IUA CONFERENCES ADVISORY COMMITTEES UCL

CONNECTIONS

Spring 1996, No. 32

YEAR 2000 Support in CA-IDMS: A DBA's Perspective

by Chris Hoelscher, Hoelscher Consulting

YEAR 2000 support (also referred to as "century support") is on the minds of many MIS shops these days. Both application developers and database administrators have a huge part to play in completing a successful conversion to reflect the addition of the century to existing dates within the database, application code and non-database files. Many decisions need to be made by each MIS shop. While there are few clear-cut "right" vs. "wrong" solutions to these decisions, I would like to share Liebert Corporation's solutions as a guide to approaching YEAR 2000 support.

Liebert supports four application databases: manufacturing/accounting/order entry (CA-CAS release 1.3X, no maintenance, heavily modified); general ledger (CA-CAS release 1.3X, also no maintenance, not as heavily modified); bill-of-material archive; and order-tracking. Each application database exists in Production, QA (full production size) and Development (10-15% production size).

Our first "policy" decision was to establish a project timeline. Due to business requirements and staffing resource requirements, we chose December 21, 1996 through January 1, 1997 (a normal plant shutdown period) as the planned YEAR 2000 production implementation window. However, additional high priority projects have pushed the production implementation dates to July 4-7, 1996

Another decision concerned what the format of the "new" date fields should be. One potential solution was to convert all 6-byte YYMMDD fields to 6-byte pseudo-Julian format dates (CYYDDD). Another was to create new 8-byte element names (of which the first two bytes would be filler, the remaining six would be existing date elements defined as subordinate). In the end, we decided to leave all date elements with their current name and expand the size from six bytes to eight bytes (two bytes to four bytes for year-only fields).

Our next decision was where to make the changes. We opted to create a "clone" database, dictionary, PDSs and other non-database files for our CA-CAS 1.3 manufacturing environment. This requires that as

Continued on page 3

INSIDE THIS ISSUE

IUA Advances into Cyberspace.....2

IMC96 at CA-World '96 Is Your One-Stop Resource for "The Future of Information Technology".....3

TAC Tip of the Trade.....6

IUA 1995 Annual Report....7

You Got Your SQL in My DML!.....9

CASAC/ADAC Update and Survey.....10

Monitoring Run Unit Progress Tip of the Trade.....13

Pass it on!

1 _____

2 _____

3 _____

4 _____

5 _____

Look inside for your first glimpse at CA-World '96!

IUA Connections

IUA Connections Chairperson
Terry Wright, Ross Laboratories

Technical Consultant
James Bradshaw, Clemson University

Technical Consultant
Karen Langley, Lockheed Martin Energy Systems

Technical Consultant
Janet Schwartz, Alamo Rent-A-Car, Inc.

Technical Consultant
Joe Wodushek, Waukesha Engines

Editor
Kathleen Kennedy, IUA Staff

Editorial Assistant
Carrie Collins, IUA Staff

CORRESPONDENTS

CASAC/ADAC
Humbert Piscitelli, Loral Defense Systems - Akron

EUDTAC
Carol Silberstein-Locker, First InterState Bank

IUA Board
James Winn, Williams International

TAC
Chris Hoelscher, Hoelscher Consulting

Secretary/Treasury/Finance/VUTAC
Karen Langley, Lockheed Martin Energy Systems

Special Correspondent
Dick Brinson, DBSoft

Special Correspondent
Ken Reiff, Computer Associates (CA)

Special Correspondent
Walter Remski, Williams International

IUA Connections *is a quarterly publication of the Information User Association (IUA), designed to promote its members' objectives.* IUA Connections *is not responsible for the opinions expressed by its writers and editors.*

Information User Association
401 North Michigan Avenue
Chicago, IL 60611
312/321-6827
Fax 312/527-6783
Internet: iua@sba.com

Board James Winn

IUA Advances into Cyberspace: New World Wide Web Home Page Provides Direct, Immediate Access

The IUA is beginning to wrap up a very successful 1995-96, during which new and exciting initiatives were implemented, some of which I would like to share with IUA's membership. Recognizing members' need to access timely information on an ongoing basis, we are proud to announce IUA's new Home Page, which will be up and running soon. Cybersurfers can tap into IUA Headquarters via the World Wide Web and easily access membership information, upcoming meeting information and important dates, IUA Board updates and volunteering opportunities. IUA's Home Page is also a great opportunity to reach an expanded audience, as well as a targeted forum to advertise IUA member benefits. Watch for IUA's new World Wide Web address in upcoming issues of *IUA Connections* — and conveniently access IUA Headquarters directly 24 hours a day, seven days a week.

IUA Membership Directory

The 1996 IUA Membership Directory recently mailed to IUA's expanded membership; if you have not received your copy, or if IUA Headquarters needs to update any information listed in the directory, please contact IUA Headquarters at 312/321-6827. Intended to become your resource for sharing information and partnering among the CA community, the directory's new format includes first-ever advertising from vendors. IUA would like to thank the following vendors for their continued support of the IUA:

- Allen Systems Group
- ARCH Consulting Associates
- Computer Associates (CA)
- HSL

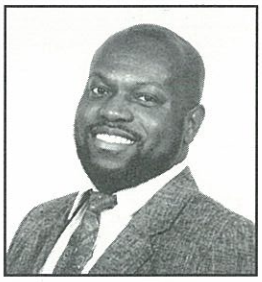

James Winn

- Harold M. Harrison
- TACT Software, a division of The A Consulting Team, Inc.
- The Project Group (TPG)

Develop New Skills and Volunteer for the IUA

IUA members do not need to invest a lot of time to receive the benefits acquired when volunteering for the IUA. In addition to enhancing your company's investment, volunteering results in: opportunities to network with CA technical representatives and influence product direction; peer-to-peer networking events; and advance news of the latest technology trends and CA product information.

Volunteering for IUA Advisory and Administrative committees is a viable means of receiving the greatest return on your investment! Please contact IUA Headquarters if you are interested in sharing your experience and knowledge.

CA-World '96

CA-World '96, August 25-29, in New Orleans, guarantees unparalleled opportunities to learn the direction of Information Technology (IT). And with keynote speakers Charles B. Wang, Chairman and Chief Executive Officer of CA, Bill Gates, Chairman and Chief Executive Officer of Microsoft Corporation, and General Colin L. Powell, USA (Ret.), this is a conference you can't afford to miss. Please refer to the Summer 1996 issue of *IUA Connections* for extended highlights and detailed registration information. With IUA Workshop '96 and CA-World '96, this is proving to be a busy and productive year for the IUA!

YEAR 2000 Support

Continued from the Cover

"normal" changes are made to the production "live" environment, they are also made to the "clone" environment (this decision dictates that often-modified entities will be the LAST to be "YEAR 2000 converted," since we did not want to re-apply YEAR 2000 changes repeatedly to the same entities because of repeated "normal change" overlays). The end result is that the "clone" dictionary and PDS datasets will replace their production counterparts at the end of the conversion style. The other databases were deemed stable enough to allow YEAR 2000 changes to be handled as normal production modifications.

The last decision was whether to change all date fields, or only those deemed "critical" (calc keys, sort fields, etc.). We chose to update all date fields for two reasons: we felt that the additional effort needed to update all date elements vs. "critical" date elements was minimal compared to the effort needed to DETERMINE which elements were "critical" vs. "non-critical;" and that as business needs (or performance demands) change, previously "non-critical" date elements may become "critical" date elements, and the whole process would be repeated over and over again.

Having made these policy decisions, I proceeded to develop a methodology to complete the database portion of this project. The remainder of this discussion describes how Liebert Corp. has implemented and will continue to implement YEAR 2000 support in a CA-IDMS database.

My first task was to identify schema record elements that should be enlarged from either two bytes to four bytes (year-only fields), or six bytes to eight bytes (Gregorian-format dates). While we have a few 7-digit Julian-format date fields, we chose to modify those in-place from OOYYDDD to CCYYDDD, not requiring physical restructure. We identified candidate schema record elements as follows: As elements with the phrase "DATE" or "YEAR" in the element name were selected, as were all elements that were defined with a usage of display and a length of two or six bytes. We ran CULPRIT reports against the database records, looking at the candidate date fields, and eliminated most 2- and 6-byte elements as not being dates. We discovered, however, that many year-only and full date elements do NOT contain the phrases "DATE" or "YEAR" in their name.

After the elements were identified, I modified them from either two to four bytes, or six to eight bytes. Had these YEAR 2000 updates been occurring in a "live" dictionary, I would have tagged them with an attribute or user-

Continued on page 4

IMC96:

CA-World '96 Is Your One-Stop Resource for "The Future of Information Technology"

by Ken Reiff, CA

The Information Management Conference (**IMC**) at **CA-World '96** is just around the corner, and, if you attended last year's conference, you know this is an opportunity that cannot be missed. Computer Associates (CA) and the IUA are once again planning an outstanding conference that will provide you with the information and education you need to make your CA-IDMS systems a success. This year's conference will be held **August 25-29** in **New Orleans**.

From database management to desktop development, IMC has the sessions that give you the latest information on the newest technologies you can use. The IMC is divided up into four distinct tracks, and within those tracks, into four separate threads. The tracks are:

Industry Trends
- gives you an idea of where our industry is headed and how it can affect your business.

Mainframe
- provides all of the material associated with mainframe topics.

Desktop
- devotes an entire track to this platform because desktop tools have become the driving force in Information Technology (IT).

Special Interest Groups
- contains all of the Special Interest Group sessions planned for the week.

The four threads within each track are: **Open Access**; **Application Development**; **Communications**; and **Data Management**. This makes it easy to locate the areas of interest and pinpoint the sessions you will want to attend.

Additional features and benefits of attending CA-World '96 include: a world pass to all of the 14 conferences being held that week; admission to the World Resource Center; pre-conference education from companies such as Amdahl, Brandon Systems, IBM, Learning Tree, Microsoft, Netscape and Novell; and special keynote addresses by Charles B. Wang, Chairman and Chief Executive Officer of CA, Bill Gates, Chairman and Chief Executive Officer of Microsoft Corporation, and General Colin L. Powell, USA (Ret.).

To register for CA-World '96 or for more information, please call 800-CA-INFO96 (224-6369) or 516-342-6600. Or visit us on the World Wide Web at: http://www.cai.com. And remember – register before May 5, 1996 to receive a discount of $200 off the registration fee.

YEAR 2000 Support

Continued from page 3

defined comment. Since YEAR 2000 updates are the only ones occurring in this dictionary, I can rely upon the DATE-LU-058 to determine which are YEAR 2000 modified elements.

Next, I created new IDD records to reflect the modified date fields. IDD does provide a nice easy way to clone as IDD record:

Add record prod-structure 2000 same as record prod-structure v 2.

I then replaced the appropriate record element (not necessarily the root element modified, but the highest group-level element to which the modified element is a subordinate):

Modify record prod-structure v 2000.
Rep rec ele prod-structure-sort-key.

In the previous example, EFF-FROM was the element modified; but is a subordinate to ITEM-SORT-KEY, which is itself a subordinate to PROD-STRUCTURE-SORT-KEY.

Having created YEAR 2000 records, I created a schema for them. The cloning ability that exists for records does not appear to exist for schemas, so I punched the "live" schema, and modified the appropriate record statements to reflect the version 2000 records. I validated the schema, and created and generated a "global" subschema.

If, at this point, you are concerned about free space in your data areas, this would be an opportune place to enlarge areas before the restructure.

With both "before" and "after" schemas available, I ran IDMSRSTC to create the IDMSRSTT macro statements. I then modified the IDMSRSTT macro statements by adding the NUPROCS=parameter to reflect the use of a user procedure to handle date manipulation at restructure time (this manipulation is discussed later).

IDMSRSTT RECNAME=PROD-STRUCTURE,NUPROCS=CSRRENTR

IDMSRSTC correctly calculated changes to the MINLEN parameters due to changed control field lengths and record lengths.

I wrote a user procedure to manipulate the date at restructure time. This program was written in COBOL II, and was link-edited with a modified IGZEOPT (RTEREUS=YES) to reduce runtime. Since the date fields at Liebert are typically alphanumeric, the restructure by default would leave the existing date in the first six bytes of the expanded date field, and initialize the last two bytes to spaces. Since dates are stored in our database as YYMMDD, the "new" two bytes should be at the beginning of the date, NOT at the end. Additionally, I want to seed those two bytes with the proper value: spaces, if the original date was spaces; zeroes, if the original date was zeroes; and the correct century otherwise (I used 40 as the cutoff year). Also, by obtaining information from dictionary records, I was able to create a "generic" database record, sized as big as the largest record to be restructured, and was able to write source code to manipulate date fields as substrings of the "generic" database record, rather than by their symbolic element name. This allowed me to NOT copy all 200+ record definitions into my user procedure (I rely on the record-id provided at runtime to determine which substrings to use).

Before restructuring, the USER-OWNED indexes whose control fields were about to be changed by the restructure were UNLOADED AND DELETED. This was done by writing in-house programs to capture set information and disconnect members from the sets (requiring alternate schemas). Alternatively, a user-written program can link to the index unload utility (as documented in the UTILITIES manual) to accomplish the same thing. After the restructure, the USER-OWNED indexes were rebuilt by another in-house program to reconnect the member records to the appropriate owner (if the pre-restructure work had been accomplished with the index utility, the indexes could have been rebuilt with the standard rebuild starting with the first sort), SYSTEM-OWNED indexes whose control field(s) were just changed by the restructure were rebuilt (using the FROM MEMBERS clause), and DATABASE AREAS whose CALC records had their CALC field(s) just changed by the restructure were UNLOADED and RELOADED. If, as a result of the restructure, a NON-INDEXED SORTED set had both 19- and 20-century dates, that set must be MANUALLY rebuilt, as an UNLOAD will NOT alter a record's position with a set. After restructuring, I updated my "base" schema and subschemas to reflect the restructured status.

At this point in the testing phase, the "clone" database was turned over to the application developers for modifications to software entities. At this point in the final implementation, the dictionary and PDS files can be copied to the production counterparts, and the implementation is complete.

Concerning the methodology devised by our application developers in regards to YEAR 2000 support, all assembler and COBOL date routine programs were re-written with NEW names. ADS/O and DC-COBOL modifications will be assigned and handled by application; JCL, PROC, parm, CULPRIT and batch COBOL modifications will be assigned and handled by PROC (i.e., one person will be responsible for modifying everything that pertains to or is used by that proc). In cases where a program, CULPRIT or parm is used by multiple procs, the first person to get to it modifies it.

Where is Liebert Corp. within the conversion process now? Liebert has implemented in production a small stand-alone pilot system within our CA-CAS manufacturing database. The remainder of our CA-CAS manufacturing database has been restructured and rebuilt, and is undergoing application development modifications. The remaining application databases will be handled as standard project migrations to be completed in 1996.

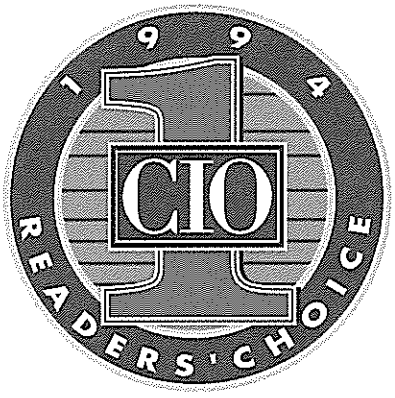

Who Says You Can't Win 'Em All?

CA-Unicenter® just did.

Again. CA-Unicenter is the first (and only) product to ever win the CIO's Readers' Choice Award three years in a row.

In category after category, CIOs chose CA-Unicenter over all the competition. They overwhelmingly voted CA as their preferred choice for Security. They said CA-Unicenter was easier to use. They said it was backed with the best service and support. Offered the best features. Came with the best documentation. And was supported by the Preferred Business Partner, Computer Associates.

No wonder over 50% of the CIOs surveyed said they plan to purchase CA-Unicenter during the next 12 months.

For Information On CA-Unicenter, Call 1-800-225-5224 Or Visit Us At http://www.cai.com

Call right now for the information that will win you over, too.

COMPUTER® ASSOCIATES
Software superior by design.

CA-Unicenter®

©1996 Computer Associates International, Inc., Islandia, NY 11788-7000. All other product names referenced herein are trademarks of their respective companies.

CULPRIT User's Guide: EUDTAC Continues to Share Documents, Prioritize Enhancement Requests

by Carol Silberstein-Locker, EUDTAC/IUAXX/Workshop '96 Chair

The End User and Development Tools Advisory Committee (EUDTAC) discusses and supports the CA products that enable end users to access their data (i.e., CA-IDMS/Server and CA-Visual Express) as well as PC based development tools (i.e., CA-IDMS/PC). EUDTAC represents business perspectives as well as user interests and expertise, promoting and uncovering topics related to End User and Development tools. The following outlines EUDTAC's ongoing agenda.

If there are features you would like to add, change or delete for your CA products, send an enhancement request to IUA Headquarters or submit a DAR to CA.

We Keep Documenting

In response to the overwhelming success of the *OLQ Menu Mode User's Guide*, EUDTAC is preparing a *CULPRIT User's Guide* for those developers and users who use CA-CULPRIT to obtain data. Look for more information and an early release of this document at IMC96 at CA-World '96. The *OLQ Menu Mode User's Guide* is available at IUA Headquarters. Look ahead! What would you like next? Tell us which documents would benefit IUA members and CA users. Do you have an idea? Outline? Rough Draft? Finished document that you would like to share? All contributions are greatly appreciated and may be submitted to IUA Headquarters.

Finally, in order to inform CA of what improvements the end users require, EUDTAC is currently preparing a ballot to prioritize and vote on the outstanding enhancement requests and DARs. If there are features you would like to add, change or delete for your CA products, send an enhancement request to IUA Headquarters or submit a DAR to CA. For more information, contact IUA Headquarters at 312/321-6827 or via the Internet at IUA@sba.com

The IUA "Tips of the Trade" series continues...

TAC Tip of the Trade

by Chris Hoelscher, TAC Correspondent

Did you ever want to see what is ***REALLY*** in your SMP/E inventory? Did you ever want a ***QUICK*** list of APARS, test fixes and USERMODS?

Execute GIMSMP with "standard" JCL, with the following changes:

```
//SMPLIST        DD       DSN=your-dsn, UNIT=SYSDA, DISP=(NEW,PASS),
// DCB=(RECFM=FB,LRECL=121,BLKSIZE=23474),
// SPACE=(23474, (10,10),RLSE)
//SMPCNTL        DD *
 SET BDY (IDMSTGT).
 LIST SYSMODS.
```

Execute CULPRIT, with "standard" JCL, with the following changes:

```
//SYS010DD       DSN=&&WORK01,DISP=(OLD,DELETE)
//SYS030DD
SYSOUT=(A,,220S),DCB=(RECFM=FBA,LRECL=081,BLKSIZE=081)
//SYSIN  DD *
 IN 121 PS          DD=SYS010
 REC I-NAME              002       007
 REC I-LIT                         012       004
 REC I-TYPE                        030       008
 01OUT 080       D NS (XXXX)
 01SORT NOSORT
 013your CORPORATION SYSMOD STATUS
 01410001  'SYSMOD'
 01420001  '  NAME  '
 01430001+'________'
 01510001  I-NAME                  SZ=007
 01410011  'SYSMOD'
 01420011  '  TYPE  '
 01430011+'________'
 01510011 I-TYPE           SZ=008
 017      IF I-LIT NE 'TYPE' DROP
 017      IF I-TYPE NE ('FUNCTION' 'USERMOD' 'APAR         ') DROP
```

1995 Annual Report

by Karen Langley, Secretary/Treasury/Finance/VUTAC Chair

In FY1994-95, IUA continued to offer practical and cost-effective services to its members. Going into the fiscal year, we knew income would be down due to the timing of CA-World. However, the volunteers and staff worked very hard to keep costs down without sacrificing products and services for the members. To help offset the loss in conference income, a change in investment strategy led the IUA to have a considerable increase in investment income in FY 1994-95.

In FY 1993-94, total revenue was $321,451and expenses totaled $305,244; excess revenue over expenses totaled $16,207. In comparison, in FY 1994-95, total revenue was $272,663 and expenses totaled $253,785; excess revenue over expenses totaled $18,878.

Conference Income

The IUA continues to provide consulting services to Computer Associates (CA) for Information Management Conference (IMC) and Business Application Conference (BAC). The disparity in conference income presented in this report is due to the fact that IMC94 and BAC94 were held in different fiscal years. (IUA's fiscal year runs July 1 - June 30). In FY1993-94, IMC94 was held in Atlanta in February 1994. However, BAC94 was held in August 1994 in New Orleans, which occurred in FY1994-95. The recent CA-World '95 in August 1995, also in New Orleans, included both IMC95 and BAC95, and will be recorded in the FY1995-96 statement.

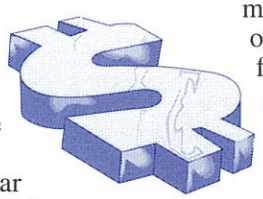

IUA Workshop

The annual IUA Practitioner's Workshop continues to grow in both attendance and offerings as reflected in both the revenue and expense numbers. The registration fee for the Workshop did not increase even though the program was expanded. The Product Direction Forums are now presented during the Workshop to give attendees the latest information direct from CA personnel. The Workshop provides members an opportunity to meet with other users, hear what other companies are doing and a chance to network with industry peers. This offering to the members has been very well received.

Membership Dues

In FY1993-94, IUA offered its members a 20th Anniversary special membership, which consisted of 18 months of membership for the 12-month price. Per acceptable accounting practices, the income generated from this greater than 1-year membership must be divided across the years. Thus the income was reduced for FY1994-95. (In FY1995-96 the IUA Executive Board voted to provide complimentary membership to all members.)

The IUA financial picture is very stable and continues to prosper.

In Sum...

The IUA financial picture is very stable and continues to prosper. With the diligent efforts of the volunteers and the staff at Smith, Bucklin & Associates, the IUA continues to effectively and efficiently provide the services and products members want and need.

Mark Your Calendars!

The Ohio Valley User Group (OVUG) Meeting will be held June 10-11, 1996 in the Cincinnati area. Please call Rob Kantor, Mead Corporation, at 513/495-3263 for more information.

• • • • • • • • • • • • • • • • •

We Want to Hear from You!

Quality Survey Reminder

Please return the enclosed Quality Survey to IUA Headquarters. The purpose of the survey is to determine the user perception of the overall effectiveness of the Quality Process for CA products and solicit changes to the Quality Process in order to achieve the highest quality.

The World's Leading High-Tech Companies Run On CA-OpenIngres.

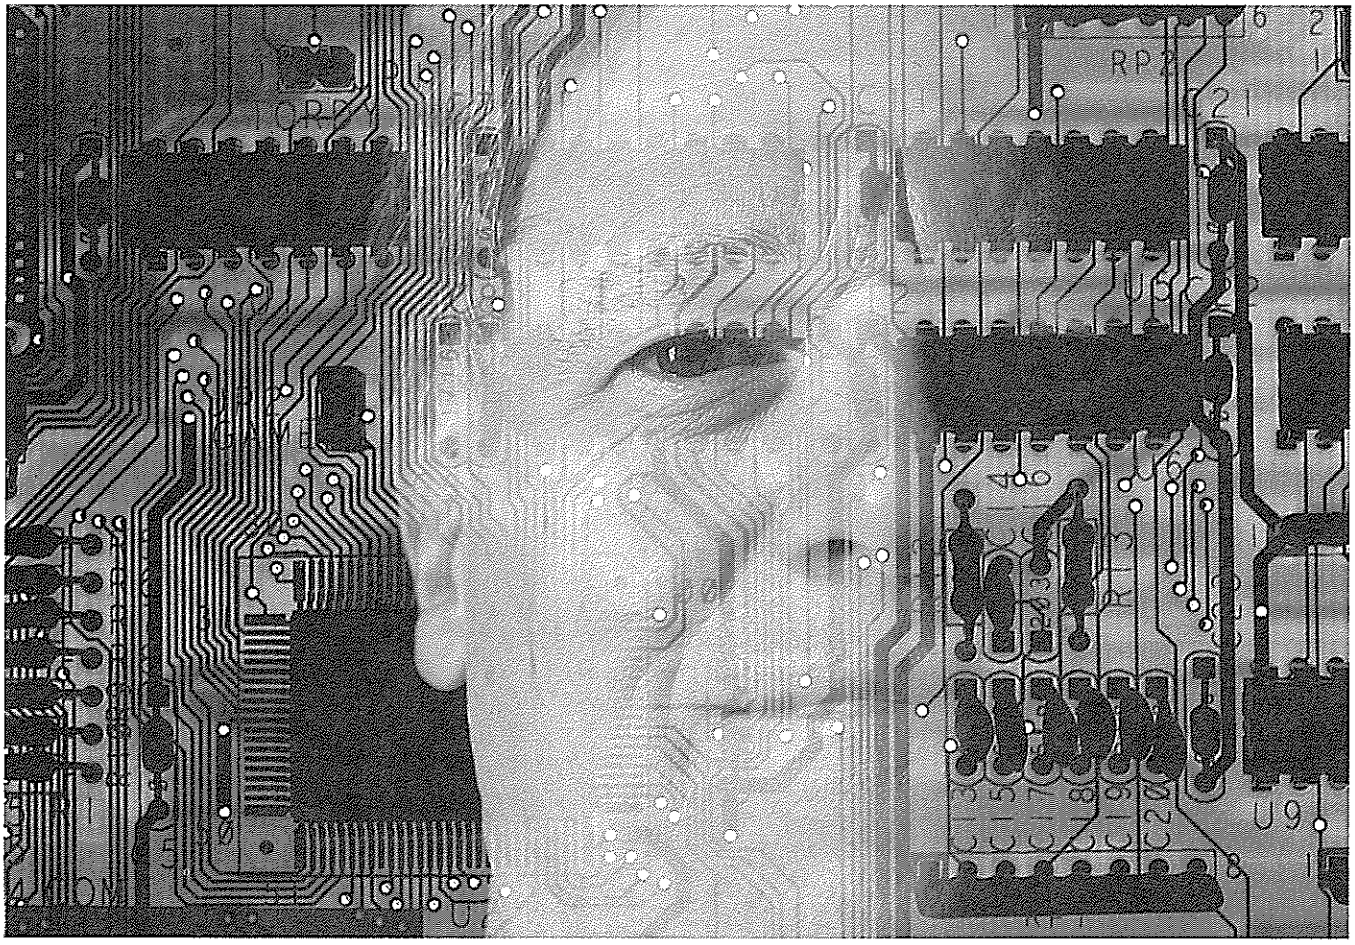

Harris Semiconductor Can Tell You Why.

Harris chose CA-OpenIngres because it's:

- **Self-manageable**
- **Mission-critical**
- **Scaleable**
- **Business process-driven**
- **Accessible anywhere, anytime**
- **Internet Commerce Enabled**

Faced with the daunting task of integrating semiconductor manufacturing plants located all over the world, Harris Semiconductor searched far and wide for the best mission-critical DBMS.

Then they put all their chips on CA-OpenIngres.®

Harris used CA-OpenIngres and CA-OpenROAD™ to implement a proactive, distributed data warehouse that many analysts believe to be the most sophisticated, state-of-the-art data warehousing system in the world. They actually created a data warehouse that spans many databases located at all of their manufacturing plants around the world. Whether the data is in Singapore, Ohio or Malaysia, it can be accessed in real-time, seamlessly and transparently.

As a result, Harris has boosted productivity while saving money and time. They're beating the competition because they found the technology that would "get them there first."

For More Information On CA-OpenIngres, Call 1-800-225-5224.

Or visit us at http://www.cai.com

COMPUTER® ASSOCIATES
Software superior by design.

Open Your Mind To OpenIngres

©1996 Computer Associates International, Inc., Islandia, NY 11788-7000. All other product names referenced herein are trademarks of their respective companies.

You Got Your SQL in My DML!

by Walter Remski, Williams International

Blast the pencils from the pockets of your horn-rimmed, Coke bottle-spectacled programming geeks. Programmatic SQL is a genie of sorts that can grant some potent wishes. At our shop he just granted access to records in entirely different schemas without as much as a subschema for us to maintain! This can be your opinion of SQL, if you respect the limits and promote the abilities of the CA-IDMS SQL option. Fortunately for the DML literate, the code shoehorns in rather gracefully.

In particular, the SQL data manipulation statements have proven fairly congruous with the familiar DML flavor of COBOL and ADS. In order to seamlessly introduce the SQL with DML you need to promote Access module awareness, Relational thinking and Transaction Management.

For those new to SQL, it is not difficult. (I wrote my first only hours after almost putting shaving cream on my toothbrush!) In the bare COBOL case you need no protocol, schema section or even a bound run unit to issue selects. The addition of a special working storage declare section to define host variables, coupled with a select statement will return data. In combination with DML you will find that a finish will close cursors, a bind run unit and obtains will not interfere with SQL statements. The primary pitfall I found was in opening cursors that reference host variables. If you don't fill the variables first, you won't get any result rows. But I'm jumping ahead.

First, create an SQL schema for your nonSQL schema. Do a DCMT display dbtable and insure the dbname that houses the segment you are "SQLing" against also has a syssql segment included in it. Next, issue the create SQL schema in OCF using the dbname parameter to point back at the target dbname. Your first wish has been granted, and you can start practicing SQL in OCF. Beware that all record names with hyphens must be in double quotes and all elements will have their hyphens translated into underscores, e.g., select part_number_001 from cassql."part-master."

Access modules are like subschemas. In our shop, creation of subschemas is through an approved request. People are more than happy knowing that by precompiling their programs with the same IDMSDMLx pre-compiler, they are no longer waiting for subschemas. The downside is that an access module must be created from the relational control module statements in the dictionary. (Create access module SQLSCHEMA. PROGRAM from DICTNAME. PROGRAM ... assuming you are using the default RCM=Program name.) This is also a good time to place the dictname into the sysidms statement of any batch jobs so they find it at runtime.

This brings us to the programming. First, dig up grandma's old COBOL shell, then declare a host variable section. Next, somewhere in the beginning of your process code place two more SQL statements, a set access module (same as program name) and a whenever sqlerror go to (some error paragraph). Then take that SQL statement that you practiced with and place it in the code. Use it in a cursor with fetches if you are working with more than one row of data. Lastly, close your cursor(s) before exiting the program. Test the code until it behaves before adding any garnish.

The simplification of process code is a real time-saver, but you need to instill a mindset foreign to IDMS. Relational joins based on set names

Continued on page 10

Is this copy of *IUA Connections* properly addressed? If not, please fill out this form and mail it to:

Information User Association
401 North Michigan Avenue
Chicago , Il 60611-4267

Current Contact Name: _____

Current Company: _____

New Address: _____

New Contact Name: _____

Company: _____

Address: _____

City, State, ZIP: _____

Phone: _____ Fax: _____

E-mail address: _____

You Got Your SQL in My DML!

Continued from page 9

and redundant data are most useful in providing "logical records" of only the required fields (no LRF infringement intended). The schema qualifier before the table name allows reference to multiple databases just by using a different prefix. (Your access module can be created in either schema.) A rule of thumb is to code your SQL to make the navigation through the database and then fetch the various elements from the simple tabular result. A first thought in the DML literate is that they need a select for each record they are using. NO! One final caution is the access plan optimizer can sometimes get lost trying to parse your statements through BOM type structures. If you run into errors constructing SQL statements and you already tried two alias' for a table, chances are you need to consult page 10-1 of the SQL reference manual.

The most important bit of advice I can give once you start mixing condensed SQL with the verbose DML is to heed the existing transaction bounding and expect to handle the two very different forms of error checking independently. The standard IDMS-STATUS has its SQL counterpart in the WHENEVER SQLERROR statement. Have a thorough understanding of the transactions underway in existing DML programs before bounding commits and rollbacks into the SQL portions or you will pay the price. Our shop has various styles of autostatus and homegrown retry logic. A safe bet is to include a new SQL error trap that deals with the specific SQL errors then route the logic through the same single exit with a message or what-not as you would in the case of a DML error. And you thought modular programming would never pay off!

CASAC/ADAC Humbert Piscitelli

BAC96 Prepares for CA-World '96; CASAC/ADAC Survey to Update Database

IUA Workshop '96 allows us to meet with CA, vendors and our peers to discuss the current and future direction of the CA-CAS products. It is imperative we participate in what will be a worthwhile gathering of individuals who share common interests. The schedule of events provides for Advisory Committee and Business Issues Forum on Sunday, April 28; PDF, Vendor Presentations, Educational Sessions and SIGs on Monday, April 29; and Educational Sessions, SIGs and Vendor Presentations on Tuesday, April 30.

Many of you have already received the CA-World '96 BAC mailing; if you have not, please contact CA to receive the necessary information. I can assure you that it will be an outstanding conference. Register early — BAC96 promises to be the best conference to date! Arrangements have been made to hold our Advisory Committee meeting on Sunday, August 25 from 10:00 a.m. - 2:00 p.m. More specific information will follow in future mailings. We will also be setting a time aside daily to provide a place for individuals to meet and discuss any CA-CAS topics. The meetings will be facilitated by A. Scott Rosales or myself; an agenda will be provided in the final agenda at registration.

It has become apparent that our CA-CAS client database needs to be updated; it presently suffers from acute data integrity. Please help remedy the illness — the prescribed medication is a survey included in this issue of *IUA Connections*. We have structured the survey so it will be painless to complete and will require only a brief moment of your valuable time. The "cured" database will be sent to all clients that complete the survey. Please be reminded of the importance of a reliable database to our daily well-being and participate in this much needed update.

Don't forget to fill out and return the enclosed ADAC/CASAC Database Update survey included in this issue of* IUA Connections. *Your feedback is greatly needed and appreciated!

IUA Quality Corner

Reprinted from IDMS/SQL News 5.1, February 1996

XA or not XA Storage POOL?

Release 12.0 is not only designed to take advantage of XA storage, but in most cases it is necessary to have XA storage pool defined to run a proper production environment under 12.0

However, for those clients running without XA storage pool (hope no one does that in 12.0), pool 0 is under heavy pressure. Typically this happens when you don't have an XA storage pool (pool 255) defined on the system statement. Some users have come across this situation when they want to link to VS COBOL in 24-bit from ADSO. Task location is ANY to load dialogs into XA reentrant pool. But there is no XA storage pool because 31 bit addressing cannot be allowed in such a situation.

One can get only 7.5 M maximum in a practical situation for Pool 0. Various clients who tried 12.0 production without defining an XA STORAGE POOL (255) had serious problems. Typically runtime messages like *DB347018 get storage failed.*

Some relief can be obtained by defining an XA STORAGE POOL (255) for the system usage. User storage allocation will still be done from Pool 0.

Storage protect OFF at system level can also release storage below the line. Using the relocatable storage feature will also release pool 0 at the expense of scratch overhead.

Of course, the ideal way is to take everything above the line by defining an XA STORGAE POOL 128. VS COBOL can address 31-bit if they are linked AMODE 31 RMODE ANY. Also optional PTF CS82390 allows one to load VS COBOL programs in XA pools.

JOURNAL Transaction Level

If the JOURNAL TRANSACTION LEVEL is set too low at 3, journal bottlenecks can occur. Either make it very high (20+) or turn it off (making it 0). To turn it off, sysgen the Journal Transaction Level to 0. Then you must do a DCMT VARY JOURNAL TRANSACTION 0 to update to journal header. A simple sysgen of the JTL to 0 doesn't update the journal. A second way is to sysgen the JTL to 0, shutdown cv, format the journals, and startup again.

How it works?

If the current number of actively journaling transactions is equal to or greater than the system JOURNAL TRANS LEVEL writing of the journal buffer is deferred until it is full. This also means that the COMMIT or FINISH is also deferred...the journal buffer is ALWAYS flushed before database buffers to insure database integrity.

If the CV crashes at this point, the journal buffer will not have been written to the journal, the database buffers will not have been written to the database...the run unit will be seen as active by WARMSTART and will be rolled out.

Storage Protect at System Level

There have been recommendations made for heavy users in 10.2 and in 12.0 (but not in 12.01) to have PROTECT ON at system level and OFF at program level. This way storage SCEs are chained off task TCE and search is faster. With protect off, storage SCEs were chained off CSA giving rise to large chains and 'uneconomical searches.' But this is true only in 10.2 and 12.0., ***not in 12.01***.

From 12.01 onwards, storage allocation, chaining and search have changed and the search is economical now even with storage protect OFF at system level. In short, there is no reason from 12.01 onwards to have storage protect ON at system level and OFF at program level. (Of course, you need PROTECT ON if you really want to use it on certain programs.)

Protect ON at system level will also result in more storage usage:
PROTECT ON at system level - 4K chunks are allocated
PROTECT OFF at system level - 128 bytes pieces are allocated
This is true in 10.2 and 12.0 - i.e., no change.
So it is reasonable to see that STORAGE PROTECT OFF will give better utilization of storage. This makes a difference when the 24-bit pool 0 is at a premium.

Editor's Note: See PIB GI 76847; CA has recently stated that PROTECT ON at the system level continues to provide improved performance.

Suddenly, ICE Is The Hottest Thing On The Internet.

Everybody's been talking about doing business on the Net. And now, you can. Thanks to a new product from Computer Associates called CA-Unicenter®/ICE™. ICE stands for "Internet Commerce Enabled.™" That's a fancy way of saying, this software is designed specifically to handle the unique requirements and challenges of the Internet.

CA-Unicenter/ICE is a distributed management software solution that secures and manages Windows NT and UNIX servers on the Internet and other TCP/IP networks. It's the first product that addresses all the management needs of Web servers and Web clients by providing security, event management, help-desk, storage management, resource accounting and database monitoring. So now you can have a secure, reliable and manageable Internet infrastructure.

Call 1-800-225-5224 for more information.

Visit us at http://www.cai.com and find out how a little ICE can lead to a lot of new business.

COMPUTER® ASSOCIATES
Software superior by design.

Introducing New CA-Unicenter®/ICE™

© 1996 Computer Associates International, Inc., Islandia, NY 11788-7000. All products referenced herein are trademarks of their respective companies.

The IUA "Tips of the Trade" series continues...

Monitoring Run Unit Progress

by Dick Brinson, DBSoft

An SQL query is doing a join on two tables, but seems to be taking a lot longer than expected. With your knowledge of the database structures involved, you might think there is an efficient access strategy. Maybe you left one of the join criteria off of your WHERE clause, and CA-IDMS is doing a full cross-product (for each row in one table, sweep the area containing the other table looking for a match). Maybe the optimizer did not identify the efficient path that you identified. Maybe there is a lot more data than originally thought. Or maybe the system is just really slow that day. It sure would be nice to see what CA-IDMS is really doing.

You have a batch program (running against a Central Version) that you think should have finished by now. There is no JES output that you can examine to determine the program's progress. It is not using excessive CPU time, and it is still doing I/Os. Maybe it is in a loop and you should cancel it. Maybe it is almost done and you should let it run. It sure would be nice to see how far along the program is.

One useful tool to determine what is going on is the W-DB screen of the OPER task. It shows the verb number last executed for each run unit. (Appendix D of the DML Reference-COBOL lists the meaning of each number.) More information would be helpful, though.

In both scenarios, a look into the currency tables would be useful. In the first scenario, if we observe the DBKey of one record remaining constant and the DBKey of the other record is increasing rather quickly (not jumping around), we know that we are sweeping for the second record for each occurrence of the first. In the second scenario, if we knew the DBKeys of the records that were current, we might be able to identify how much of the input file had been processed or how much of an area had been swept.

The following program examines currency blocks for all active run units in a CV and shows the DBKey that is current for each record type. After assembling it, link it with IDMSBALI. By executing it repeatedly, you can monitor the progress of a run unit. Instead of typing in the program, you can download it from the new, upcoming IUA Web Page.

```
CURR      START
          #MOPT CSECT=CURRCS,ENV=USER,AMODE=ANY,RMODE=ANY
CURREP    EQU   *
          USING CSA,R10
          ENTRY CURREP
          USING CURREP,R12
          LR    R12,R15
          #GETSTG TYPE=(USER,SHORT),PLIST=*,LEN=WORKSTGL,ADDR=(R2),      X
                COND=NO,INIT=BLANK
          USING WORKSTG,R2
          #LINEOUT OUTLEN=80,LRB=PARMS,OUTAREA=HEADING1,COND=(ALL),      X
                HDR=1,OPTNS=(NEWPAGE,NOWAIT)
          #LINEOUT OUTLEN=80,LRB=PARMS,OUTAREA=HEADING2,COND=(ALL),      X
                HDR=2,OPTNS=(NOWAIT)
          L     R3,CSACCEA                    GET ADDR OF CCE
          USING CCE,R3
          L     R8,CCEFTBKA                   GET FIRST TBK
          USING TBK,R8
VIBLOOP   L     R9,TBKFVIB                    GET VIB
          USING VIB,R9
          CLC   VIBTSKID,=F'0'                IF TASK ID ZERO,
          BE    NEXTVIB                             GO GET NEXT VIB
          MVC   OUTLINE,=CL80' '
          MVC   OUTPNAM,PRPROGM               GET PROGRAM NAME
          L     R4,VIBRUNID                   GET RUNUNIT ID
          CVD   R4,DBLWORD
          MVC   WRKFLD,MASKP
          ED    WRKFLD,DBLWORD+2
          MVC   OUTRUID,WRKFLD+5
          L     R4,VIBTSKID                   GET TASK ID
```

Continued on page 14

Tips of the Trade

Continued from page 13

```
                CVD     R4,DBLWORD
                MVC     WRKFLD,MASKP
                ED      WRKFLD,DBLWORD+2
                MVC     OUTTSKID,WRKFLD+5
                L       R7,VIBFIBA              GET ADDR OF FIB
                USING   FIB,R7
                L       R6,FIBSRN               GET ADDR OF SUBSCHEMA RECORD
                USING   FSR,R6
FSRLOOP         CR      R6,R7                   IF LAST SUBSCHEMA RECORD,
                BE      NEXTVIB                         GO GET NEXT VIB
                MVC     OUTRNAM,FSRNAME         GET RECORD NAME
                L       R5,VIBVSEA              GET BASE OF VSR
                A       R5,FSRDYNPG             ADD DISPLACEMENT TO VSA
                USING   VSA,R5
                L       R4,VSALOPG              GET LOW PAGE
                CVD     R4,DBLWORD
                MVC     WRKFLD,MASKP
                ED      WRKFLD,DBLWORD+2
                MVC     OUTLOPG,WRKFLD+2
                L       R4,VSAHIPG              GET HIGH PAGE
                CVD     R4,DBLWORD
                MVC     WRKFLD,MASKP
                ED      WRKFLD,DBLWORD+2
                MVC     OUTHIPG,WRKFLD+2
                DROP    R5
                L       R5,VIBVSEA              GET BASE OF VSR
                A       R5,FSRDYN               ADD DISPLACEMENT TO VSR
                USING   VSR,R5
                L       R4,VSRCDBK              GET CURR DBKEY OF REC TYPE
                DROP    R5
                L       R1,VIBKYFMT
                C       R4,=F'-1'
                BE      NEXTFSR
                LA      R5,0                    CONVERT DBKEY TO DISPLAY
                SRDL    R4,0(R1)
                CVD     R4,DBLWORD
                MVC     WRKFLD,MASKP
                ED      WRKFLD,DBLWORD+2
                MVC     OUTCRPG,WRKFLD+2
                LA      R4,0
                SLDL    R4,0(R1)
                CVD     R4,DBLWORD
                UNPK    OUTCRLN,DBLWORD+5(3)
                OI      OUTCRLN+4,X'F0'
                MVI     OUTCRLN,C' : '
SHOWIT          #LINEOUT OUTLEN=80,LRB=PARMS,OUTAREA=OUTLINE,COND=(ALL),      X
                        OPTNS=(NOWAIT)
NEXTFSR         L       R6,FSRIBN               GET NEXT SR

                B       FSRLOOP
NEXTVIB         L       R8,TBKNCCE              GET NEXT TBK
                CR      R8,R3                   IF NEXT IS THE CCE,
                BE      ALLDONE                         GO TO ALLDONE
                B       VIBLOOP
ALLDONE         MVC     OUTLINE,=CL80' '
                #LINEOUT OUTLEN=80,LRB=PARMS,OUTAREA=OUTLINE,COND=(ALL),      X
                        OPTNS=(NOWAIT)
                XR      R15,R15
                #RETURN
BLANK           EQU     X'40'
MASKP           DC      X'402020202020202020202120'
HEADING1        DC      CL25'TASK ID     RUN ID PROGRAM      '
                DC      CL18'    RECORD NAME    '
                DC      CL20'  LO PAGE   HI PAGE    '
                DC      CL17'  CURRENT DBKEY   '
HEADING2        DC      CL25'------- ------- ------- '
                DC      CL18'----------------'
                DC      CL20'-------- --------'
                DC      CL17'----------------'
                LTORG
WORKSTG         DSECT
OUTLINE         DS      0CL80
                ORG     OUTLINE
OUTTSKID        DS      CL7
                DS      CL1
```

Continued on page 16

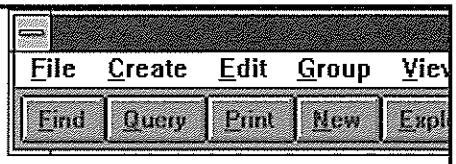

Hitch A Ride To Second-Generation Client/Server On CA-OpenROAD.™

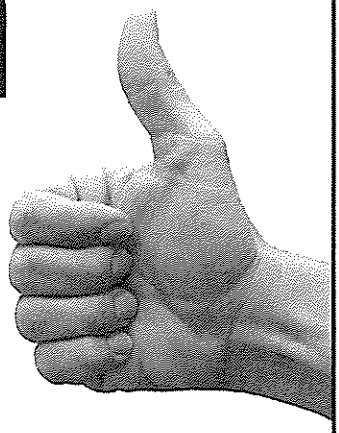

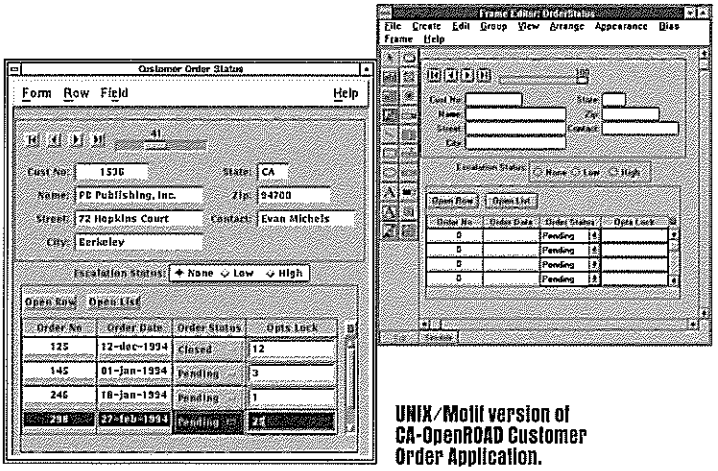

UNIX/Motif version of CA-OpenROAD Customer Order Application.

Paint your applications with CA-OpenROAD's visual development tools.

The path to second-generation client/server has finally been cleared. Introducing CA-OpenROAD.™

With its advanced, powerful, repository-based architecture, applications can be constructed accurately and with ease. And with its multi-platform GUI support, it doesn't matter whether you're targeting UNIX/Motif or Windows.

Mission-Critical
Native Database Access
Object-Oriented
Fully Scalable
Template-Driven
Repository-Based
Multi-Platform

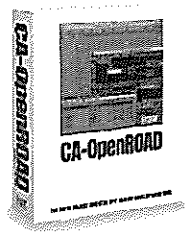

And since CA-OpenROAD is the only product of its kind that supports triggers and stored procedures with identical source code across databases, you don't have to be a database expert to use it.

What's more, CA-OpenROAD provides both template-driven application generation and complete object-orientation including encapsulation, inheritance and polymorphism.

For More Information, Call 1-800-225-5224, Dept. 13831.

So phone today to arrange a demonstration of new CA-OpenROAD. It's a test-drive you'll find truly exhilarating.

New CA-OpenROAD

COMPUTER ASSOCIATES®
Software superior by design.

© 1994 Computer Associates International, Inc., Islandia, NY 11788-7000. All products referenced herein are trademarks of their respective companies.

Tips of the Trade

Continued from page 14

```
OUTRUID     DS    CL7
            DS    CL1
OUTPNAM     DS    CL8
            DS    CL1
OUTRNAM     DS    CL18
OUTLOPG     DS    CL10
OUTHIPG     DS    CL10
OUTCRPG     DS    CL10
OUTCRLN     DS    CL5
            ORG   OUTCRPG
OUTCRKY     DS    CL15
DBLWORD     DS    D
WRKFLD      DS    CL12
PARMS       DS    20F
WORKSTGL    EQU   *-WORKSTG
            COPY  #CSADS
            COPY  #CCEDS
            COPY  #TBKDS
            COPY  #VIBDS
            ORG   VIBWORK
            DS    0F
PRSSNAM     DS    CL8
PRPROGM     DS    CL8
            #FIBDS TYPE=DSECT
            #FSRDS TYPE=DSECT,VARS=YES
            #FSADS TYPE=DSECT,VARS=YES
            END
```


IMC At

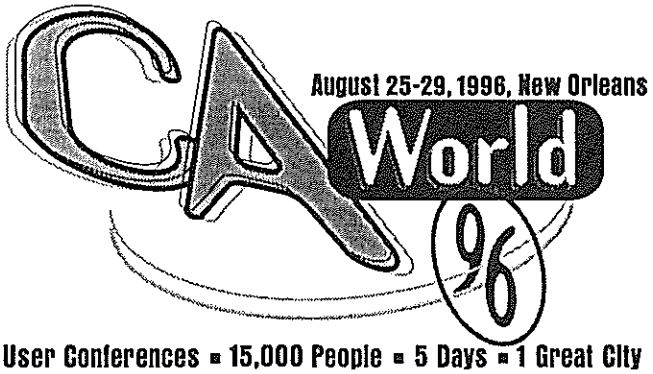

14 User Conferences ▪ 15,000 People ▪ 5 Days ▪ 1 Great City

The Future Of Information Technology

To receive a preliminary agenda or more information, call 1-800-CA INFO96 (224-6369).
Visit us on the Web at http://www.cai.com.

COMPUTER® ASSOCIATES
Software superior by design.

Make your plans now! Be a part of CA-World'96—the most informative and educational event of the year.

CA-World'96 highlights include:

- FREE Formal Education Courses
- IMC and Cross-Conference User Sessions
- The World Resource Center featuring state-of-the-art Test Drive areas and hundreds of exhibitors demonstrating their latest solutions and enterprise-wide technology
- Key CA development and support staff providing new product demonstrations, hands-on training, topical presentations and one-on-one meetings
- CA-World'96 and conference-specific keynotes

This year's **Information Management Conference** has it all. From mainframe to desktop, IMC'96 will provide you with the information necessary to bring your CA-IDMS® Database into the next millennium. It is essential in today's business environment to be able to integrate information technology across all products and all platforms. Come and find out what is going on inside CA-IDMS to help you resolve these important business issues. You will also discover what is brewing in other new technologies. This year's IMC will help you implement successful plans for today and assist you in planning for continued success tomorrow. Don't miss it!