converting to this newest version. We have had very few problems with 2.0 and it corrected several minor problems from 1.1.1. We want to encourage all users to upgrade to 2.0 as soon as possible to take advantage of the improvements, help ease support issues and free up developers for further extensions to the product.

In **SECURITRE** 2.0, the functionality of ADABAS Security and Utility Security have not changed significantly. Even if you are not quite ready to implement the new Security for NATURAL feature, you can still upgrade to 2.0 since all of the security components for NATURAL can be implanted independent of the rest of **SECURITRE**.

Version 2.0.1

We have begun work on the next **SECURITRE** release. We are always looking for help with beta testing. If there are any **volunteers** lurking in the shadows, we would like to hear from you. Version 2.0.1 will include fixes for any minor problems in 2.0, and we are including several enhancements users have requested:

- STRFNR parameters will be viewable from the RTM.
- PURGE and SCRATCH will be supported as part of Security for NATURAL.
- STEPLIB support has been included to allow you to make the default steplib for a library something besides SYSTEM.
- It will no longer be necessary to set the NATURAL parameter RELO=OFF.
- Utility Security has been enhanced to optionally write out DDCARD and DDKARTE statements to a sequential dataset.
- A UTMODE parameter has also been added to Utility Security to allow utilities to run in DORMANT, WARN, or FAIL, independent of the MODE setting for file security.

As 2.0.1 wraps up, coding will continue on 2.1.0. The release date for 2.1.0 is still a little loose at this point, however we are looking at some time during the fourth quarter (of **THIS** year).

Program Pathing

We plan to add Program Pathing support in 2.1.0. We have mentioned Program Pathing in the past, but we feel it is worth mentioning again, because nothing else is quite like it.

Currently, if a user needs to update the Payroll file, the Payroll file must either be unsecured, or the user needs to have explicit access to update this file. One of the problems with giving a user update access to the Payroll file is that the user can now update it from ***ANY*** program. Often, great effort is expended to make sure the user is not allowed to run other update programs, or cannot break out of the application to get to an editor, build a program, and update the Payroll file, etc. It is often difficult to create an environment which guarantees the user could not possibly make unauthorized updates.

In order to eliminate this problem, we are implementing Program Pathing. Program Pathing will allow you to give a specific ***program*** access to update a given file, say, Payroll. The user is granted access to run the update program only. The user does ***NOT*** have access to update the Payroll file directly. Since the program has the capability to update the file, the necessary updating occurs. If the user tries to update the Payroll file by any other means, the access will fail.

Increased Flexibility, Security

We have received requests to permit even greater flexibility in pseudo-dataset names. In 2.1.0 you will optionally be able to include items such as JOBNAME, CICS TRANID or SMFID in the pseudo-dataset name. This will allow you to control if user FRED can access FILE 13 on the production database from CICSTEST, from transaction NATX, or even from a particular CPU.

We've met with **Computer Associates** and discussed **SECURITRE for DOS**. This, and a lot more features are scheduled for 2.1.0 and subsequent releases. Be sure to check the next TREETIPS for all the particulars. If you have any suggestions for future releases, be sure to give us a call. **You can talk directly to the developers.**

TREEHOUSE SOFTWARE

Minimizing ADABAS Increase Outages

This article was contributed by an author who wishes to remain anonymous.

Performing an ADABAS increase is a relatively simple thing for most folks. However, this simple exercise can be a source of irritation for users who demand very high availability. Using the standard method, a typical increase can take from a few hours to a few days, depending upon the size of the existing storage as well as the number of cylinders to be added. With an increasing number of users demanding near continuous availability, many technical folks are learning that they must use outage time sparingly and efficiently.

The increase method recommended by Software AG is outlined in the utilities manual along with the ADAFRM and ADADBS utilities. The recommended method is as follows:

1. backup the ADABAS you intend to increase
2. format the new storage specifying DISP=MOD and a primary storage request of 0
3. uncatalog the dataset being increased and recatalog specifying the old and new volumes
4. run an ADABAS with the INCREASE function

With the use of protection logs, the backup can be completed while ADABAS is up. However, the other steps must be done during an outage. For an increase of 3,500 data cylinders, the format alone could run an hour on 3380 type devices (your mileage may vary). And let's not forget the "implied" potential restore which is the invisible step 5. It is reasonable to believe that whenever you are requested to do a backup, there is some chance that you will need to do a restore. In this case, the need for a restore will arise if something goes wrong with your format job which damages your existing data. As a result, the actual outage required to perform this increase, assuming you do a parallel backup, is about an hour for the format, a few minutes for the catalog and ADADBS steps and some small chance of a multi-hour restore. When obtaining outage time from your users, you might ask for 70 minutes and mention that there is an ever so slight chance you might need substantially more or you could just request the 70 minutes and **hope that everything goes well.** In either case, an outage of this magnitude is **not acceptable for shops which demand high ADABAS availability.**

Fortunately, there is another method available which will facilitate an increase of any size, and can be be performed during an **outage of only about 1 minute.** The general process is to allocate and format, in advance, the storage you will be adding and then ZAP the associated VTOCs to get everything pointing as it should. Outage time is reduced to about a minute regardless of the number of cylinders being added since the only remaining functions are to recatalog with the new storage and run the actual ADADBS INCREASE job.

Let's assume that you want to add packs FRED01, FRED02, and FRED03 to ADABAS.DATA.PROD which currently exists on ADA101, ADA102, and ADA103. The following process will substantially decrease your increase time under MVS. If you run another operating system, make the appropriate modifications.

1. Run an ADAREP being sure to specify MSGLEVEL=(1,1). The ADABAS file information is not important but the physical allocation map and the JES messages are.
2. Allocate and format the new space on the desired volumes. Use a name similar but not identical to the storage you will be adding to and do not catalog. In this case, we will use ADAFRM to allocate and format ADABAS.DATA.PROD1 on volumes FRED01, FRED02, and FRED03.
3. On the volume containing the last extent of the dataset which is to be increased, dump the VTOC entry of the dataset. AMASPZAP does that quite nicely. In our example, the job would look something like this:

```
//STEP01          EXEC PGM=AMASPZAP
//SYSPRINT        DD SYSOUT=*
//SYSLIB          DD DISP=SHR,VOL=SER=ADA103,
//                DCB=KEYLEN=44,UNIT=SYSDA,
//                DSN=FORMAT4.DSCB
//SYSIN           DD *
  ABSDUMPT 0000000110 0000000130
/*
```

(continued on next page)

Minimizing ADABAS Increase Outages

(continued from previous page)

The above example presumes that the CCHHR address of the VTOC, which can be discovered through FILE AID as well as a host of disk utilities, is known but the exact address of the DSN in question, or offset into the VTOC, is not.

4. On all volumes containing the newly formatted dataset, dump the VTOC entry of each new extent. The same job used in step 3 can be used here; however, the VOLSERs will be FRED01, FRED02, and FRED03 and the CCHHR address will need to be adjusted for the specific volume. The ADAREP and VTOC dumps done thus far are required to set up the remaining steps. However, they will also be referred to in the event that you elect to "undo" this change or accidently fat finger. For this reason, it might be prudent to print and safely stash them until your increase is completed.

5. At this point, we are ready to alter the VTOCs on the new volumes. There are two things that must be accomplished. First, we must change the names of the new extents to match that of the dataset which is being increased. Since the operating system has this DSN allocated to the currently active ADABAS, we must make this change in such a way that the operating system does not know anything is going on. This may be accomplished using the following job:

```
//STEP01          EXEC PGM=AMASPZAP
//SYSPRINT        DD SYSOUT=*
//SYSLIB          DD DISP=SHR,DCB=KEYLEN=44,
//                UNIT=SYSDA,DSN=FORMAT4.DSCB,
//                VOL=SER=FRED01
//SYSIN           DD *
CCHHR 0000000112
VER 10 F1
REP 10 40
```

MVS is equipped with a self-defense mechanism which prevents ZAPs to the VTOC without confirmation from the console. Simply replay 'Y' to the message that appears. The CCHHR address for this dataset is obtained from the VTOC dump we took of this volume. Since the DSN starts at offset 00 in the VTOC, the address to be modified is quite simple to calculate. Once this job is done, the name of our dataset on volume FRED01 will be altered from ADABAS.DATA.PROD1 to ADABAS.DATA.PROD (F1 changed to 40) without complaint from the operating system, even though the DSN is allocated to an active job. Rerun this job for FRED02 and FRED03 after making VOLSER and address changes.

6. Offset 34 in each extent contains an extent sequence number which helps chain the extents together. The extents on ADA101, ADA102, and ADA103 contain 01, 02, and 03 respectively. However, the extents located on FRED01, FRED02, and FRED03 were allocated as a new dataset using DISP=NEW. As a result, they also contain 01, 02, and 03 respectively rather than the 04, 05, and 06 that we need. Again, we will need to ZAP the VTOCs to make the desired modifications.

```
//STEP01          EXEC PGM=AMASPZAP
//SYSPRINT        DD SYSOUT=*
//SYSLIB          DD DISP=SHR,DCB=KEYLEN=44,
//                UNIT=SYSDA,DSN=FORMAT4.DSCB,
//                VOL=SER=FRED01
//SYSIN           DD *
CCHHR 0000000112
VER 34 01
REP 34 04
```

Rerun this job for FRED02 and FRED03 after making VOLSER, CCHHR address and VER/REP modifications.

7. Offset 5D of each VTOC entry contains a bit which, when set, indicates that this is the last extent of the dataset. This bit is currently reset on ADA101, ADA102, FRED01, and FRED02. However, at allocation time, ADA103 and FRED03 each contained the last extent of their respective datasets. Since our desired result is to have one dataset spanning all six volumes, we want this bit reset on all volumes except FRED03. The following job will make the desired change:

```
//STEP01          EXEC PGM=AMASPZAP
//SYSPRINT        DD SYSOUT=*
//SYSLIB          DD DISP=SHR,DCB=KEYLEN=44,
//                UNIT=SYSDA,DSN=FORMAT4.DSCB,
//                VOL=SER=ADA103
//SYSIN           DD *
CCHHR 0000000112
VER 5D 82
REP 5D 02
```

Again, the CCHHR address will be obtained from the dump of the VTOC. Everything we have done thus far may be safely done hours or even days in advance. At this point, our storage is fully allocated, completely formatted and the VTOCs correctly

(continued on page 16)

DASD Migration Using RESTRUCTUREDB

By Bill Gohs, National Education Corporation

This article was contributed by
Bill Gohs
Database Administrator
National Education Corporation
Corporate Data Center
1751 W. Diehl Road, Naperville, IL 60563-9099
Phone (708) 369-3000

Our Environment

National Education Corporation (NEC) runs six databases for three divisions, primarily using ADABAS 5.1.6 with NATURAL 2.1.4.

One of these databases was spread over several 3380 disk drives, including one 3390 device which had been SYSGENed in 3380 track compatibility mode. Compatibility mode mimics the track size of 3380s, but wastes approximately 25% of the storage capacity. We wanted to correct this situation, but the only way ADABAS would function on these devices was to use a new and "unproven" ADABAS 5 zap for 3390 devices. Further, because 3390 devices were just released from IBM, we did not want to be the first site to place its Production databases on them. We decided to wait until the next available opportunity.

An Opportunity Arises

That opportunity came several months later when the lease had expired and DASD was swapped. We now had the room to move our most active Production database entirely onto 3390s in native mode.

Previously, we had migrated an ADABAS 4 database from 3350s to 3380s. This exercise took over 20 hours, not including backups, etc., because we had to unload and reload every file. With the new ADABAS 5 Utilities, it took just over four hours.

The process involved running an ADABAS RESTRUCTUREDB out to 3380 tape cartridges and running a store function back on to newly formatted 3390 devices (we used the 8390 dummy device type to get larger blocksizes). The RESTRUCTURE function to put our old database (5168 cylinders) out to tape took approximately three hours. It took just over one hour to store this database (now 4200 cylinders) onto a newly formatted database on the 3390 drives. The size of the database in MB and the available free space were unchanged (Actually, free space increased slightly due to the larger track sizes on the 3390s.).

Planning Ahead

Before we scheduled the conversion, we needed to know how much time we would need and how much space to allocate on the new database. We ran an ADAORD RESTRUCTUREDB on a small database when it was not active and estimated the time it would take for the real database based on size comparisons. We also converted this small database to the 3390s to be sure the MPM would have no problems.

Estimating size was somewhat more involved. Although we had the required sizes calculated in the RESTRUCTUREDB run of the small database, these did not seem accurate. Pages 8-32 and 8-33 of the Utilities manual give formulas for conversion based on the blocksizes of the various device types which seem to work out fairly well. Unfortunately, the RESTRUCTUREDB Utility will calculate only the space needed to store a file's data, but without any free space. In addition, the utility's estimate for UISIZE is too low to store this index. Fortunately, running the ADAORD STORE solves the problem by taking as much space for a file as was used on the old device type, including free space.

This still left us without the size needed on the new device before actually running the STORE. The following chart gives an after-the-fact look at the figures for one file:

	NISIZE	UISIZE	DSSIZE
Extents from Restructure Report	79800	1710	99900
Unused Block Count	32976	579	15571
Net Blocks Used (Restructure Report)	46854	1131	84329
Store Estimate Given by Restructure	25524	239	62674
Space Used by STORE Function	46284	991	72927

(continued on next page)

DASD Migration Using RESTRUCTUREDB

(continued from previous page)

Note how the actual space used on the STORE is close to the net number of blocks used from the Restructure Report, while the STORE estimate given is much lower. This estimate would remove all free space from the file. A pleasant side effect of the RESTRUCTUREDB and STORE is that all files are reordered, as if running a series of Reorder Files.

Summing up the RESTRUCTUREDB and STORE Functions

The RESTRUCTUREDB STORE estimate gives an incorrect estimate of the number of blocks needed, eliminating all free space in the process. There may not be enough time to run RESTRUCTUREDB to get the estimate anyway, since the database must be down during this operation.

The solution was to write my own program to estimate the proper sizes needed. This program lists space currently allocated to a file and calculates the space needed to load the file onto the new device type. The program leaves the unused space in the file. Totals needed in blocks and cylinders are also given.

Input to the program is an ADAREP listing as Work File 1 and a parameter card in CMSYNIN designating the old device and the new device. The device types will translate into block sizes for ASSO and DATA to be used in the conversion calculations later. The program only shows one device type for simplicity, but any number of devices can be used. One of the other routines reads the ADAREP page containing the current padding factors and places them in a table (255 files max). Then, the Allocated/Unused portion of the ADAREP are used to do the actual conversion calculations and print the results for each file. Various totals are added and printed.

RESTRUCTUREDB and STORE for File 33

```
****************
* FILE   33    *  - EXTENTS BEFORE REORDERING
****************

DATAFFAC = 3%, ASSOFFAC = 3%

TOP ISN = 2978740, MAX ISN EXPECTED = 3999315
```

I FILE I LAY- I OUT	DEV I TYPE	LIST I TYPE	ALLOC I SPACE (BLOCKS)	FROM RABN	TO RABN	UNUSED I SPACE (BLOCKS)
I ASSO	3380	AC	5987	76057	82043	0
I ASSO	3380	UI	1710	82044	83753	579
I ASSO	3380	NI	79800	83754	163553	32946
I DATA	3380	DS	99900	87618	187517	15571

```
A D A O R D

FILE    33 (SAS-GRAD-REQMNT ) RELEASED FROM ADAORD PROCESSING.
FUNCTION BEING PROCESSED WAS RESTRUCTUREDB.

 TO STORE THE FILE SUPPLY THE FOLLOWING SPACE PARAMETERS:

    NISIZE= 25524 BLOCKS          ASSOFFAC= 3%
    UISIZE=   239 BLOCKS
    DSSIZE= 62674 BLOCKS          DATAFFAC= 3%
    2967329 DATA STORAGE RECORDS PROCESSED.
    LARGEST RECORD HAS 990 BYTES.

A D A O R D

FILE    33 (SAS-GRAD-REQMNT ) RELEASED FROM ADAORD PROCESSING.
FUNCTION BEING PROCESSED WAS STORE.

 REQUIRED SPACE :

    NISIZE= 46284 BLOCKS          ASSOFFAC= 3%
    UISIZE=   991 BLOCKS
    DSSIZE= 72927 BLOCKS          DATAFFAC= 3%
```

RESTRUCTUREDB and STORE for File 38

Note Secondary Extent on Padding Factor Page

```
****************
* FILE   38    *  - EXTENTS BEFORE REORDERING
****************

DATAFFAC = 3%, ASSOFFAC = 3%

TOP ISN = 65, MAX ISN EXPECTED = 667
```

I FILE I LAY- I OUT	DEV I TYPE	LIST I TYPE	ALLOC I SPACE (BLOCKS)	FROM RABN	TO RABN	UNUSED I SPACE (BLOCKS)
I ASSO	3380	AC	1	39893	39893	0
I ASSO	3380	UI	7	39901	39907	1
I ASSO	3380	NI	7	39894	39900	2
I DATA	3380	DS	15	59256	59270	7

```
A D A O R D

FILE    38 (SAS-SCHOOLS     ) RELEASED FROM ADAORD PROCESSING.
FUNCTION BEING PROCESSED WAS RESTRUCTUREDB.

 TO STORE THE FILE SUPPLY THE FOLLOWING SPACE PARAMETERS:

    NISIZE=     5 BLOCKS          ASSOFFAC= 3%
    UISIZE=     1 BLOCK
    DSSIZE=     5 BLOCKS          DATAFFAC= 3%
    65 DATA STORAGE RECORDS PROCESSED.
    LARGEST RECORD HAS 713 BYTES.

A D A O R D

FILE    38 (SAS-SCHOOLS     ) RELEASED FROM ADAORD PROCESSING.
FUNCTION BEING PROCESSED WAS STORE.

 REQUIRED SPACE :

    NISIZE=     4 BLOCKS          ASSOFFAC= 3%
    UISIZE=     4 BLOCK
    DSSIZE=    10 BLOCKS          DATAFFAC= 3%
```

(continued on next page)

DASD Migration...

(continued from previous page)

ADAREP Report after STORE

FILE	NAME	BLOCKS ALLOCATED NI	UI	AC	DATA/CYL	BLOCKS UNUSED NI	UI	DATA/CYL
3	NAT2-FUSER	580	29	55	11068/92	347	23	2588/21
4	NAT2-FNAT	165	10	28	4310/35	39	8	403/3
9	NAT2-FSEC	403	17	23	109/0	310	12	28/0
11	GEN-TABLE-FILE	116	5	17	127/1	113	3	115/0
12	NAT2-AF	14	15	5	605/5	4	9	65/0
13	NAT2-FDIC	1720	191	23	511/4	1055	141	133/1
19	CHECKPOINT	5	3	3	20/0	5	3	19/0
20	SAS-CLASS-SCHED	1863	27	11	737/6	1819	24	683/5
31	SAS-DATES	379	8	31	194/1	121	2	22/0
32	SAS-GRADE-BOOK	28927	937	873	49275/410	9137	758	12308/102
33	SAS-GRADE-REQMNT	46284	991	3490	72927/607	20754	749	10252/85
34	SAS-LOG-CLASS	23	2	5	25/0	10	0	1/0
35	SAS-LOG-QTR-MOD	110	5	7	94/0	33	0	9/0
36	SAS-PHYS-CLASS	4743	81	42	4542/37	1612	55	47/0
37	SAS-PROGRAMS	23	5	3	25/0	8	2	8/0
38	SAS-SCHOOLS	14	14	1	10/0	9	8	4/0
39	SAS-STUDENTS	8350	452	130	9544/79	2756	398	427/3
40	SAS-TABLES	217	11	23	303/2	81	8	32/0
41	SAS-TRANSACTION	41	5	9	73/0	35	2	60/0
42	SAS-TRANSCRIPTS	34	5	9	36/0	77	7	77/7

Padding Factor Page from ADAREP after STORE

```
*********************************************
*                                           *
* CONTENTS OF DATABASE   197 (EDCOS-PROD-DB    ) *
*                                           *
*********************************************
```

FILE	NAME	LOADED	TOP-ISN	MAX-ISN	EXTENTS NUA	D	PADDING AA	DE	IND ACIS	
3	NAT2-FUSER	90.08.13	55974	63029	111	1	10	10	NNNS	
4	NAT2-FNAT	89.11.06	25583	32087	121	1	3	3	NNNS	
9	NAT2-FSEC	89.11.06	3572	26357	111	1	3	3	NNIS	
11	GEN-TABLE-FILE	89.01.15	13936	19491	111	1	3	3	NNNS	
12	NAT2-AF	89.11.06	2412	5729	121	1	3	3	NNIS	
13	NAT2-FDIC	89.11.06	18197	26357	111	1	3	3	NNIS	
19	CHECKPOINT	90.11.18	1350	3437	111	1	10	10	NNNN	
20	SAS-CLASS-SCHED	89.06.07	10221	12605	111	1	3	3	NNIS	
31	SAS-DATES	89.01.15	31078	35525	111	1	3	3	NNIS	
32	SAS-GRADE-BOOK	90.11.04	544728	1000457	111	1	3	3	NNNS	
33	SAS-GRAD-REQMNT	90.11.04	2978740	3999539	111	1	3	3	NNNS	
34	SAS-LOG-CLASS	89.01.15	3898	5729	111	1	3	3	NNNS	
35	SAS-LOG-QTR-MOD	89.01.15	6419	8021	111	1	3	3	NNNS	
36	SAS-PHYS-CLASS	89.06.08	43444	48131	111	1	3	3	NNNS	
37	SAS-PROGRAMS	89.01.15	1196	3437	111	1	3	3	NNNS	
38	SAS-SCHOOLS	90.04.13	65	1145	221	1	3	3	NNNS	
39	SAS-STUDENTS	90.01.14	157502	158147	111	1	3	3	NNNS	
40	SAS-TABLES	90.11.08	21777	26357	111	1	3	3	NNNS	
41	SAS-TRANSACTION	89.04.30	8961	10313	111	1	3	3	NNIS	
42	SAS-TRANSCRIPTS	89.04.30	5771	10313	111	1	3	3	NNNS	
43	SAS-MESSAGES	89.01.15	132	1145	111	1	3	3	NNNS	
44	SAS-SECURITY	90.01.14	30653	50423	111	1	3	3	NNIS	

Program Output Based on ADAREP

```
        DATABASE SPACE CALCULATIONS FOR CONVERTING 3380 TO 8390 DEVICE TYPE
                              197 (EDCOS-PROD-DB    )
FNR   NAME                3380 BLOCKS ALLOCATED          8390 BLOCKS TO ALLOCATE
                            NI     UI    AC     DATA       NI     UI    AC     DATA
  3   NAT2-FUSER          1000     50    93    15163      583     29    54    11213
  4   NAT2-FNAT            285     15    48     5905      166      9    28     4367
  9   NAT2-FSEC            696     31    28      150      405     18    22      111
 11   GEN-TABLE-FILE       200     10    29      175      117      6    17      129
 12   NAT2-AF               25     10     8      829       15      6     5      613
 12   NAT2-FDIC           2967    330    38      700     1728    192    22      518
 19   CHECKPOINT-FILE        5      5    11       32        3      3     6       24
 20   SECURITY-FILE          2      2     1        2        1      1     1        1
 30   SAS-CLASS-SCHED     3213     47    18     1010     1872     27    10      747
 31   SAS-DATES            654     15    53      266      381      9    31      197
 32   SAS-GRADE-BOOK     49875   1616  1496    67500    29055    941   872    49916
 33   SAS-GRAD-REQMNT    79800   1710  5987    99900    46488    996  3488    73875
 34   SAS-LOG-CLASS         41      5     8       35       24      3     5       26
 35   SAS-LOG-QTR-MOD      190     10    11      130      111      6     6       96
 36   SAS-PHYS-CLASS      8179    140    71     6223     4765     82    41     4602
 37   SAS-PROGRAMS          41     10     5       35       24      6     3       26
 38   SAS-SCHOOLS            7      7     1       15        4      4     1       11
 39   SAS-STUDENTS       14398    780   236    13075     8388    454   137     9669
 40   SAS-TABLES           375     20    38      416      218     12    22      308
 41   SAS-TRANSACTION       72     10    14      101       42      6     8       75
 42   SAS-TRANSCRIPTS       60     10    14       50       35      6     8       37
 43   SAS-MESSAGES           5      5     1       10        3      3     1        7
 44   SAS-SECURITY        1095     15    74      255      638      9    43      189
        .
        .
        .
TOTAL BLOCKS IN FILES: 363966          8433  18314  428359     212031     4917  10672  316766
TOTAL CYLINDERS IN FILES:              1371         3173                    1012         2640
AND FREESPACE NEEDED:                   374          250                     276          208
TOTAL CYLINDERS IN DB:                 1745         3423                    1288         2848
```

TREEHOUSE SOFTWARE

Minimizing ADABAS...

(continued from page 13)

point to each other with correct names. The only problem is that neither the operating system nor ADABAS know anything about it. To remedy this, the two remaining steps need to be run while ADABAS is **down**.

8. Although the VTOCs contain all the correct information, the catalog is as of yet unenlightened. You will need to uncatalog the dataset and recatalog reflecting both old and new volumes. IEFBR14 can be used to accomplish this task.

```
//STEP01      EXEC PGM=IEFBR14
//OLDDATA     DD DSN=ADABAS.PROD.DATA,
//            DISP=(OLD,UNCATLG)
//*
//STEP02      EXEC PGM=IEFBR14
//NEWDATA     DD DSN=ADABAS,PROD.DATA,
//            DISP=(NEW,CATALG),
//            UNIT=SYSDA,
//            VOL=SER=(ADA101,ADA102,
//            ADA103,FRED01,
//            FRED02,FRED03)
```

9. ADABAS keeps track of the number of RABNs he has in each of his components. In order to inform ADABAS that he has some new storage, you need to run a quick ADADBS INCREASE job specifying the new storage.

Once the INCREASE job is done, the only thing left to do is to bring ADABAS up and put him back into service. You might want to run ADAREP to satisfy yourself that the storage is really there. If you accidentally do something you did not mean to do, all your changes can be safely undone with the help of the ADAREP and VTOC dumps you produced earlier. However, you will find that accidents are few, generally harmless, and easily detected. Additionally, after practicing this procedure on your test database a time or two, you will find yourself looking forward to your next increase opportunity.

TREEHOUSE SOFTWARE

Altering ADABAS Block Sizes

By David Williamson, State of Utah

This article was contributed by David Williamson of the State of Utah Department of Administrative Services. David describes the move from 3380 disk drives to 3390 disk drives, and is primarily dealing with changing block sizes.

Why Alter?

There are several reasons to change block sizes. A few of these are better performance, changing disk types, or to more fully utilize the available space per track. While any one of these are good reasons to change your block sizes, the focus of this discussion will be on changing from 3380 type disk drives to 3390 disk drives. The steps required to change block sizes are similar no matter what the reason.

Under ADABAS 411 the only way to go from one block size to another was to unload each file individually and reload it into a new database that had been defined using the new block sizes. This wasn't bad if you had a small database or only a few files. However, if you had a large database or many files, it became a very long process.

ADABAS 5 Improvements

With ADABAS 5, a new utility, **ADAORD** was added. Two of the functions in this utility are designed to make the process of changing block sizes much easier. The **RESTRUCTURE** function creates a sequential file from the data on the database. The **STORE** function reads the sequential file created by the restructure function and adds the whole database or selected files into another database with the same or different block sizes.

While this is a **big improvement over Version 4**, there are some fine points about the process you need to understand to make the process completely successful. The **security and checkpoint files will not be processed** by the utility even though the restructure function prints out messages about these files as if they are being processed. The store function must be run against an existing database so if the move is to a newly allocated database you must run not only ADAFRM, but you must also define the database with ADADEF.

If you are not using ADABAS security, you don't have to worry about the security file. However, if you are using ADABAS security, you will have to use ADALOD to allocate the security file and rebuild the data in the file yourself. The ADAULD and ADALOD utilities will not process the security file. Those sites which maintain the ADABAS security information in some repository besides the security file will have an easy time. Those who don't may find this to be the most difficult part of the whole process.

Checkpoint File

The ADAULD and ADALOD utilities will not process the checkpoint file either. This is not quite as difficult a problem as the security file but it does need some attention. The ADABAS checkpoint data on the checkpoint file is not accessed by any utilities. The only use for this information is if you need it to create some control cards to do some sort of recovery. This means that all you really need is a hardcopy listing of the checkpoints which can be produced by running ADAREP with the parameter CPEXLIST.

If you are storing user data on the checkpoint file, you need to decide if it is worth the trouble of copying the user data off and restoring it to the new file or if you can get along without it. To copy the user data you will need to write two small programs (**the source listings are available for the asking from TSI**). The first program V5ETREAD reads the user data records from the original checkpoint file using the RE command and writes the user-id and the data to a sequential file. The second program V5ETWRIT reads the sequential file and issues an OPEN command using the user-id obtained from the sequential file. A CLOSE command is then issued using the user data obtained from the sequential file. The OPEN and CLOSE commands are repeated for each record on the sequential file. Note that even though the names for these programs are similar to those used in the ADABAS 4 to ADABAS 5 conversion, they are different. The user data is obtained differently from a Version 5 database than from a Version 4 database.

(continued on next page)

Altering ADABAS...

(continued from previous page)

The Steps

Now that you know what to expect, you can actually go through the process. Following are the steps required to move a database (including security and checkpoint files) from 3380 to 3390 disk drives.

1- ADAFRM allocate and format with new dataset names on 3390 disks
2- ADADEF define database and checkpoint file
3- ADAREP database report with CPEXLIST for database on 3380 disks
4- V5ETREAD copy checkpoint file user data to sequential file
5- Bring down the database
6- ADAORD run using the RESTRUCTURE function, MODE=SINGLE
7- Rename 3380 database datasets
Rename 3390 database datasets to old 3380 names
8- ADAORD run using the STORE function, MODE=SINGLE
9- ADARUN CARDS change DEVICE=3390 or ZAP the default; calculate and change PLOG number of blocks
10- Start the database
11- V5ETWRIT reload user data to checkpoint file
12- ADALOD load empty security file
13- Rebuild data in security file
14- ADAREP database report

By allocating and formatting the new 3390 database with new dataset names, this step can be accomplished while the 3380 based database is still up and running. For a large database **this can save a significant amount of time.** Another benefit of using new names and renaming the datasets in step 7 is that if for any reason you have a problem with the move, you can go back simply by renaming the datasets to their original names.

It took approximately 5 hours to move a database consisting of 1768 cylinders of associator and 2984 cylinders of data storage. **Two years ago I spent over 20 hours** to move the same database using unloads and reloads. However, two years ago there were fewer records in the database.

In ADABAS 5 Software AG has provided what they call pseudo device types. These are device types that start with an 8. For example, the pseudo device type for 3390 disk types is 8390. The pseudo device type is an easy way to use larger block sizes. For the 3390 disk the 8390 pseudo device block sizes use an additional 3 percent of the available space on each track. Another benefit of a larger block size is that more records will fit in a block so you should do fewer I/Os to accomplish the same task. It will also increase the maximum size allowed for a record. If you are planning on moving from one disk type to another anyway, you should consider using alternate block sizes (pseudo device types).

I used the 8390 pseudo device type in my move from 3380 disks to 3390 disks. I have seen the **buffer efficiency increase** from approximately 12 to 15 while using a buffer pool of the same size. The number of commands processed in the last threads has dropped considerably. The **average duration** for commands **has decreased.**

I had reservations about going to 3390 disks because they hold so much data on a single pack, and I was afraid that performance would suffer. However, **I have been pleased with the results** I have encountered.

TREEHOUSE SOFTWARE

Editor's Sproutings

(continued from page 2)

Please Keep Up With Us!

Since each release of a product contains fixes as well as enhancements, it is important for customers to upgrade to the current version as soon as possible. This will often eliminate a problem you may be experiencing. Upgrading is also important because we do not support the previous release of a product as easily once newer releases are available.

The following chart displays the product versions you should be running.

Treehouse Product	Version(s) Supported
AUDITRE	1.1.0 (1.2.0 beta)
AUTOLOADER	1.2.0
DynaMacs	3.2.0
N_2O	2.2.1
SECURITRE	2.0.0
TRIM	4.0.0 / 5.0.3

TREEHOUSE SOFTWARE

ADABAS/NATURAL Publications

NATURAL 2 Tips & Techniques Answers your questions on the how and why of NATURAL functions along with the essential do's and don'ts privy only to the NATURAL 2 guru. There's info for the NATURAL Administrator; plenty on data structures with examples; a section on enhancements; the Map Editor; batch NATURAL processing guidelines and Structured Programming notes. 210 pp.

NATURAL 2 Developers Handbook has taken the best from our training manuals and combined it with plenty of hands-on experience into a very well rounded guide to using NATURAL. We've detailed the editors with plenty of examples and workthrus for each facility. A perfect companion to the Tips & Techniques publication. It comes with an excellent set of appendices with working examples. 476pp.

NATURAL 2 Self Study Guide Is a basic working guide to using NATURAL 2. The Study Guide covers all the editors and a very large part of the programming statements with hands-on workthrus. This NATURAL 2 Study Guide is intended to get entry-level programmers, analysts, and anyone interested in working with the leader in 4GL languages started with NATURAL 2 fundamentals. 402pp.

Inside ADABAS Is designed for the developer of applications utilizing ADABAS. Associator and Data storage are detailed, file and field structures are examined, commands are explored with an eye towards performance. Reviews of the data dictionary, utilities and security, and reporting and performance monitors are included. The appendices include ADASQL and Direct Call programming to round out the developer's background in ADABAS. 226pp.

Introduction to PREDICT Takes the reader through a session building a conceptual file introducing the editor, file and field concepts with chapters on verification rules, PREDICT menus, keywords, documentation, object retrieval and direct commands. 194pp.

All these books are bound, soft-cover editions in an 8 1/2 by 11 inch format. Each is complemented by a comprehensive index and the last three have exercises intermixed to better reinforce many of the concepts presented.

European customers interested in these books should contact MaK DATA SYSTEM, the Treehouse Software affiliate in Germany. Others interested in purchasing these books should contact WH&O. For ordering information, call, FAX, or write:

WH&O International
(A division of Meristem Systems Corp.)

P.O. Box 82355
Wellesley, MA 02181
Telephone: 1-800-553-6678
Fax: 1-617-239-0827

Publication of this advertisement does not indicate support or endorsement by Treehouse Software or its affiliates.

SHOE Reprinted by permission: Tribune Media Services

Treehouse Software's Affiliate List

Area Served	Affiliate	Phone Number
Africa	Syntactica	27-11-880-4976
Australia	Computer Consultants Australia	61-3-416-3377
Japan	Nissho Iwai Corp.	81-3-3588-4598
France	Fairware	331-39-65-0688
Germany	MaK DATA SYSTEM	49-431-3995-134
Brazil, Argentina	MSS do Brasil Ltda.	55-61-225-1158
United Kingdom & Eire	OCS Software Limited	44-81-207-5434

ADABAS, NATURAL, NATURAL Security, PREDICT, PREDICT CASE, and CON-STRUCT are all products of Software AG. VAX is a product of Digital Equipment Corp. (DEC). CA-7, CA-ACF2, and CA-TOP SECRET are products of Computer Associates. DB2, RACF and CICS are products of IBM. Ventura is a product of Xerox.

NFV

Treehouse Software products and services include:

DynaMacs - a family of productivity tools for NATURAL 2
N_2O - a NATURAL application Change Management System
SECURITRE - a Centralized Security Administration package for ADABAS/NATURAL
TRIM® - an ADABAS/NATURAL performance monitor
AUDITRE - an ADABAS auditing tool
AUTOLOADER - an ADABAS file automatic unload/reload/dump utility

Consulting on ADABAS, NATURAL, and associated topics

Customized on-site classes on a variety of ADABAS/NATURAL related subjects, including:
ADABAS (Concepts and Facilities through Advanced Topics)
ADABAS Direct Calls
NATURAL 2 (Beginning, Intermediate, Advanced, and End User)
PREDICT
COM-PLETE
TSI Products (DynaMacs, N_2O, SECURITRE, TRIM, AUDITRE, AUTOLOADER)
Custom Classes on ADABAS/NATURAL Designed to Meet Unique Needs

Phone:
(412) 741-1677

Fax:
(412) 741-7245

Treehouse Software, Inc.
400 Broad St., Suite 206
Sewickley, PA 15143

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